

The Processes of Scientific Research: The Strategy of Experimentation

Deepak S. Kulkarni

Department of Computer Science

Carnegie-Mellon University

Pittsburgh, Pennsylvania 15213

This thesis is submitted to Carnegie-Mellon University in partial fulfillment of the requirements for the degree of Doctor of Philosophy. This research was sponsored by the Defense Advanced Research Projects Agency (DOD), ARPA Order No. 4976 under Contract F33615-87-C-1499 and monitored by: Avionics Laboratory, Air Force Wright Aeronautical Laboratories, Aeronautical Systems Division (AFSC), Wright-Patterson AFB, OH 45433-6543. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or of the U.S. Government.



Abstract

In this research, we produce a program KEKADA capable of carrying out intelligent experimental programs on problems similar to those faced by a number of experimental scientists. KEKADA has a set of experimentation strategies, that were detected from the traces of the behaviors of actual scientists. KEKADA strategies include : focusing on a surprising phenomenon, characterizing the surprising phenomenon by general strategies such as magnification, applying divide-and-conquer, determining the scope of phenomenon, factor-analysis, relating to similar phenomena, and domain-specific strategies and hypotheses. The domain-specific heuristics in KEKADA are efficient and practical instantiations of general strategies such as - controlled experimentation, determination of complexity of a process, testing of a causal chain, componential analysis, differencing and divide-and-conquer. We also analyze the reasons why KEKADA in its present form would not be a good experimental scientist.

Acknowledgements

With Herb Simon as an advisor, with Allen Newell, David Klahr and Tom Mitchell as thesis committee members, and with the intellectual, computational and social environment of CMU, I had the best environment to study the subject of scientific discovery. Although due to my bounded rationality I could learn only a few skills in four years, I would like to take this opportunity to thank all those who made this a great learning experience.

First of all, I would like to thank my advisor Herb Simon for his invaluable guidance and collaboration. This work has benefitted much from his deep insights in cognitive science. Working with him has been an exhilarating experience. I would also like to thank other members of my thesis committee- Allen Newell, David Klahr, and Tom Mitchell. Discussions with them have always been thought-provoking and have exposed many new issues.

I would like to thank Langley, Simon, Bradshaw, and Zytkow, whose previous work provided the basic questions in scientific discovery on which this research focuses. Many thanks are due to Prof. Frederic L. Holmes whose work on Hans Krebs' research has provided some of the basic data used in this research. I would like also like to thank the Royal Institution, London and Thomas Martin for publishing Faraday's diaries. This work builds upon the analysis of Ryan Tweney, who has been a pioneer in taking a cognitive science approach to the analysis of Faraday's diaries.

I would like to thank Raul Valdes-Perez and Peter Jansen for many stimulating discussions. I have also benefitted from comments of the following people on an earlier paper and a draft of this dissertation- John Modell, David Steier, Kurt Van-Lehn, Kenneth Schaffner, Eugene Miya, Yumi Iwasaki, Peter Jansen, Craig Knoblock, and Elaine Kant. Chen Ho, David Hackney, Uday Shenoy and Natarajan Ganesh have helped me with chemistry. I would also like to thank the participants in the Machine Learning seminar for the lively discussions on the topic of discovery.

My CMU years have been made particularly enjoyable by my friends: Milind Tambe, Sunil Issar, Yumi Iwasaki, Wei-Min Shen, Harry Printz, Anthony Stentz, Raul Valdes-Perez, Peter Jansen, Uday Shenoy, Raju Ramanujan, Inderpal Bhandari, Thad Polk, Ravi Mirchandani, Craig Kaplan, Craig Knoblock, Carlo Tomasi, Angelika Zobel, and Yoshimasa Goto. I would like to express my indebtedness to my parents for instilling the importance of education in me.

Table of Contents

1. Introduction	3
1.1. Perspective	3
1.2. The Strategy	4
1.3. Overview of the research	4
1.4. Organization	5
2. Task Analysis	7
2.1. The Strategy	7
2.1.1. The Choice of the problems	8
2.1.2. Methodology	8
2.2. Derivation of the test-problems	9
2.2.1. Hans Krebs' work on urea synthesis.	9
2.2.2. Initial State: Background of the Discovery	9
2.2.3. Final state: Discovery of Ornithine Cycle	10
2.2.4. Transition from Initial state to final state	11
2.2.5. Problem Statement	11
2.2.6. Degree of similarity to actual problem	11
2.2.7. Sources of Historical Data	12
2.2.8. Problems	12
2.2.9. Sensitivity analysis of the problem statements	13
2.3. The Procedure followed	13
3. KEKADA the program	15
3.1. Scenario	15
3.2. Dual Search in KEKADA	16
3.3. Production System	17
3.4. Representation of data and processes	18
3.4.1. Representation of Data	18
3.4.2. Representation of Confidence Measures	19
3.4.3. Representation of processes/heuristics	20
3.5. Goal of KEKADA	20
3.5.1. The goal of the program	20
3.5.2. Focusing on surprise as a search-control strategy	20
3.6. KEKADA control structure	21
3.6.0.1. Interaction of Heuristics	23
3.6.1. Implementation of the 'surprise' heuristic	24
3.6.2. Expectation-setters	24
3.6.3. Problem-generators	26
3.6.4. Problem-Choosers	26
3.6.5. Hypothesis or Strategy Choosers	26
3.6.6. Experimenters	26

3.7. Strategies KEKADA uses to characterize a surprising phenomenon	26
3.7.1. Characterization of a surprising phenomenon	27
3.7.2. Magnification of the phenomenon	27
3.7.3. Divide-and-conquer strategy	28
3.7.4. Determination of the scope of the puzzling phenomenon	28
3.7.5. Factors analysis	29
3.7.6. Relating similar phenomena	29
3.7.7. Gather-data strategy	30
3.7.8. Activation of previously known hypotheses	30
3.7.9. Domain-specific Strategies	31
3.7.9.1. Testing a causal chain	31
3.7.9.2. Controlled Experimentation / Factor Analysis	32
3.7.9.3. Mill's Difference Principle	32
3.7.9.4. Reasoning about structural components	33
3.7.9.5. Determination of the complexity of a process	34
3.7.9.6. Successive Refinement	34
3.7.9.7. Conservatism	35
3.7.9.8. Direct Verification	35
3.7.10. Decision-makers	35
3.8. Summary	36
4. Program Behavior	37
4.1. Behavior in the biochemistry of urea metabolism	37
4.1.1. Initial Working memory of the program	37
4.1.2. Overview of the KEKADA behavior	38
4.1.3. The Ornithine Effect Discovery	38
4.1.4. Determination of Scope	41
4.1.5. Discovery of the Reaction Path	42
4.1.5.1. Discovery of Citrulline as an Intermediate	43
4.1.6. Summary of the results produced	43
4.1.7. Analysis	44
4.1.7.1. Prior hypotheses factor the problem space	44
4.1.7.2. Focus on a surprise factors the problem space	44
4.1.7.3. Effectiveness of gradual refinement strategy	44
4.1.7.4. Heuristics do not always work	45
4.1.8. Trace of KEKADA behavior on urea synthesis	45
4.1.9. Differences in KEKADA behavior and Krebs research	49
4.1.10. Sensitivity analysis of the problem statement	49
4.1.11. Summary	50
4.2. Behavior in the biochemistry of glutamine metabolism	50
4.2.1. Background of the discovery	50
4.2.2. Initial working memory of the program	51
4.2.3. Overview of KEKADA's behavior on amino acid metabolisms	51
4.2.4. Characterization of Ornithine-in-kidney Effect	52
4.2.5. Results produced of the program	57
4.2.6. Analysis	57
4.2.6.1. Focusing on a surprise as a search-control strategy	57
4.2.6.2. Determination of scope	57
4.2.6.3. Gather Data Strategy	58
4.2.7. Domain-specific strategies	58
4.2.8. Trace of KEKADA behavior	58
4.2.9. Sensitivity analysis of the Problem Statement	61
4.2.10. Summary	61

4.3. Behavior in the ether reactions	62
4.3.1. Background of the discovery	62
4.3.2. Initial state	62
4.3.3. Discovery of alcohol structure	62
4.3.4. Trace of program behavior	64
4.3.5. Results produced by KEKADA	65
4.3.6. Sensitivity analysis of the problem statement	65
4.3.7. Summary	66
4.4. Behavior in magneto-electricity	66
4.4.1. Background of the discovery	66
4.4.2. Initial Working Memory	67
4.4.3. Overview of KEKADA's research on magneto-electricity	67
4.4.3.1. Observation of Make-Break Effect	67
4.4.3.2. Characterization of Make-Break Effect	69
4.4.3.3. Observation of induction from Magnetic Motion.	71
4.4.4. Analysis	71
4.4.4.1. Role of focusing on a surprise	71
4.4.4.2. Magnification Strategy	72
4.4.4.3. Factor Analysis	72
4.4.4.4. Relating similar phenomena	72
4.4.5. Trace of the program behavior	72
4.4.6. Results produced by KEKADA	74
4.4.7. Sensitivity analysis of the problem statement	75
4.4.8. Summary	75
4.5. Behavior on KEKADA surprise	75
4.5.1. Introduction	75
4.5.2. Initial state	76
4.5.3. KEKADA's research on cognitive science	76
4.5.4. Trace of program behavior	78
4.5.5. Results produced by KEKADA	79
4.5.6. Summary	80
4.6. Other Cases from History of Science	80
4.6.0.1. Priestley's work on oxygen	80
4.6.0.2. Discovery of X-rays	81
4.6.0.3. Discovery of radioactivity	81
4.6.0.4. Discovery of radioactive chains	81
4.6.0.5. Discovery of Photo-electric Effect	82
5. Analysis	83
5.1. Generality of the heuristics	83
5.2. Applicability of the heuristics in multiple tasks	84
5.3. Empirical evidence of the effectiveness of KEKADA	84
5.4. Analysis of the program performance	85
5.4.1. Role of dual search of KEKADA	85
5.4.2. Focusing on surprise as a search-control strategy	86
5.4.3. Role of KEKADA control structure	87
5.4.4. Strategies used to understand the surprising phenomenon better	87
5.4.5. Magnification Strategy	88
5.4.6. Determination of scope	89
5.4.7. Divide and conquer	90
5.4.8. Factor analysis	90
5.4.9. Gather Data Strategy	91

5.4.10. Related phenomenon	91
5.4.11. Domain-specific strategies as efficient and practical instantiations of General strategies	91
5.4.12. The effect of the intervention of the user	92
5.4.12.1. Urea synthesis problem	92
5.4.13. The role of the missing processes	94
5.4.14. Role of the Architecture	95
5.5. KEKADA as a model of Hans Krebs	95
5.6. The class of problems over which KEKADA is applicable	95
5.7. Summary	95
6. Conclusion	97
6.1. Summary of the Ideas in the Chapters	97
6.2. Central Thesis of this research	97
6.3. Contributions of the work	98
6.4. Relation to other work	98
6.5. Directions for future research	101
6.5.1. Comparison of plan-tweaking with experimentation	101
6.5.2. Problem-choosers	101
6.5.3. Creation of discovery assistants	101
6.5.4. Representation Change	102
6.5.5. Integrated models	102
Appendix A. Glossary	103
Appendix B. KEKADA Heuristics In An Ordered Form	105
B.1. Expectation-Setters	105
B.2. Problem-generators	105
B.3. Problem-Choosers	106
B.4. Hypothesis or Strategy Choosers	106
B.5. Experimenters	106
B.6. Experiment-Proposers	108
B.7. Hypothesis and Confidence Modifiers	109
Appendix C. A detailed trace of KEKADA	111
C.1. The generation of hypotheses and strategies in the face of Ornithine Effect	111
C.2. Pieces of knowledge in the working memory	111
C.3. Changes in the working memory	113
C.4. A firing of a production	115
Appendix D. KEKADA.1 simulation of the Discovery of Urea Cycle	117

List of Figures

Figure 2-1: Ornithine Cycle	10
Figure 3-1: Two-space Model of Learning	17
Figure 3-2: Interaction of heuristics	22
Figure 4-1: Discovery of Ornithine Effect	39
Figure 4-2: Discovery of Ornithine Cycle	42
Figure 4-3: Ornithine Effect in Kidney	52
Figure 4-4: Study of deamination reaction	53
Figure 4-5: Study of glutamic acid effect	55
Figure 4-6: Discovery of Glutamine reaction	56
Figure 4-7: Discovery of Alcohol Structure	63
Figure 4-8: Study of Induction of electricity	68
Figure 4-9: Study of induction of electricity	69
Figure 4-10: Discovery of Induction from motion of a magnet	70
Figure 5-1: Generality of heuristics	84
Figure 5-2: Applicability of heuristics in multiple tasks	85
Figure 5-3: Number of Experiments carried out	93

Chapter 1

Introduction

One thing I have learned in a long life: that all our science, measured against reality, is primitive and childlike - and yet it is the most precious thing we have.
- Albert Einstein. (MacGibbon, 1973)

Lenat and Feigenbaum (1987) argue that the ability to learn by discovery is needed in future AI expert systems for two reasons. After handcoding some initial knowledge, it will be cheaper to let the program learn by discovery. Secondly the program can discover new knowledge that was not known by the program designers. Similarly, Simon in an earlier analysis of machine learning research also came to the conclusion that machine discovery is an important research area (Simon, 1982).

Apart from its central importance for the AI programs, machine discovery also holds promise as a means of producing new scientific knowledge in various disciplines. This potential has already been demonstrated by the Meta-DENDRAL program (Buchanan & Feigenbaum, 1978). Some of Meta-DENDRAL discoveries were published in a chemical journal.

1.1. Perspective

Previous AI work on scientific discovery includes the work of Buchanan and others on Meta-DENDRAL (Buchanan & Feigenbaum, 1978), of Lenat on AM (Davis & Lenat, 1980) and EURISKO (Lenat, 1982), of Friedland (1979) on MOLGEN and of Langley, Simon, Bradshaw, and Zytkow (1987) on the BACON series. The BACON series of programs produced heuristics which are capable of discovering new terms and new laws from data. The question of what processes are capable of producing such data was left largely unanswered. This research aims to advance the state of the art by producing computational strategies of experimentation applicable over a wide variety of domains.

The last few years have also seen a number of historians, philosophers, and sociologists of science advocating the cognitive science approach to the study of scientific research (Gruber 1974, Tweney 1987, Giere 1988). There have been analyses of historical accounts from the cognitive science point of view

(Tweney 1985). Research reported in this thesis has benefited from some of the fine-grained data and analyses produced by the historians (Holmes 1980).

1.2. The Strategy

Scientific research involves a wide variety of processes, and a large number of questions can be asked about the nature of scientific research. One would tend to think that some of these processes (for example, the use of mental imagery) could not be fully automated in the near future. The strategy used here to deal with the complexity of scientific research is to focus on a specific aspect of it and try to answer only a few questions in one study.

This research has focused on identifying some strategies of experimentation that are likely to be useful in empirical scientific research.

1.3. Overview of the research

In this research, we produce a program KEKADA capable of carrying out intelligent experimental programs on problems similar to those faced by a number of experimental scientists. KEKADA has a set of experimentation strategies that were detected from the traces of the behaviors of scientists. KEKADA strategies include : focusing on a surprising phenomenon, characterizing the surprising phenomenon by general strategies such as magnification, applying divide-and-conquer, determining the scope of phenomenon, factor-analysis, relating to similar phenomena, and domain-specific strategies and hypotheses. The domain-specific heuristics in KEKADA are efficient and practical instantiations of general strategies such as - controlled experimentation, determination of complexity of a process, testing of a causal chain, componential analysis, differencing and divide-and-conquer.

We defined five different problems similar to those faced by some experimental scientists. KEKADA was shown capable of producing interesting research results on these problems. However KEKADA lacks many of the processes a human researcher employs in his research, and thus the program would not be a good experimental scientist in its present form.

1.4. Organization

Chapter 2 will describe the methodology used in this research. Chapter 3 describes KEKADA. Chapter 4 describes the behavior of KEKADA in various task-situations. Chapter 5 analyses the behavior of KEKADA. Chapter 6 summarizes the conclusions drawn from the work. There is a glossary at the end of the dissertation to help a reader who is not familiar with the specific subject-domains in which KEKADA operates.

Chapter 2

Task Analysis

If you try to do the task, you may or may not succeed. But if you don't, you certainly won't.
-Anonymous

2.1. The Strategy

If a human researcher makes a discovery worthy of getting a patent, we would be justified in judging him to be capable of taking a position of a research scientist in a particular discipline. Such a judgement would not be appropriate in the case of computer programs. Human researchers possess a wide range of processes and are capable of producing fruitful investigations on a very wide class of problems. It is not within the current state of art in AI to produce a computer program having 100,000 productions a human mind may have. Therefore, a reasonable research strategy would be to create a computer program that possesses certain processes such as those capable of making a discovery of a law from numerical data, but not many others. *Such a computer program will lack many of the processes a human researcher possesses, and thus it will be capable of making discoveries on a class of problems which is significantly smaller than the class of the problems on which a human researcher can make discoveries.* As we make progress in identifying various processes involved in scientific research, we will have programs capable of working on a larger range of problems. As we mentioned earlier, this research aimed at identifying some strategies useful in experimental research and at creating a program possessing them. We needed to choose a set of problems to test the effectiveness of this program.

2.1.1. The Choice of the problems

The ideal choice for such problems would be some of today's open research problems. However testing a computer program directly on such research problems would present a number of difficulties. There is no guarantee that a particular research problem is likely to be solvable within a short period of a few months even by a human researcher. There would be an even smaller chance that a computer program having only a limited number of processes will be capable of making a discovery on an arbitrarily chosen problem. Furthermore, experimentation would be costly and would require help of researchers working on that problem. Lastly, the scientific community relies on the time-consuming process of peer review of evaluating new research results. For these reasons, it is a risky and costly strategy to test a computer program directly on open research problems.

Most engineering disciplines and medicine deal with such a difficulty by creating an artificial system closely resembling the real one and testing the mechanism first on this system. In our own work, we will set KEKADA with some problems closely resembling those faced by some scientists in the past. However when one runs a program on a problem from the history of science, one might know both the solution of the problem and possibly the pathway to the solution before running the program. Methodology we use in evaluating the ability of the computer program has some safeguards against the effects of knowledge of the solution of the problem on the construction of the program.

2.1.2. Methodology

The methodology used in this research is three-fold:

- * Observe the behavior of scientists as evident in diaries and retrospective interviews. This is a primary rich source from where we uncover methods.

- * Express these methods in the form of a running program.

- * Run the program on a number of different tasks.

Consider a heuristic in the program that satisfies two properties First, it is not specific for the particular problem such as urea synthesis (A). It could be a heuristic applicable generally in chemistry or biochemistry or it could be domain-independent heuristic. Secondly, it is used in solving more than one problem (B).

Such a heuristic would not be specific for a particular problem on which the program was run. Thus it is likely to be useful for a class of problems. If we create a program having such heuristics and show

that it is sufficient to solve a number of different problems, we can safely conclude that it will be sufficient to solve a wider class of problems. The assumption in this methodology is that the above two constraints (A and B) are good safeguards against the bias introduced by prior knowledge of the solution of the problem.

2.2. Derivation of the test-problems

In this section we will define a set of problems on which we will run our program. First we will examine the work of Hans Krebs on the urea synthesis and define a problem similar to the one faced by Hans Krebs.

2.2.1. Hans Krebs' work on urea synthesis.

Hans Krebs' discovery of urea synthesis in 1932 was an important discovery. In the next subsection, we will first describe the initial state of knowledge with which Krebs began his research on July 26, 1931. We will then describe the research results he had produced by April 13, 1932. My account is based largely on the historical investigations of Krebs' work by F.L. Holmes (Holmes1980). Then we will describe a problem which is quite similar to the research problem Krebs seemed to be facing.

2.2.2. Initial State: 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. It had then been hypothesized that amino acids might be producing urea with ammonia as an intermediate or amino acids might be combining with ammonia in a mechanism producing urea. Both of these were *incomplete* hypotheses about the mechanism of urea synthesis.

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

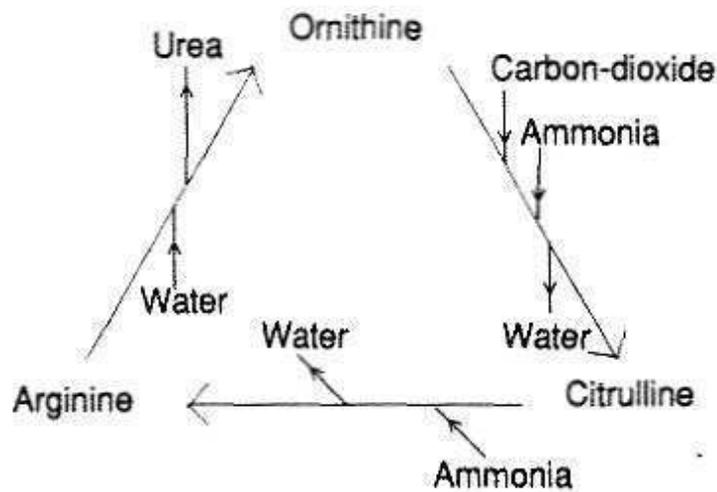


Figure 2-1: Ornithine Cycle

2.2.3. Final state: Discovery of Ornithine Cycle

Krebs studied urea synthesis using the tissue slice method. By April 13, 1932, he had shown that the reaction proceeds by a cyclic mechanism shown in the figure 2-1. One molecule of Ornithine combines with one molecule of ammonia and one molecule of carbonic acid producing citrulline and water. In a second reaction, one molecule of citrulline combines with a molecule of ammonia producing arginine. In a third reaction, arginine hydrolyses to produce ornithine and urea. The ornithine cycle as it was understood and depicted in 1932 is shown in Figure 2-1. Other researchers have since further elaborated the steps in the cycle, and the ornithine cycle as we understand it today is somewhat more complex (See Lehninger 1982).

2.2.4. Transition from initial state to final state

Let us now examine the nature of the transition of the state of knowledge of Hans Krebs. The initial state of knowledge of Krebs was mainly a set of reactions and a set of hypotheses about the mechanism of urea synthesis. Thus it included such hypotheses as that amino acids produce urea by combining with ammonia or by producing ammonia as an intermediate. Krebs' final state was again a set of reactions, now including the reactions involved in the ornithine cycle. Thus by April, 1932, Krebs had more detailed knowledge about the mechanism of urea synthesis than he had in November 1931. Thus his scientific research produced progress from one knowledge state to a better knowledge state.

2.2.5. Problem Statement

Now we will define a problem P1, which is similar to the problem faced by Hans Krebs.

Problem 1 (P1):

Given:

A surprising phenomenon: that alanine produces very little urea on liver tissue slices.

Two previously postulated hypotheses about urea synthesis: amino acids might somehow combine with ammonia to produce urea, or amino acids may deaminate producing ammonia which may further convert to urea.

Some Basic Facts: Structural formulae of various substances, set of values of pH, ordered-lists of amino-acids, amines, and carboxylic acids, stimulators and inhibitors associated with processes. A set of previously known reaction including the arginine reaction.

Produce: some of the reactions involved in the urea synthesis.

2.2.6. Degree of similarity to actual problem

Problem P1 is not exactly the same as the problem faced by Hans Krebs. P1 has been specified at a certain level of abstraction. Krebs had to carry out the experiment by setting up the apparatus. Planning an experiment given the top-level goal is a research problem on which some AI work has been done. (For example, Friedland, 1979) Thus problem P1 is aimed at capturing the difficulty Krebs faced in coming up with experiments such as "Carry out an experiment on Ornithine and Ammonia in liver tissue slices under certain experimental conditions." However post-hoc knowledge has affected our definition of the problem statement. While it is quite clear that the problem P1 has strong resemblance with the problems Krebs faced, it is beyond the scope of this thesis to make an evaluation of how simplified this problem is with respect to Krebs' actual problem.

2.2.7. Sources of Historical Data

We will use a number of different historical sources for defining problems similar to the problems some scientists faced and also as sources of uncovering new mechanisms.

Biochemistry of urea synthesis: Hans Krebs' discovery, in 1932, of the urea cycle was a major event in biochemistry. Holmes' reconstruction of this discovery from published papers, laboratory notebooks, and diaries constitutes a magnificent body of data (Holmes, 1980).

Amino-acid biochemistry: After his discovery of the urea cycle, Hans Krebs continued to work on amino acid metabolisms. Holmes has now provided us with data on this later research pathway which led to the discovery of glutamine synthesis.

Magneto-electricity : In 1831, Faraday carried out research that resulted in the discovery of induction of electricity from the motion of a magnet. The Royal Institution has published Faraday's diaries, which describe how Faraday carried out his research (Martin, 1932). Faraday's personal communications have also been published (Williams, 1971). A number of Faraday historians have analyzed the data leading to the discovery of induction (Tweney, 1987).

Williamson's work on alcohol structure: In 1850, Williamson carried out a series of experiments that allowed him to infer the molecular structure of alcohols. In his classic paper on etherification, Williamson reports how his research progressed (Leicester, 1952).

2.2.8. Problems

Similar to the way we derived the problem P1, we have derived three more problems corresponding to Krebs' work on amino-acid biochemistry, Faraday's work on magneto-electricity, and Williamson's work on alcohol structure. We derived an additional problem based on a research problem in cognitive science at a high level of abstraction.

Problem 2 (P2):

Given

A surprising phenomenon: that Ornithine produces ammonia on kidney tissue slices.

Some Basic Facts: Structural formulae of various substances, a set of values of pH, ordered-lists of amino-acids, amines, and carboxylic acids, stimulators and inhibitors associated with processes. Knowledge about how inhibitors normally affect a reaction.

Previously postulated hypotheses: Amino acids might be producing ammonia in liver by an oxidative, hydrolytic or reductive reaction.

Produce: new facts about amino acid metabolisms.

Problem 3 (P3):

Given

A surprising phenomenon: that common alcohol, ethyl iodide and potash together produce ether.

Facts: Structural formulae of common alcohol, ethyl iodide and other chemicals.

Prior hypotheses: Two hypotheses about the structure of common alcohol, one proposing existence of an ethyl group, other proposing existence an etherin group.

Produce: new phenomena and make other inferences.

Problem 4 (P4):**Given**

A surprising phenomenon: On switching on electric current in a coil around a cylinder, a temporary current is produced in another nearby coil.

Knowledge: Various methods of measurement of electric current. Knowledge about the value-sets and classes of various attributes associated with the experiment. Arago's effect.

Produce: Interesting phenomena.

Problem 5 (P5):**Given**

Surprising phenomenon: When a program X is run on the urea synthesis problem, the number of state transitions of the program that match with the behavior of Krebs is much larger than expected.

Domain Facts: Names of components of X by Kulkarni and Simon (1988), names of various computers, and domains ordered by cost and availability.

Produce: new phenomena and hypotheses about them.

2.2.9. Sensitivity analysis of the problem statements

Problems P2, P3, and P4 have close resemblance with the problems faced by the corresponding scientists. But it should be noted that it is beyond the scope of this thesis to make an evaluation of how simplified these problems are. In forth chapter, where we would describe the behavior of KEKADA, we would make some comments on how KEKADA would behave on slightly different problems.

2.3. The Procedure followed

We started with data on Krebs work on urea synthesis. We detected certain heuristics in this data and built a program which is capable of working on the problem P1 (Kulkarni&Simon1988). We then tried to run the program on the problems P2, P3, P4, and P5 successively. For the program to run successfully on these new problems, we needed to switch to a more general representation and we also needed to add a few new heuristics.

Chapter 3

KEKADA the program

The whole of science is nothing more than a refinement of everyday thinking.
- Albert Einstein (1950)

In this chapter, we describe the program KEKADA. In the first section, we will describe a scenario to give the reader an idea about KEKADA's abilities. In the second section, we will describe the role of dual search in KEKADA. The third section will describe the representation of data and processes in KEKADA. We will describe the role of surprise in KEKADA in the fourth section. KEKADA control structure, which embodies the dual space search and the strategy of focusing on surprises, will be described in the fifth section. In the sixth section we will describe the strategies KEKADA uses to characterize a surprising phenomenon. The final section is a summary.

3.1. Scenario

Below is a partial trace of KEKADA's behavior on a particular surprising phenomenon it encounters. Some of the details have been omitted for the sake of clarity. A more detailed description can be found in the chapter 4.

-
- Focus attention on a surprising phenomenon that Ornithine produces ammonia in kidney.
 - Find out if other amino acids can also produce ammonia.
 - When other amino acids are also found to produce ammonia in kidney, recognize this as the deamination reaction.
 - Consider the oxidative, hydrolytic and reductive hypotheses.
 - Carry out experiments to verify the oxidative hypothesis.
 - Experiments confirm the oxidative hypothesis.
 - Gather more data on other amino acids.
 - Notice that glutamic acid produces an unusual reaction.
 - Focus attention on this unusual reaction.
-

KEKADA begins with its attention on a surprising phenomenon that ornithine produces ammonia in kidney. One of the strategies it then follows is to find out if other amino acids also behave in a similar

manner. When experiments are carried out on other amino acids, it is found that the phenomenon seems to be not specific to ornithine, but is also exhibited by other amino acids. Now KEKADA recognizes that this reaction could possibly be the deamination reaction. It then considers various known hypotheses about the mechanism of the deamination reaction. One of the hypotheses is that it could be an oxidative reaction in which an amino acid is oxidized to produce a keto acid and ammonia. An experiment, carried out to test this, confirms it. Next KEKADA decides to gather more data about the deamination reactions. While carrying out experiments on various amino acids, it notices that glutamic acid produces an unusual reaction. It then decides to focus attention on this unusual reaction.

Thus in this scenario KEKADA showed that a deamination reaction takes place in kidney and by a particular oxidative reaction. Scientific research is a continuous process and at the point where the above scenario ends, KEKADA has found a new puzzle to attend to. As we describe KEKADA, we will use this scenario as an example on a number of occasions. Now we will describe the overall organization of KEKADA.

3.2. Dual Search in KEKADA

The basic source for new knowledge in KEKADA is nature. KEKADA carries out experiments on the external environment to gather new information and to modify confidences in the existing information. Thus it explores a space of rules containing hypotheses and strategies, and a space of experiments containing experiments and results. 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. This organization is shown in figure 3-1.

Let us consider two examples of how KEKADA uses experiments. If KEKADA is studying the hypothesis that a specific substance, ornithine, is acting as a catalyst in a given reaction, it may decide to carry out a specific experiment to verify this. Thus it may experiment to see if a sufficiently large amount of urea can be produced in a reaction in the presence of a small amount of ornithine. If large amount of urea can be produced, it would constitute confirmative evidence about the catalytic hypothesis. Experiments could also be used purely for the purpose of gathering new information. For example, in the scenario in the introductory section, KEKADA decides to gather data about a reaction in which amino acids deaminate. KEKADA carries out experiments on various amino acids one by one. The experiments reveal that glutamic acid produces an unusual reaction.

Now we will proceed to give further description of the problem spaces and strategies KEKADA

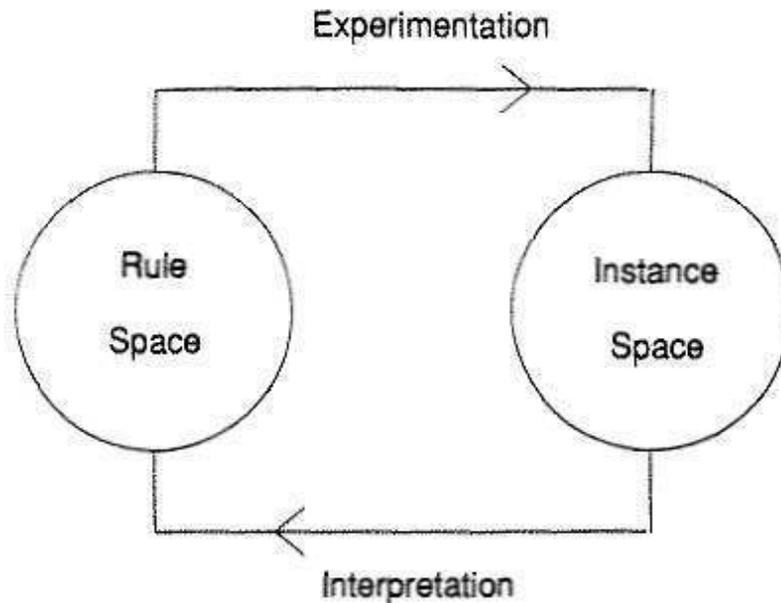


Figure 3-1: Two-space Model of Learning

uses. First we will describe the representation of data and processes of KEKADA. Later we will describe the operators and heuristics it has and how they are effective in controlling the space in which it is searching.

3.3. 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.

3.4. Representation of data and processes

3.4.1. Representation of Data

Working memory elements are represented as attribute-value pairs. Below we describe some of the important categories of working memory elements: experiments, hypotheses/strategies, and confidences in hypotheses/strategies.

Experiment. An experiment has a set of independent entities. An independent entity is any entity over which we have control to decide whether to introduce it in the given experiment. Consider a chemical reaction in which ornithine combines with ammonia to produce urea. Ornithine and ammonia are independent entities, as we have control to decide whether to introduce them in the experiment or not. On the contrary, a dependent entity is any entity over which we do not have a direct control, but which is produced as an effect in the given phenomenon. In the above reaction, urea is a dependent entity.

An entity has a set of variables associated with it. These variables can be independent, dependent or apparatus variables. For example, the entity ornithine may have concentration as an independent variable, and the rate of consumption as a dependent variable. Apparatus-variables are those that play an auxiliary role in the process under consideration and thus are not considered direct causes of the results of the experiments. But the distinction between an apparatus-variable and an independent variable can at times be subjective. Experiments may have one or more goals associated with them.

Below we give an example of how an experiment may be represented.

```
var-of-expt ^type independent ^no 1 ^attribute name ^value ornithine
var-of-expt ^type independent ^no 1 ^attribute conc ^value medium
var-of-expt ^type apparatus-var ^no 0 ^attribute aerobicity ^value yes
var-of-expt ^type dependent ^no 1 ^attribute name ^value urea
expt-goal ^goal g0001
```

Semantics of the example:

The experiment is carried out with medium concentration (conc) of the substance ornithine under aerobic condition. The goal of the experiment is indicated by a symbol (g0001), which is a pointer. It is observed that urea is produced in this reaction.

A Process: Like an experiment, a process may have independent entities, dependent entities, and corresponding independent, dependent, and apparatus variables. A process may further have sub-processes. A surprise, an experiment, an expectation, a process, and a hypothesized process have similar representations.

A Hypothesis/Strategy: A hypothesis or strategy is represented by an internal working memory element called 'description.' This description has a *type* and it also allows the specification of reactants and groups involved. A *type* may thus be 'determine-scope', 'is-catalyst', or 'donates-group.' The specification of a specific hypothesis may also involve a number of reactants. For example:

(description ^name g001 ^type catalyst ^reactant1 ornithine)

This would represent that ornithine might be acting as a catalyst.

(description ^name g002 ^type donates-group ^group amino ^reactant1 ornithine ^reactant2 urea)

This would indicate that ornithine might be donating an amino group to urea.

3.4.2. Representation of Confidence Measures

The confidence in a hypothesis is represented by a 5-tuple:

1. **Success:** the number of experiments that have verified the universal hypothesis about a class or a hypothesis in general.
2. **Failure:** the number of experiments that have falsified the hypothesis.
3. **Failed-effort:** the amount of effort spent unsuccessfully 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.

Below is an example of how confidence in an hypothesis may be represented.

(description ^class-name amino-acid ^name g00004 ^type scope-over-class ^member-value ornithine ^member-attribute name ^member-no 1)

(confidence ^success 0 ^name g00004 ^implied-success 0 ^implied-failure 0 ^fail 0 ^failed-effort 3)

Semantics of the example: This working memory element represents the confidence in the hypothesis that the observed phenomenon may not be limited to ornithine, but similar phenomena may also be exhibited by other members of the class amino-acid. The program has failed to verify the hypothesis on three attempts.

3.4.3. Representation of processes/heuristics

The KEKADA system is implemented in the production system language OPSS (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.

3.5. Goal of KEKADA

3.5.1. The goal of the program

As we discussed in our task analysis section, the goal of experimental scientific research is to produce interesting phenomena. KEKADA employs a set of strategies to control its search in its attempt to achieve the goal.

3.5.2. Focusing on surprise as a search-control strategy

It is well-known that surprises have played a central role in many important discoveries. An example of this is Priestley's work. In the course of years of work producing many important research results, Priestley observed that "the first hints, at least of almost everything we have discovered, of much importance, have occurred to me in this manner (as unexpected phenomena)." (Conant, 1957) Some of the previous AI work has produced evidence of the utility of focusing on the unexpected (Quinlan, 1983, Lenat, et al, 1983). Schank (1982) has argued for a central role for expectation failures in learning. He argues that we would not be able to notice errors in our view of world without setting expectations. KEKADA attends to surprises, thereby searching the part of problem space that is likely to be dense with interesting phenomena. To do this, it associates expectations with each experiment. After the experiment is carried out, if the expectations have been violated, it focuses its attention on the surprise.

3.6. KEKADA control structure

The control structure of the program is based on two of the ideas introduced earlier: the dual search space and the strategy to focus on the surprises. The program has a set of heuristic operators which allow it to carry out this search and which allow it to focus on surprises.

These heuristics fall into the following processes.

1. **Experiment-proposers**, which propose experiments based on existing hypotheses.
2. **Experimenters**, which carry out experiments.
3. **Hypothesis or strategy proposers**: When the system has decided to focus on a particular problem, these decide which hypothesis to focus on or which strategy to adopt for the work on the problem.
4. **Problem-generators**, which propose new problems or subproblems on which the system can focus attention.
5. **Problem-choosers**, which choose the task the system should work on next.
6. **Expectation-setters**, which set expectations for the experiments to be carried out.
7. **Hypothesis-generators**, which generate new hypotheses about unknown mechanisms or phenomena.
8. **Hypothesis-and-confidence-modifiers**, which modify the hypotheses on the basis of new evidence; and which modify confidences about hypotheses on the basis of the interpretations of experiments.
9. **Decision-Makers**, are the 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 alternatives.

In an appendix section, we list the program with the full set of heuristics ordered in each category. Readers may want to refer to it while reading the traces of KEKADA runs. Now we will describe functionality of the program in more detail.

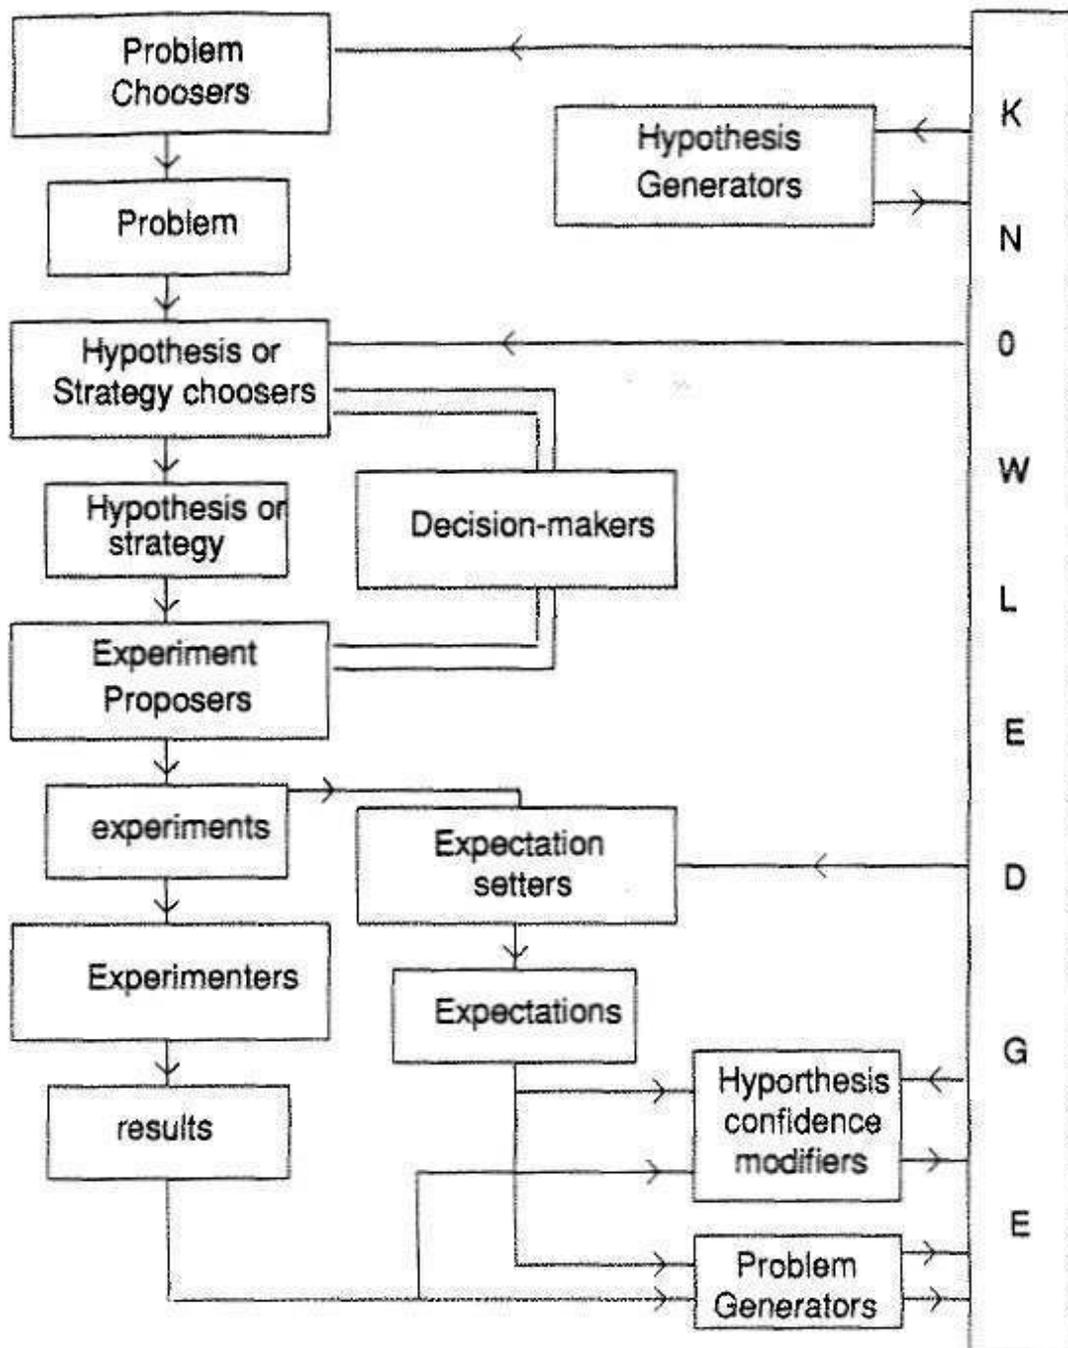


Figure 3-2: Interaction of heuristics

3.6.0.1. Interaction of Heuristics

At any stage the system is trying to understand a given surprising phenomenon. *Hypothesis-generators* create hypotheses when faced with a new problem. Thus at any given stage a certain number of hypotheses or strategies with varying confidences are present in the working memory.

Hypothesis-or-strategy-choosers will choose one or more of the hypotheses or strategies 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 experiments are interpreted by the *hypothesis-and-confidence-modifiers*. When applicable, *problem-generators* may add new problems to the agenda and preempt the system to focus on a different problem. When the system focuses on a new surprising phenomenon, *hypothesis-generators* generate various hypotheses and strategies about this phenomenon.

Thus this control structure allows KEKADA to carry out a two-space search and to focus on a surprising phenomenon opportunistically.

The following trace of KEKADA illustrates how the control structure of KEKADA operates.

[Throughout this thesis, we will be using the following short-forms: HSC: Hypothesis and Strategy Choosers, DM: Decision Makers, EP: Experiment Proposers, E: Experimenters HCM: Hypothesis and Confidence Modifiers, PG: Problem Generators, PC: Problem Choosers, ES: Expectation Setters.]

HSC1, DM1	Decide to verify the oxidative reaction.
EP2	Decide to carry out experiments on alanine and oxygen together.
E1, HCM6	Alanine and oxygen together produce keto-acid with q-rate of 10 and ammonia is also produced at q-rate of 10. Oxygen is consumed at the q-rate of 5. Conclude that the oxidation hypothesis of deamination is correct.
HSC1, DM1	Decide to gather data on deamination of amino-acids.
EP10, DM2, ES*, E1	Try out experiments on various amino-acids.
	(Last one is glutamic-acid: q-rate of oxygen consumption 5, q-rate of production of keto-acid is 0.)
PG1, PC8	Make the glutamic acid reaction the focus of attention.
HG2	There might be a missing inhibitor.

The semantics of this trace were discussed earlier in this chapter. The example illustrates how

different types of heuristics interact in KEKADA. We started with HSC and DM heuristics choosing a particular hypothesis, an EP heuristic suggested an experiment. An E heuristic carried out the experiment, and a HCM heuristic modified the confidence in an existing hypothesis. In the latter half we see how an ES heuristic can be used to set expectations for an experiment and PG and PC heuristics can be used to focus attention on a surprising phenomenon.

3.6.1. Implementation of the 'surprise' heuristic

Whenever an experiment is to be carried out, expectation-setters associate expectations with the experiment. Experimenters carry out the experiment. Problem-generators check if expectations are violated. Problem-choosers choose a specific problem and decide to focus attention on it.

3.6.2. Expectation-setters

Knowledge of previous experiments is used to set expectations for new experiments. This knowledge is stored in summary elements called 'expectation' and 'var-of-expectation.' As we discussed earlier, every experiment has a set of independent and dependent variables associated with it.

In the summary element, expectations are associated with some of the dependent variables. If the attribute is non-numeric, then the expectation is a symbolic value stored in 'expected-value.' Thus for an experiment on an amino acid, the expected output-reagent is a non-numeric attribute and would have a value such as 'urea' associated with it. If the attribute is of the type numeric, then the expectation has two numeric values -- the lower bound and upper-bound. For example, the rate of formation of a reagent is a numeric attribute and therefore has a lower-bound and an upper bound. 'no-of-expt' gives the number of experiments which were used in the derivation of a specific expectation. When the expectations are derived purely from analytic considerations or using methods inferior to methods being used currently, the no-of-expt is set to 0.

Expectation summaries are stored in the form of a tree of cues to limit the search. (For convenience of implementation, we use a simplified EPAM model of memory.) A number of summaries may match against the experiment to be carried out. For example, we may expect an amino-acid in general not to produce urea in liver, we may expect a specific amino acid, arginine, to produce a large amount of urea. The summaries associated with the amino-acid class and with arginine both would match against the experiment involving arginine. Two rules are used to resolve the conflicts - summaries with the larger number of cue-matches and with more specific cue-matches are preferred. Thus the

expectation about arginine would be used to set expectations about an experiment with arginine. Once an expectation summary element is chosen according to these conflict resolution rules, the following rules are used to set the expectations of the experiment.

[ES1] Expected lower-bound is set to be the lowest of the 'lower-bounds' of matching expectation-summaries. if the attribute is numeric.

[ES2] Expected upper-bound is set to be the highest of the 'higher-bounds' of matching expectation-summaries. if the attribute is numeric.

[ES3] For an attribute with symbolic value, expected value is set to be the the expected-value associated with the matching summary element.

After an experiment is carried out, a match is carried out against the tree of cues, and expectation-summaries in the matching elements are modified according to the following rules.

[ES4] If the previous expectations were not based on experimental evidence and the attribute is numeric, then 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. The expected-value with a symbolic attribute is that observed in the experiment.

[ES5] If the previous expectations were based on at least one experiment, and the present experiment violates the bounds or value in the expectation-summary, update it.

To illustrate how expectation-setters work, consider the following example. When KEKADA decides to work on the urea synthesis problem, it carries out experiments with alanine and a number of other amino acids. These amino acids produce urea at a very low rate (Q-rate = 1). Now assume tolerance factor to be 4. After the experiment on alanine is completed, ES4 would set the expectation summary with a lower-bound on Q-rate of 0 and a higher of 5. After the alanine experiment, experiments are carried out on a number of other amino acids. In each case expectations would be associated with the experiment by ES1-3: that the q-rate be in the range 0-5 and the output be urea, so that whenever an experiment is carried out on an amino-acid, the resultant q-rate is expected to be in this range. One of the last amino acids is ornithine. Ornithine produces urea at a much higher rate (Q-rate = 9). Hence, ES5 heuristic resets the upper-bound on the expected output for amino-acids to be 9. Now if an experiment is later carried out with an amino acid on liver slices, KEKADA will expect the output to be urea with q-rate in the range from 0 to 9.

3.6.3. Problem-generators

KEKADA tests for any violations of expectations, as the experiments that violate expectations, provide a pointer to a part of problem space dense with interesting phenomena.

[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.

For example when KEKADA carries out an experiment on glutamic acid in kidney, it finds that q-rate of production of keto-acid is 0, which is lower than the minimum expected amount. The PG1 heuristic in such a case would detect a surprise and add it to the agenda.

3.6.4. Problem-Choosers

At present we have a computationally simple scheme for choosing a new problem. KEKADA focuses on any surprise it encounters.

[PC1] 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.

3.6.5. Hypothesis or Strategy Choosers

[HSC-1] Evaluate the alternative hypotheses or strategies and choose one or more of them for consideration using the decision-making rules.

3.6.6. Experimenters

In the current system, there are no experimentation heuristics.

[E1] The outcomes of experiments are stored in certain working memory elements and are directly copied from them.

3.7. Strategies KEKADA uses to characterize a surprising phenomenon

3.7.1. Characterization of a surprising phenomenon

In the previous section we discussed how KEKADA uses a strategy to focus on a surprise to attend to the part of the problem space rich with interesting phenomena. A surprising phenomenon indicates that the prior knowledge was deficient in some dimension. KEKADA employs following set of strategies to characterize a surprising phenomenon.

1. Magnify the phenomenon by varying apparatus variables.
2. Divide and conquer: the surprising effect may depend on one of the subprocesses.
3. Assess the scope of the surprising phenomenon.
4. Determine if all the independent entities are necessary to produce the surprising phenomenon.
5. Try to find a relation between similar phenomena.
6. Gather more data about a surprising phenomenon.
7. Use domain-specific strategies.

We will now discuss these strategies one by one, and describe how they are implemented.

3.7.2. Magnification of the phenomenon

KEKADA may try to magnify an effect by carefully changing the setup of the apparatus and the methods of measurement, that were used to produce the surprising phenomenon. This increases the chances of making crucial observations on further experimentation with the surprising phenomenon. After magnification, the surprising phenomenon may become far more interesting and important. When Roentgen found new kinds of rays which could penetrate cardboard, that was surprising. The next step Roentgen took was to see if these rays could also penetrate the walls, and found that they do. This amazing penetrating power of these rays immediately made clear the importance of the phenomenon. The following productions are used in KEKADA's implementation of the strategy of magnification.

[HG9] *If a phenomenon is found to be surprising, then a possible strategy is to attempt to magnify the effect (or make it conveniently visible) by changing the apparatus variables associated with the phenomenon.*

[EP9] *If the chosen strategy is to magnify a certain effect, then carry out experiments varying the value of each apparatus variable over the set of values associated with this variable. Note that a variable may be associated with the method of the measurement of a dependent variable.*

[HCM-9] *If the goal of the experiment was to magnify the effect or make it more visible, and it is observed that by changing the value of some apparatus variable the phenomenon is magnified, decide to carry further study of the surprising phenomenon with the new value of the apparatus variable.*

3.7.3. Divide-and-conquer strategy

If the surprising phenomenon is known to contain subprocesses, KEKADA may focus on one of the subprocesses, assuming that the surprising result may depend on it. This is a specific implementation of the well-known 'divide-and-conquer' strategy. The following productions are related to the implementation of this strategy.

[HG3] If there is a hypothesis that a phenomenon has subprocesses and the phenomenon is noted as surprising, hypothesize that the surprising result depends on one of the subprocesses (divide-and-conquer strategy).

[EP7] If the chosen strategy is divide-and-conquer, carry out each of the subprocesses of the phenomenon under various conditions.

3.7.4. Determination of the scope of the puzzling phenomenon

KEKADA may try to assess the scope of the surprising phenomenon using domain-specified taxonomies. When Fleming observed *Penicillium* had the property of killing particular bacteria, the questions which were later pursued were: Can other molds kill bacteria? Can *Penicillium kili* similar other bacteria? Following productions are related to the implementation of this strategy.

[HG1] If an independent variable <at> associated with an independent entity of a surprising phenomenon has value <v> and <c1> is a class of values of <at> containing <v>, then consider the strategy to determine the scope of the phenomenon over the range of values specified by <c1>.

[EP1] If the preferred strategy is to assess the scope of a surprising phenomenon over a class of values of an attribute of an independent entity, then use the decision-makers to choose a value A in that class, and decide to study the phenomenon with A as the value of the variable.

[HCM-5] If the goal of the experiment currently carried out was to assess the scope of a surprising phenomenon; then check the similarity between the resulting phenomenon and the surprising phenomenon and accordingly modify the confidences in the hypothesis guessing the scope of the phenomenon over a class.

[HCM-7] If there is a hypothesis that the surprising phenomenon may have scope over a class <c> And the success-slot in the confidence of this hypothesis exceeds the threshold value for generalization, generalize the surprising phenomenon and transfer the control to the hypothesis-generators.

[HCM-3] If the amount of effort spent on an existential hypothesis reaches a specified high value (which we have assigned to be 3), make the hypothesis inactive.

3.7.5. Factors analysis

Given a surprising phenomenon, one may ask which variables are actually causing the surprise. Some scientific disciplines use "control experiments" to determine whether a particular factor is in fact necessary for the production of an observed phenomenon. Determining exactly which of the factors are necessary to produce the phenomenon is important in gaining a better understanding of it. The following productions are related to the implementation of the factors analysis strategy.

[HG7] If the phenomenon has two or more independent entities, then consider the strategy of deciding whether all entities are necessary to produce the phenomenon.

[EP3] If the preferred hypothesis is that the phenomenon has A and B as 2 independent entities, carry out experiments on A and B in combination and on A and B separately. An independent variable associated with A may be dependent on some variables associated with B and vice versa. In such a case, the experimenter should add additional apparatus to give causal support to the variable.

3.7.6. Relating similar phenomena

KEKADA may look for phenomena that are similar to the surprising phenomena in some way and then try to find a relation between these closely related phenomena. In particular, if two anomalous effects are being observed and both include the same variable, then there is some chance that there is a common hidden mechanism. This is due to our belief in the principle of uniformity of nature. Thus when Hans Krebs was working on the ornithine effect, there were two odd reactions. First ornithine was the only amino acid producing a lot of urea in the presence of ammonia. Second a reaction was also known in which arginine hydrolysed to produce ornithine and urea, a reaction peculiar to arginine. It was reasonable to guess that there might be some relation between these two peculiar reactions. In this case there was in fact the relation that is shown in fig 2-1. KEKADA uses both domain-specific and domain-independent strategies to guess such a relation. One domain-independent strategy conjectures that there might be a whole class of substances exhibiting a common effect. The following productions are related to the implementation of this strategy.

[HG6] If the goal is to study a puzzling phenomenon and another phenomenon and the surprising phenomenon contain two common dependent or independent entities, then create a hypothesis that the other phenomenon may be related to the surprising phenomenon.

[EP8] If the preferred hypothesis is to study the relation of another phenomenon to a surprising phenomenon, then create the following hypotheses and add them to the hypothesis set:

(a) If the surprising phenomenon and the related phenomenon have a common dependent entity $\langle u \rangle$ and a common independent entity $\langle v_{na1} \rangle$; and $\langle v_{na1} \rangle$ has an attribute $\langle at \rangle$; and the surprising phenomenon and the related phenomenon have values of $\langle at \rangle$, of $\langle va \rangle$ and $\langle va2 \rangle$ respectively, a possible strategy is to try to find out if there is a set of values of that attribute containing both $\langle va \rangle$ and $\langle va2 \rangle$ which will exhibit the same phenomenon.

(b) If the surprising phenomenon and the related phenomenon have a common dependent entity $\langle u \rangle$ and the surprising phenomenon has an independent entity $\langle va \rangle$ and the related phenomenon has an independent entity $\langle va2 \rangle$, a possible strategy is to try to find a set of entities containing both $\langle va \rangle$ and $\langle va2 \rangle$ that will exhibit the same phenomenon.

(c) If the phenomenon is a chemical reaction, the surprising reaction and the related reaction have a common output $\langle u \rangle$ and the related phenomenon has an input $\langle i \rangle$, then create the hypothesis that $\langle i \rangle$ is an intermediate in the surprising reaction.

3.7.7. Gather-data strategy

When one finds a novel and unusual phenomenon, even systematic collection of data about it can be of great interest to the scientific community. Furthermore while collecting such data, surprises can turn up. HG12 and EP10 are the heuristics related to this strategy. In the scenario in the beginning of the chapter, we saw how KEKADA comes across the glutamic acid effect, while gathering data on the deamination reaction.

[HG12] If one of the variables of the surprising phenomenon is a class, then a possible strategy is to gather more data by carrying out experiments on the various members of this class.

[EP10] If the preferred strategy is to gather data about a phenomenon and an attribute of an independent entity is a class, then use the decision-makers to choose a value A in that class, and decide to study the phenomenon with A as the value of the attribute.

3.7.8. Activation of previously known hypotheses

If the surprising phenomenon is thought to be a previously known type of process, then one should consider previously known hypotheses about that process. Thus in the scenario in the section 3.1, when KEKADA posits that the surprising phenomenon in which amino acids are producing ammonia in kidney as the deamination reaction, it activates known hypotheses about the deamination, namely, oxidative, reductive, and hydrolytic processes.

[HG11] Activate any previously known hypotheses about the surprising phenomenon.

3.7.9. Domain-specific Strategies

The domain-specific strategies in KEKADA are actually specializations of following general strategies:

- * Testing a causal chain.
- * Controlled experimentation / Factor Analysis
- * Mill's Difference Principle
- * Determining the complexity of a process
- * Successive Refinement
- * Conservatism
- * Direct Verification

3.7.9.1. Testing a causal chain

One strategy proposes experiments to test a causal chain $A \rightarrow B \rightarrow C$ where A, B, C are three events. The fact that B is an intermediate event implies that B should be able to cause C; furthermore B should be able to cause C no more slowly than A can. In domains where we can produce B as an independent event under experimental conditions and where means are available to measure the rates of occurrence of C, a possible strategy would be to see if B can produce C at all. Further if differential speeds are meaningful (e.g. the causation is not instantaneous), then one would also measure the rates of formation from A and from B and compare them. The domain of metabolic biochemistry in the early 1900s satisfied these properties. KEKADA uses the following productions to implement this strategy.

[EP2] If there is a hypothesis that the chemical reaction under study contains 2 subreactions PR1 and PR2 one followed by the other, and I1 is the set of inputs to PR1, and I2 the set of inputs to PR2, then study these two reactions, measuring the rates of formation of the outputs.

[HCM-1] If the domain is chemistry And the goal of two of the experiments currently carried out is to study the hypothesis that B is an intermediate in the reaction from ACL to C And these two experiments measure the rates of formation of C from A and from B, And A is the member of the class ACL, modify the implied-success or implied-failure slot in the confidence about the above hypothesis depending on whether there is faster formation from B or from A.

3.7.9.2. Controlled Experimentation / Factor Analysis

Earlier we discussed the role of controlled experimentation and factor analysis in the study of complex phenomena. In some domains, there are many background variables and processes. Thus two entities A and B under study may both be producing another entity C at certain base rates independently of each other. We may want to ask: do these entities interact at all when we employ both of them? If in the domain there is a property associated with C which will allow us to answer this question by measuring P from A, B separately and together, then one may resolve the matter. For example in chemical reactions, we might look at the rates of formation of C. Then if A and B together produce more than the sum of what is produced by A and by B separately, there is an interaction between A and B. Thus say ornithine produces urea at the rate of 1, ammonia at the rate of 4, and both together at the rate of 9, one should conclude that both reactants are involved in the reaction. If the output rate were only 5, one can not make such a conclusion. The following productions allow such an interpretation. Experiments needed for such interpretations are suggested in general by the factors-analysis heuristics.

[HCM-4] If the domain is chemistry ; And the goal of three of the experiments currently carried out is to study the hypothesis that ACL and B react together to form C; And these three experiments measure the rates of formation of C from A, and from B, and from A and B together; And A is a member of the class ACL; modify the implied-success or implied-failure slot in the confidence about the above hypothesis depending on which of the following two is greater: the rate formation from A and B together, or the sum of the rates from A and from B. (Allow for an error tolerance factor.)

3.7.9.3. Mill's Difference Principle

Mill's Difference Principle states that if under conditions C1 phenomenon P occurs and under C2 it does not occur, then the differences between C1 and C2 are causally responsible for the phenomenon P. For example we know that urea forms in the body from alanine, but under experimental conditions of a particular experiment it does not. We may then conclude that certain substances are missing from the experimental setup. In the metabolic reactions in the tissue slices, the suspected missing substances would be energy-producing stimulators such as glucose.

When a surprising phenomenon occurs, one possible hypothesis is the presence of other interfering phenomena. If we know means of removing these interferences, we would use such means to eliminate them. Such strategies would mostly be domain-specific. In metabolic biochemistry, one could use substances called inhibitors to remove these interfering phenomenon.

This strategy is implemented by the following productions.

[HG2] If substances previously known to influence the phenomenon were absent from the surprising phenomenon, then hypothesize that the absence of such an activator/inhibitor is the causal factor behind the surprise. (Set priority equal to be 2.)

[EP6] If the chosen hypothesis is that the reason for a surprising outcome is in the absence of some entity, choose one of the entities that earlier experiments seem to have associated with the given class of processes and study the effects of adding this entity to the independent entities associated with the surprising phenomenon.

[HCM-2] If the goal of the experiment is to study the hypothesis that the cause of the surprising phenomenon lies in the absence of an independent entity And in the experiment which was just carried out, the entity currently guessed to be missing did not have any effect on the phenomenon; increase "failed-effort" slot in the confidence that an independent entity is missing, by 1.

3.7.9.4. Reasoning about structural components

The STAHL and DALTON programs (Langley, et al, 1987) incorporated a set of general heuristics to generate componential and structural models of substances from a set of reactions. In the problems KEKADA faces, it uses the same reasoning in the reverse situation. It knows the components and structures of the molecules, but it may not have a complete understanding of the observed reaction. Thus it may know that ornithine and ammonia produce urea, but it may not know if auxiliary reactants are involved and how many molecules of ornithine and ammonia are involved in the reaction. KEKADA uses structural reasoning to constrain in such a situation. Thus if ammonia has an amino group and urea also an amino group, it would create an hypothesis that ammonia contributes its amino group towards the urea. Thus it is reasoning similar to STAHL in a slightly reversed situation. (STAHL would identify components of urea from a reaction such as "2 molecules of ammonia and one of CO₂ form urea." In the situation KEKADA faces, it has the componential models, but it does not have the exact description of the reaction.)

[HG4] If a chemical reaction produces some output (with q-rate or quantity above a minimum threshold), create hypotheses asserting which reactant donates which group to the output substance and if there is more than one reactant, then a reactant may be a catalyst. If the reaction has only one known input and only one output, then guess that there must be other auxiliary inputs among substances around or other auxiliary unknown outputs. Note it is possible that more than one structural formula has been hypothesized about a given input or output. (Set priority to be 6.)

[HG8] If the surprising phenomenon is a chemical reaction and the input reactant <i1> is hypothesized to have the formula <ii1> or <ii2>, the group <gr> is present in the formula <ii1> but not <ii2>, none of the other input reactants contains the group <gr>, and one of the output reactants contain the formula <gr>, conclude that the reactant <i1> has the formula <ii1> and not <ii2>.

[HCM-8] If in a chemical reaction, a small amount of an input can produce large amounts of the output, conclude it acts as a catalyst and there exists an intermediate in the catalytic reaction, and apply HCM-10.

[EP5] If the chosen hypothesis is that the reactant A in an experiment is a catalyst, then carry out the experiment over long periods but with very low initial quantity and concentration of A. Measure final quantities of all outputs.

[EP4] If the phenomenon under consideration is a chemical reaction and the preferred hypothesis is that <in> donates the group <gr> to the <out>, then carry out the reaction making a special effort to measure the rate of consumption of <in> and the rate of formation of <out>.

3.7.9.5. Determination of the complexity of a process

If a process is a simple one-step process, one or more properties would be true of it. (These properties often could be derived directly from the domain's definition of "one-step" process.) If these properties are not true of a process P, then it can be inferred to be a multi-step process.

[HG5] If the given phenomenon is a chemical reaction and a one-step reaction from inputs to outputs of a reaction is found not to be possible with only two inputs, then create the hypothesis that an intermediate exists. (Set priority to be 6.)

3.7.9.6. Successive Refinement

Science follows the strategy of successive refinement of its theories as a practical necessity. For example, a science that aims to understand the metabolic reactions in the body would not be able to hypothesize a complete theory of pathways in the body and be able to propose experiments to test such a theory. KEKADA uses the successive refinement strategy in its work on metabolic reactions.

[HCM-10] If the preferred strategy is to verify the existence of an intermediate in an experiment, 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) that has been evaluated most likely to be an intermediate in the reaction. If in one of the sub-reactions there is more than one input, conclude that there is yet another intermediate. Ask user to carry out literature survey and recursively apply HCM-10.

3.7.9.7. Conservatism

If phenomenon P violates model M, a conservative strategy hypothesizes that model M is correct, but there is an additional process causing the violation. For example, recently some experiments have reported results that deviate from the predictions of Newton's Gravitational laws. In response to these anomalies, it has been proposed that there may be an additional component force to gravity (Poole, 1988). Conservatism would not suspect the validity of previous knowledge if a relatively simple hypothesis can explain an anomaly. This strategy has been implemented by following productions.

[HG10] If the phenomenon under study is a chemical reaction and the incremental q-rate of an output rises unexpectedly on adding an inhibitor, then conclude that this reactant is being consumed in a chemical reaction. (Thus it may be either degrading or reacting with one of the other reactants.) (Set priority to be 6.)

3.7.9.8. Direct Verification

The following heuristics are used to suggest an experiment to test a hypothesis directly.

[EP11] If the hypothesis under consideration is that reactant <r1> may be involved in an unknown reaction and there exists a hypothesis that <r1> and <v> react together and <v> is a reactant in the surprising phenomenon, carry out a reaction with <r1> and <v> as the reactants.

[HCM-6] If the goal of the experiment last carried out was to verify whether a hypothesized process description is the correct process description of a surprising phenomenon, And the results of the experiment confirm this, conclude that this hypothesized process description is the correct description of the surprising phenomenon.

3.7.10. Decision-makers

Decision-makers are used to choose (1) between alternate hypotheses/ strategies (2) between a number of substances. Each strategy and hypothesis in KEKADA has a preference level associated with it. KEKADA uses the preference levels to choose between the applicable strategies.

Magnification: (Set priority to be 1.) Divide and conquer: (Set priority to be 3.) Scope: (Set priority to be 4.) Factor analysis: (Set priority to be 5.) Relating similar phenomena: (Set priority to be 6 or 7 depending on the number of common variables.) Gather-data strategy (Set priority to be 8.) Domain-specific strategies (A few domain-specific hypotheses have high preferences, but most of them are considered after the general strategies are used.)

In the situations in which KEKADA has been run, this preference structure creates a reasonable ordering among the strategies. The behavior of the program is however not overly sensitive to the

ordering. When we compare KEKADA's ordering with the behavior of specific scientists, we find that the scientists follow an order which is slightly different from KEKADA in some cases. In such cases we allow the user to reorder the hypotheses.

DM1 production implements this strategy.

[DM1] The order in which the hypotheses or strategies are considered is based on the user-specified priorities, which closely resemble the program-defined priorities.

Furthermore at times KEKADA needs to choose among alternative members of a class. e.g. it may have to choose among a number of amino acids. The DM2 heuristic implements this.

[DM2] The substances are stored in the form of an ordered list. (It is assumed that this list has been ordered by cost and availability criteria.)

3.8. Summary

In this chapter we have described KEKADA. Some of its important features of the program are:

- * 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.

- * The program is written in the production system language OPS5.

- * The global goal of the program is to attend to puzzling phenomenon (interesting phenomenon) and try to understand them better.

- * The program employs the following processes: Hypothesis-generators, Hypothesis-or-strategy-choosers, experiment proposers, decision-makers, expectation-setters, hypothesis-or-confidence-modifiers, problem-generators, problem-choosers.

- * KEKADA has some domain-independent and some domain-specific strategies. Domain-independent strategies include magnification of the phenomenon, divide-and-conquer, assessing the scope, factor-analysis, relating similar phenomenon. Domain-specific strategies are specializations of general strategies such as controlled experimentation, determination of complexity of a process, testing a causal chain, componential analysis, differencing, and divide-and-conquer.

Chapter 4

Program Behavior

A man, viewed as a behaving system, is quite simple. The apparent complexity of his behavior over time is largely a reflection of the complexity of the environment in which he finds himself.
- Herb Simon(1981).

In this chapter, we will describe the behavior of KEKADA on the research problems defined in the second chapter and try to analyse the reasons for its success in producing important research results on these problems.

4.1. Behavior in the biochemistry of urea metabolism

Hans Krebs' discovery, in 1932, of the urea cycle was a major event in biochemistry. The problem that Krebs attacked, to discover how urea was synthesized in living animals from the decomposition products of proteins, had been investigated for many years with very limited success. The general nature of the catalytic process discovered served as a prototype for much subsequent research and theory on metabolic phenomena. (The historical account of this discovery is derived mainly from Holmes, 1980.)

4.1.1. Initial Working memory of the program

KEKADA is given the following knowledge: Structural formulae of various substances, a set of values of pH, ordered-lists of amino-acids, amines, and carboxylic acids, stimulators and inhibitors associated with processes. We run KEKADA with its focus of attention on the following surprising phenomenon: that alanine produces very little urea in liver tissue slices. KEKADA also has two previously known hypotheses about how urea might be synthesized. These hypotheses are: amino acids might somehow combine with ammonia to produce urea, or also that amino acid may deaminate producing ammonia which may further convert to urea.

4.1.2. Overview of the KEKADA behavior

We divide our account into three phases: discovery of the ornithine effect, the determination of scope, and the discovery of the reaction path.

1. **The ornithine effect.** KEKADA begins with its focus of attention on a puzzling phenomenon that alanine can't produce much urea in liver tissue slices. It tests the efficacy of various amino acids in producing urea, with generally negative results. When it carries out the experiment with ornithine (one of the less common amino acids) and ammonia, an unexpectedly large amount of urea is produced. It then focuses on the ornithine effect.

2. **Determination of scope.** KEKADA next follows a standard strategy: if a given compound exerts a particular action, check if homologues and other similar compounds have a similar action. KEKADA carries out tests on amino acids, amines similar to ornithine. But none of these substances has effects comparable to ornithine.

3. **Discovery of reaction path.** KEKADA now seeks to elucidate the mechanisms of the ornithine effect. Concluding from the quantitative data that the ornithine could only be a catalyst, KEKADA infers that ornithine with ammonia produces arginine, which in turn produces urea and ornithine. Later experiments indicate that citrulline is an intermediate substance between ornithine and arginine.

We must now spell out the details of KEKADA's experiments and reasoning somewhat more fully.

4.1.3. The Ornithine Effect Discovery

KEKADA begins its exploration with its attention on a surprising phenomenon. In an experiment it is observed that a tissue slice with alanine produces very little urea, much lesser than expected.

There are two previously known hypotheses about urea formation.

1. Amino-acids may produce urea with ammonia as intermediate.
2. Amino-acids may combine with ammonia to produce urea.

By examining the alanine to urea reaction, it is also concluded that an intermediate exists in the reaction.

KEKADA also generates the following hypotheses and strategies in response to the puzzling phenomenon.

1. One possible strategy is to magnify the phenomenon by varying the apparatus variables.
2. Since alanine on liver tissue slice does not produce urea, and since it is assumed that

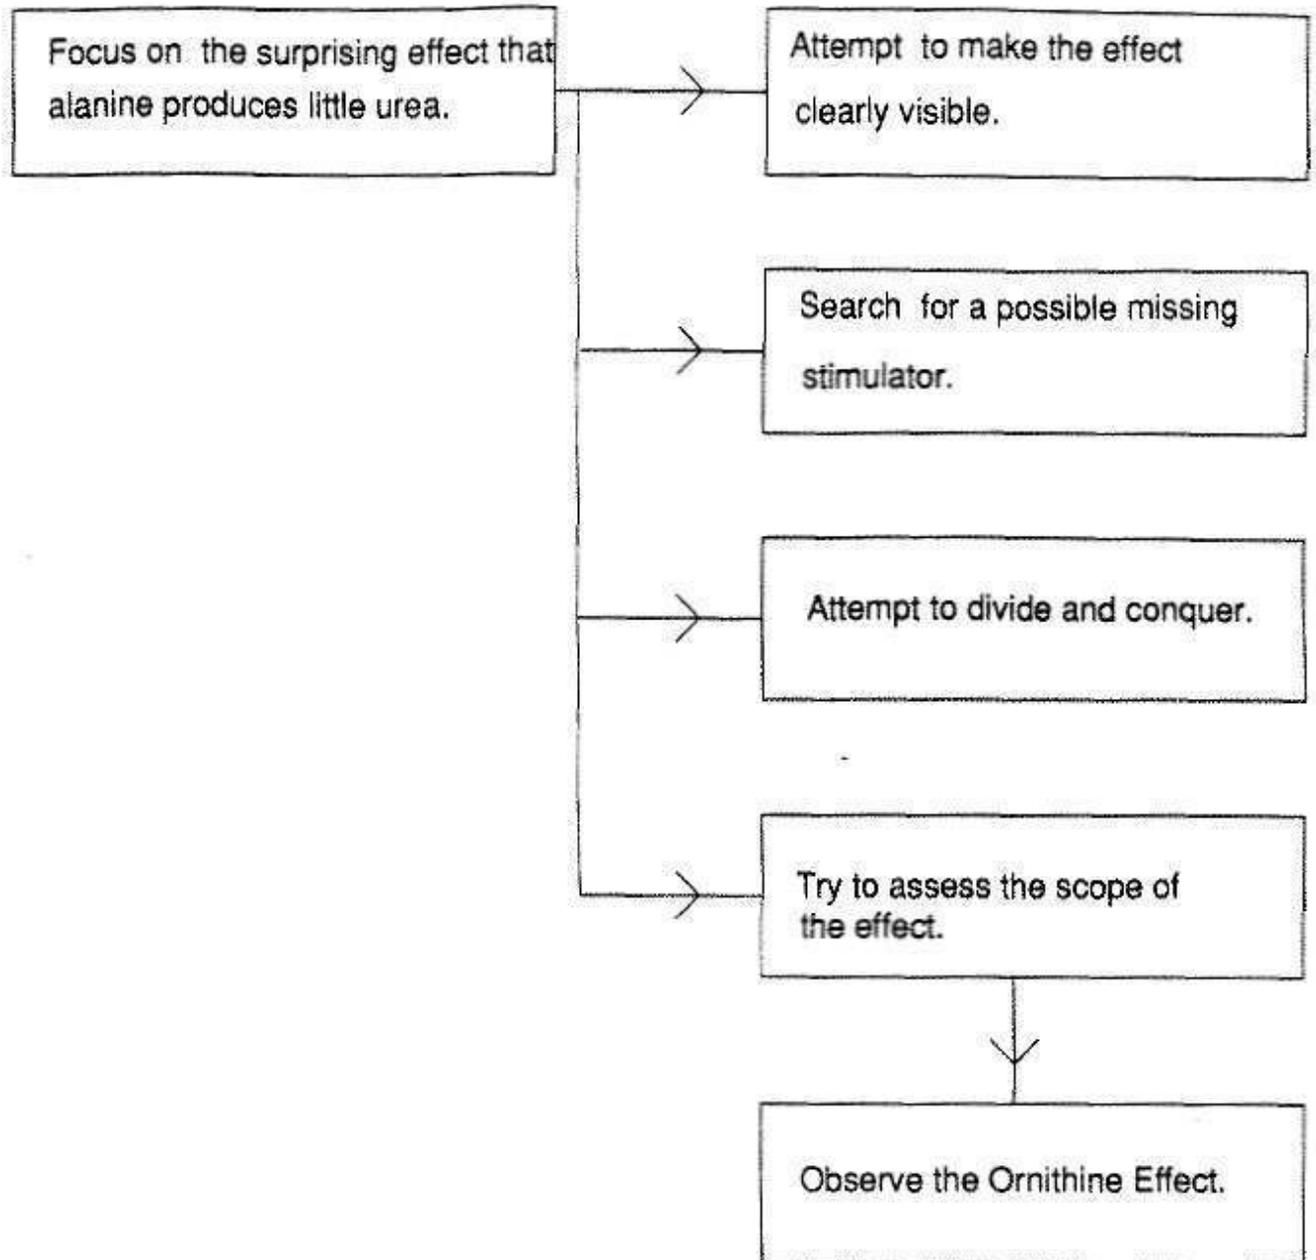


Figure 4-1: Discovery of Ornithine Effect

alanine in the living organism does produce urea, there must be some stimulator, present in the organism, that is missing from the tissue slice preparation.

3. 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.
4. A possible strategy is to determine if the phenomenon is valid only for alanine, or is valid for other amino-acids, amines, or carboxylic acids.

The various experiments that the system now carries out are driven by these hypotheses or strategies. At the beginning, the system has internally defined priorities about various hypotheses and strategies. The system-defined priorities at times differ from the order in which the scientist considered these hypotheses. To give a better fit to the data about a specific scientist, we allow the user to reset the priorities.

The first strategy chosen is to try to magnify the phenomenon by varying the apparatus variables associated with the phenomenon. Absence of production of urea might be due to improper physiological conditions. By changing them one might be able to get alanine to produce urea. However such attempts fail.

In response to the possibility that there is a stimulator in whose presence alanine produces urea, the system tries to identify the stimulator. KEKADA adds such substances as glucose and fructose, without any change in outcome. These results do not falsify the hypothesis that there might be a substance in whose presence alanine would produce urea, but they do reduce confidence in the hypothesis. Each failed guess of the stimulator 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 divide-and-conquer hypothesis leads KEKADA to study each of the subprocesses individually. The two subprocesses involved are the formation of urea from ammonia and the the formation of ammonia from the amino-acid. KEKADA repeats experiments on formation of urea from ammonia 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. Similarly, KEKADA studies the other subprocess involved.

The strategy next considered is to try to determine the scope of the surprising phenomenon. While trying to determine the scope, KEKADA tries to explore the validity of the two hypotheses about the mechanism of urea synthesis: that an amino acid and ammonia may combine to form urea or that an amino acid may produce urea with ammonia as an intermediate. KEKADA now carries out experiments with different amino acids. It carries out experiments on the amino-acid, on ammonia separately and on the two combined together. 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. The experiment shows that ornithine alone 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.

4.1.4. Determination of Scope

The ornithine effect now becomes the focus of attention. KEKADA considers a number of different possible hypotheses and strategies about it.

1. Try to magnify the effect by varying the apparatus variables.
2. Assess the scope of the surprising phenomenon.
3. It is possible that ornithine or ammonia may be acting as a catalyst.
4. An intermediate exists in the chemical reaction in which ornithine and ammonia are inputs and urea is the output.
5. Both the inputs are necessary for the reaction to occur.
6. Possible hypotheses about which input reactant donates which group to which output reactant.
7. This effect may be related to the arginine reaction.
8. Problem may be in one of the sub-reactions of the process.

The first strategy chosen is to try to magnify the phenomenon by varying the apparatus variables associated with it. In this case the fact that alanine was observed not to produce urea, might be due to the fact that physiological conditions such as pH are not proper. By changing them one might be able to get alanine to produce urea. However these attempts fail to turn up anything interesting.

The next strategy considered is to try to assess the scope of the puzzling phenomenon. KEKADA consider three possible groups of substances that may exhibit the ornithine effect: (1) certain carboxylic acids, (2) certain amino acids, and (3) certain amines.

A whole series of experiments is carried out with these 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.

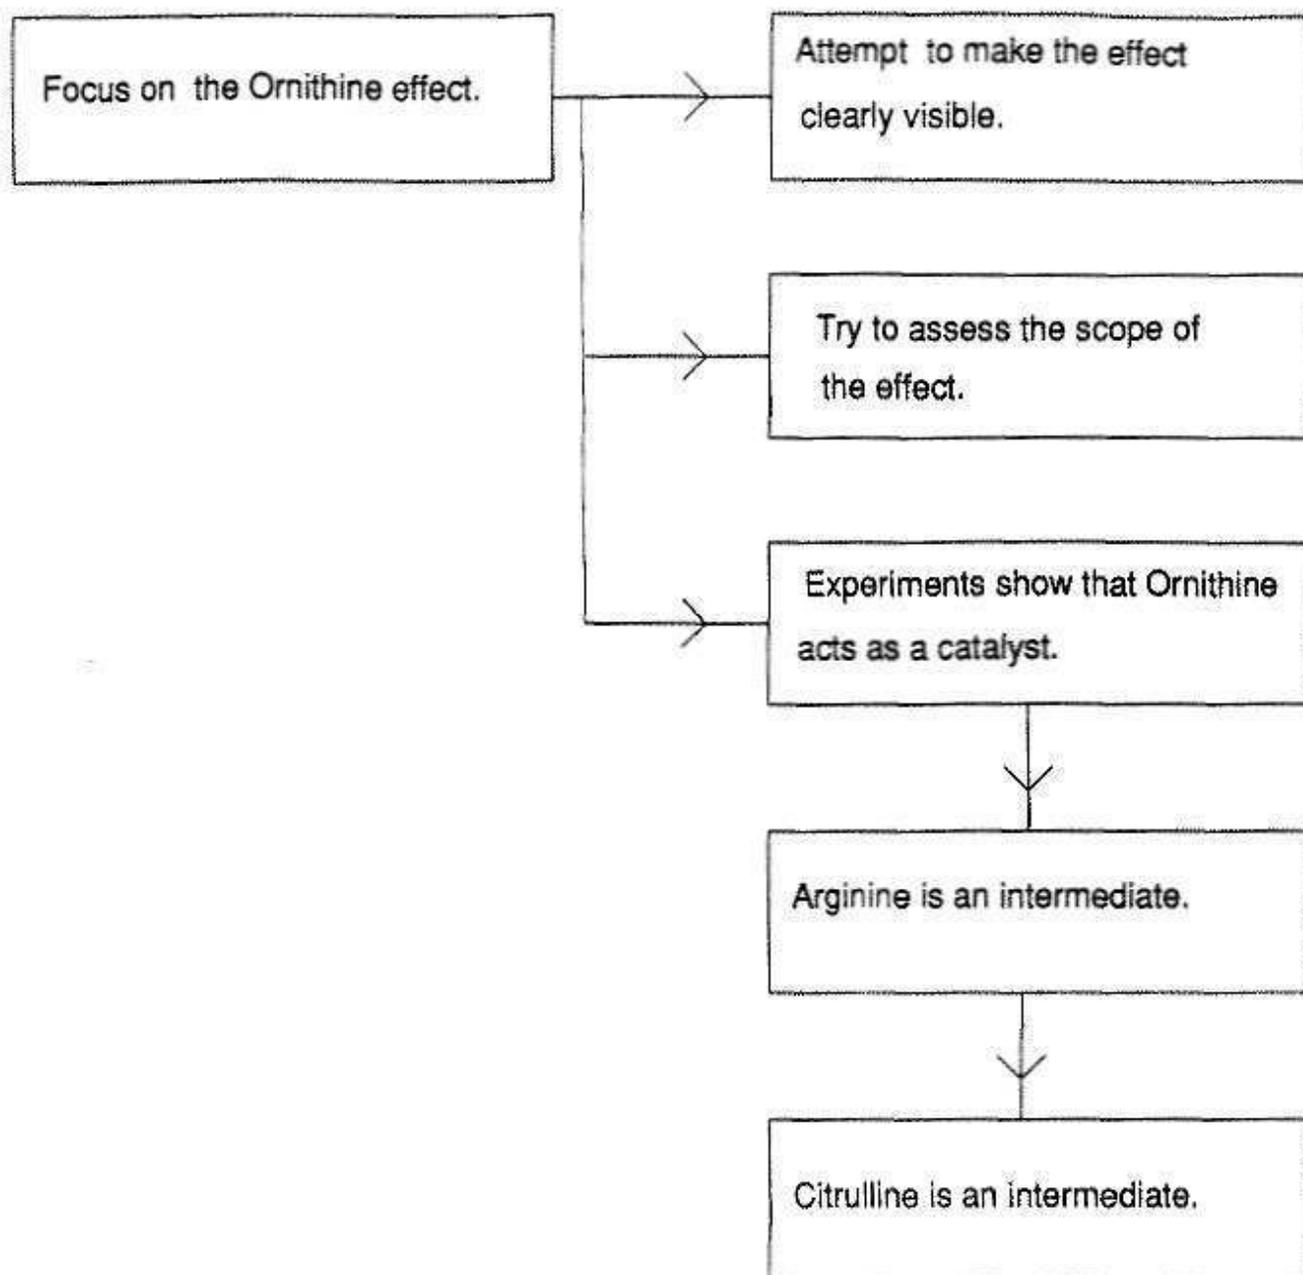


Figure 4-2: Discovery of Ornithine Cycle

4.1.5. Discovery of the Reaction Path

At this stage, 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 further guidance, the number of possible reaction paths is large and the system is able to consider only very, incomplete process descriptions. These hypotheses are at a high level of abstraction, where all the details need not be specified. Some such 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.

Next, KEKADA decides 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.

4.1.5.1. 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 an 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 2-1*.

4.1.6. Summary of the results produced

In the research discussed in the previous section, KEKADA produced a number of research results. It showed that most amino acids do not produce urea in liver slices. It also showed ornithine enters a cyclic pathway in which ammonia is converted to urea. In the next section we analyze what strategies helped KEKADA to succeed on this problem.

4.1.7. Analysis

4.1.7.1. Prior hypotheses factor the problem space

KEKADA did not start from scratch, but it started with two previously known hypotheses about urea synthesis. These hypotheses kept KEKADA in the part of the problem space containing interesting phenomena. The hypothesis that amino acids and ammonia may be combining together to form urea led KEKADA to try combination experiments. Previous hypotheses led KEKADA to work on the part of the problem space containing amino-acids, ammonia as reactants and with liver tissue as the locus of the process.

4.1.7.2. Focus on a surprise factors the problem space

KEKADA's decision to focus its attention on a surprise led it to a part of the problem space that turned out to have some interesting phenomena. Ornithine effect was surprising to KEKADA because KEKADA did not have any knowledge that would allow it to assign a special role to Ornithine. Thus the surprising effect indicated a deficiency in its knowledge.

4.1.7.3. Effectiveness of gradual refinement strategy

We noted earlier that today's textbooks show a more detailed mechanism of urea synthesis than the one Krebs discovered. Krebs did not have cognitive resources to solve the problem in that much detail. For Krebs, like any other scientist, it was not possible to know by foresight to what degree of detail available cognitive and other resources would allow him to solve the problem. Thus we may look at stages in the development of the problem:

1. Ornithine effect
2. Ornithine as a catalyst
3. Arginine as intermediate
4. Citrulline as intermediate
-
- X. Detailed mechanism in today's textbook.
-

There are two benefits of gradual refinement, it allows the program to generate reasonable alternatives, and hence to make progress to the point permitted by its resources. For example, while examining the role of ornithine as a catalyst, it is possible for the program to generate reasonable alternatives about an intermediate, but it would not be possible for the program to have generators for complete process models.

Second, a scientist could only in hindsight know for sure how detailed a solution to the problem his

cognitive resources would permit it to discover. This calls for a strategy of gradual refinement with an agenda that allows opportunistic shift of attention.

4.1.7.4. Heuristics do not always work

Many contemporary scientists had an additional heuristic: synthetic reactions can only studied *in vivo* with intact cell structure. Such a heuristic would have factored the problem space in an effective and useful manner for many of the other research problems. However in the case of urea synthesis it eliminated the part of the problem space containing phenomena necessary for understanding the synthesis. It thus led Warburg and others to believe that it was not a wise idea to study urea synthesis using tissue slices. Krebs, a newcomer, was free of these biases. Heuristics often work, but may not always work.

4.1.8. Trace of KEKADA behavior on urea synthesis

We present here the log of the run of KEKADA on the urea synthesis problem. An asterisk (*) denotes repeated application of a set of heuristics. *Seq/* names the sequence of firings of heuristics that is enclosed in the following pair of dashed lines.

The fact that alanine produces very little urea is a surprising phenomenon on which the system is focusing attention.

Heuristics	Results
HG11	Amino-acids may produce urea with ammonia as intermediate. Amino-acids may combine with ammonia to produce urea.
HG5	Studies alanine to urea reaction, decides that intermediate exists.
HG2	Some stimulator might be missing from the tissue slice preparation.
HG3	The reason for surprise may be one of the sub-reactions: either in the reaction that produces ammonia from the amino acid or in the reaction that produces urea from ammonia.
HG1	Study scope of the phenomenon by varying the "conc" attribute of variable alanine over the class "conc-range" of typical values of conc.
HG1	Study scope of the phenomenon over the class 'amino-acid.'
HG1	Study scope of the phenomenon over the class 'carboxylic-acid.'
HG1	Study scope of the phenomenon over the class 'amines.'

HG9 Effect needs to be magnified by varying apparatus variables.
 Gets user-assigned priorities from the user (identical to the system-defined priorities.)
 HSC1,DM1,ES*,EP9,E1(*)
 Tries magnifying the reaction by varying the values of pH, place, and aerobicity. The attempts to magnify do not succeed.

begin[seq0]
 HSC1,DM1 Consider the hypothesis that an absence of a stimulator may be causing the surprise.
 EP6,DM2* Guess the stimulators which may be present: various substances involved in carbohydrate mechanism. Chooses lactate.
 EP3 Decides to carry out an experiment on lactate and alanine both together.
 ES* Set expectations for the experiment.
 E1, ES4 Ask user for output for this experiment. The experiment on alanine and lactate produces urea at the q-rate of $Q_{urea}=1$.
 HCM-2 Modify failed-effort slot in the applicable hypotheses. Modify the confidence in the hypothesis that lactate and alanine are both required to produce urea.

[End seq0]

[Repeat seq0 for various substances.]

HCM-3 Make inactive the existential hypothesis that there may be a *stimulator missing* after trying on 3 substances.
 HSC1 Evaluate the alternatives.
 DM1 Decide to consider the hypothesis that the cause of the process may be in one of the subprocesses.
 EP7, E1 Carry out experiments on the subprocesses under various conditions of pH, aerobicity and in various organs. Nothing turns up.

[seq1]

[Begin seq1]
 HSC1 Evaluates the alternatives.
 DM1 Decides to consider following hypotheses simultaneously :1. that surprising phenomenon

may have scope over the class of amino-acids. 2. that an amino-acid and ammonia may be acting together to form urea. 3.that amino-acids may produce urea with ammonia as intermediate.

EP1-3, DM2 Decides for an experiment on glycine and on ammonia. Decides for an experiment on both combined together.

ES*.E1 $Q_{urea} = 1$; for the amino-acid, $Q_{urea} = 3$ for ammonia and $Q_{urea} = 4$ for the combination;

ES1-3, E1, ES4 HCM1,4,5

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] $Q_{urea} = 9$ for ornithine + ammonia

PG1, PC1 Notices the ornithine effect and makes it the focus of attention. Creates following hypotheses.

HG4* Possibility that ornithine or ammonia is catalyst.

HG7 Creates a clue that mixed action of both the inputs.

HG4* Hypotheses about who donates what to the reaction.

HG6* Possibility of relation to similar reactions, for example, to arginine reaction

HG3 problem may be in one of the sub-reactions of the process

HG1* Possibility that the phenomenon may be common to a class of substances: namely amino-acids, amines, carboxylic-acids.

HG9 Magnify the effect by varying the apparatus variables.

HG5 Intermediate exists.

[seq2]

HSC1, DM1, EP9, E1(*)

Tries to magnify the surprising phenomenon by varying various apparatus variables.
Attempts fail.

[Begin seq2]

HSC1 Evaluates the alternatives.

DM1 Decides to study the scope of the phenomenon. Considers that the phenomenon may be common to amino-acids.

- EP1,DM2 Considers various amino-acids. Decides on leucine as the choice.
- EP1 Decides to run an experiment on the amino-acid leucine and on ammonia combined only.
- ES1-3, E1,HCM-5 Sets expectations for these experiments. Asks the 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]

- HCM-3 Removes the description that some amino acids might produce urea.

{seq3}

[Repeats {seq2} for various amines.]

- HCM-3 Removes the description that some amines might produce urea.

[Repeats {seq2} for various carboxylic acids.]

- HCM-3 Removes description that some carboxylic-acids might produce urea.

- HSC1-DM1 Chooses the possibility that ornithine is catalyst.

- EP5 Decides for an experiment to verify catalysis.

- ES*, E1 Carries out experiments to check catalysis. Low conc of ornithine + NH_4Cl for 4 hours for each molecule of ornithine, 25 molecules of urea are formed.

- HCM-8 Concludes that ornithine acts as a catalyst. Decides to detect the intermediate in the reaction.

- HCM-10 Creates candidates for intermediate. Balances the reactions. Counts the number of inputs. Evaluates the plausibility of intermediates. Chooses arginine. Further concludes that in the reaction from ornithine to arginine there exists intermediate. Ask user to carry out a literature survey.

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

- HCM-10 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.

4.1.9. Differences in KEKADA behavior and Krebs research

After the initial failure to produce much urea from amino acids on tissue slices, two of the experiments Krebs carried out were with ammonia in liver tissue-slices from starved and well-fed rats. A better rate from the slices from well-fed rats indicated that tissue-slice method worked. Here we find Krebs suspecting the effectiveness of the tissue-slice method and carrying out some experiments to test it. While KEKADA may switch to better apparatus while studying a phenomenon, it does not have the knowledge that can suggest experiments to test the effectiveness of the apparatus.

After observing the Ornithine effect, Krebs tried to make some changes in the saline solution to improve the effectiveness of the method. Thus Krebs seems to be using the magnification strategy, but below the level of abstraction of KEKADA.

In his research work on urea synthesis, Krebs carried out experiments on thymine. This shows that he also entertained a known hypothesis that pyrimidines might be precursors to urea. If KEKADA is also supplied with that hypothesis, it will also carry out the relevant experiments.

4.1.10. Sensitivity analysis of the problem statement

In this subsection, we will try to make some comments on how difficult it will be for KEKADA to work on a problem which is slightly different from the current formulation.

Hypotheses about the urea synthesis: Two hypotheses about urea synthesis stated in P1 were the dominant hypotheses in 1920s. A problem statement which has an additional hypothesis such as the one about pyrimidines being precursors to urea would result in KEKADA carrying out a few more experiments. A problem statement that has a large number of incorrect hypotheses about urea synthesis or the one that does not have any hypotheses at all would make the problem significantly more difficult than those faced by actual scientists.

One aspect of the statement which seem to be seriously affected by the post-hoc knowledge is the fact that *givens* in the problem include the arginine reaction, but not hundreds of other arbitrary reactions known at that time.

4.1.11. Summary

In this section, we saw how KEKADA can rediscover the ornithine cycle as the mechanism of urea synthesis. Prior knowledge of incomplete theories allow KEKADA to work in the problem space containing important phenomena. KEKADA's focusing attention on the surprising phenomenon result in its attending to a part of the problem space rich in interesting phenomena. Gradual refinement of the hypotheses using domain knowledge allows it to deepen its understanding of the ornithine effect.

4.2. Behavior in the biochemistry of glutamine metabolism

Krebs continued to work on amino acid metabolisms, after the discovery of the Ornithine cycle. This research produced a number of interesting results. It established that the deamination of amino acids occurs by an oxidative reaction in kidney, and not in liver as has been previously assumed. It produced data on deamination rates of various amino acids. Furthermore it showed that glutamic acid combines with ammonia producing glutamine, a substance that was not previously known to play any role in metabolisms. Discovery of the glutamine reaction opened a whole set of new questions in metabolic biochemistry. As we mentioned earlier, the basic data on Krebs' work on amino acids was supplied to us by Holmes (1986).

4.2.1. Background of the discovery

Before Krebs began his work on deamination of amino acids, experiments on deamination of amino acids had been carried out mainly using the perfusion method and diet studies. It was believed erroneously that the liver is the site of deamination. Furthermore, there were 3 possible pathways posited about the nature of the deamination reaction. The three theories argued that deamination could take place by an oxidative, reductive, or hydrolytic reaction. While a lot of studies had been carried out on deamination, none of them had revealed any special role for glutamic acid. Glutamine was not known to have a place in metabolic pathways.

4.2.2. Initial working memory of the program

While Krebs worked on the urea synthesis problem, he had carried out an experiment with ornithine on kidney tissue slices. The result was the formation of ammonia. As previous theories posited no role for kidney in deamination reactions, this was an unexpected result. Krebs carried out a long series of fruitful experiments after he started paying attention to this result.

We saw in the last section that KEKADA also decides to carry out an experiment with ornithine on kidney tissue slices. As KEKADA's problem-evaluation heuristics are not very sophisticated, we supplied KEKADA with the result that the experiment does not result in the formation of ammonia or urea. Thus KEKADA was not surprised and did not divert its attention from the ornithine effect. In the present section, we run the program with its attention focused on the surprising reaction in which ornithine produces ammonia in kidney. Background knowledge about amino acids provided to the program is almost the same as that which was provided for its run on the urea synthesis problem.

The background knowledge includes the following: Structural formulae of various substances, set of values of pH, ordered-lists of amino-acids, amines, and carboxylic acids, stimulators and inhibitors associated with processes. The system starts with its focus of attention on the following phenomenon: that ornithine produces ammonia in the kidney. The system also knows that there are 3 previously postulated reaction pathways about the deamination reaction in which an amino acid loses its amino group.

4.2.3. Overview of KEKADA's behavior on amino acid metabolisms

KEKADA's behavior on the research problem can be divided into four different stages: characterization of the ornithine-in-kidney effect, study of the deamination reaction, study of the glutamic acid effect, and the discovery of the glutamine reaction.

Characterization of ornithine-in-kidney effect: After observing that ornithine produces ammonia in kidney, KEKADA tries to characterize the surprising phenomenon and then realizes that it is a part of the general deamination reaction.

Study of the deamination reaction: KEKADA further confirms the oxidation scheme of deamination and gathers more data on the deamination reaction. In the course of gathering this data, it comes across an unusual reaction from glutamic acid.

Study of glutamic acid effect: While attempting to characterize the glutamic acid reaction, KEKADA notices that arsenite has a further surprising effect on the glutamine reaction.

Discovery of glutamine reaction: KEKADA conjectures the possibility of glutamic acid combining with ammonia. It further verifies experimentally that glutamic acid does indeed produce glutamine combining with ammonia.

We will now describe KEKADA's behavior in more detail. The description is a paraphrase of the trace described in a later section.

4.2.4. Characterization of Ornithine-In-kidney Effect

The system begins with its focus of attention on the puzzling observation that ornithine produces ammonia in kidney.

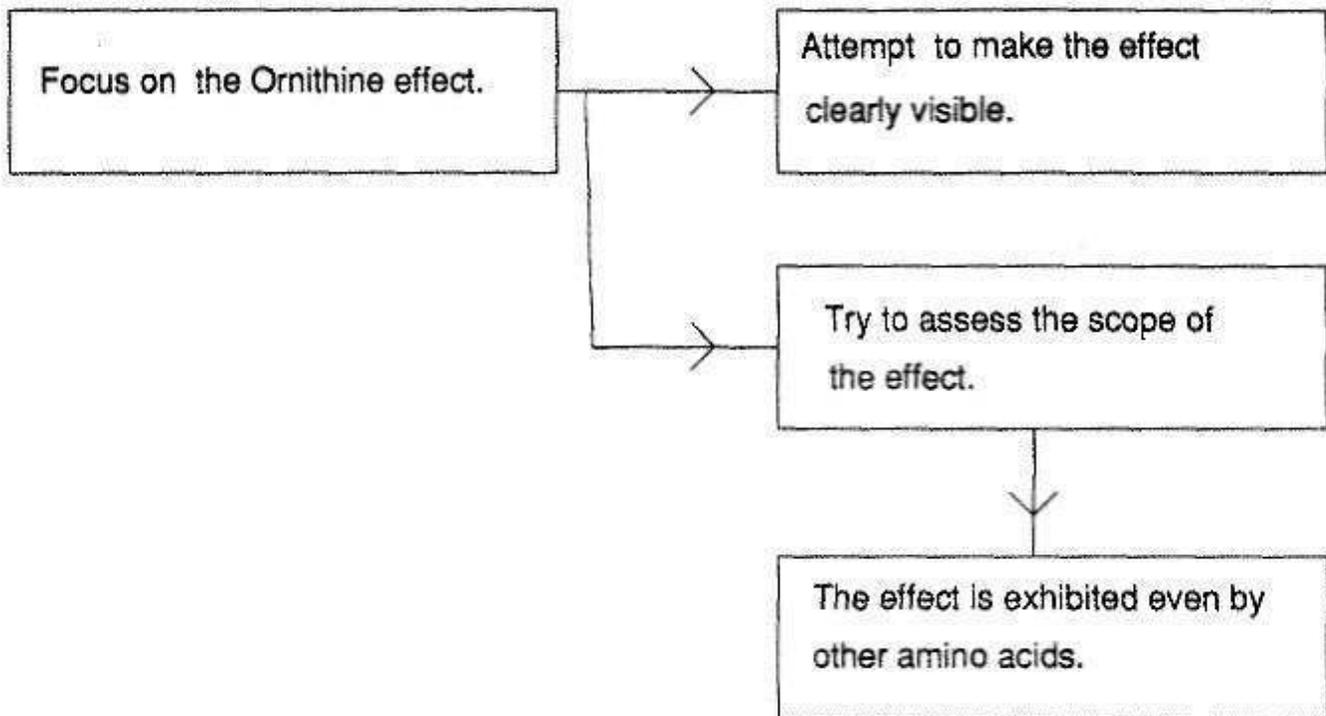


Figure 4-3: Ornithine Effect in Kidney

In response to this surprise, the hypothesis and strategy generators suggest a number of hypotheses and strategies.

1. The reaction could be an example of the deamination reaction. Three reactions have previously been hypothesized about deamination: It could be an oxidative, hydrolytic, or a reductive reaction.

2. Ornithine might be donating one of its amino groups to ammonia.

3. The effect may be common to a class of substances. (amino acids, amines, carboxylic acids)

Another strategy would be to examine the effects of varying the concentration of ornithine.

4. Magnify the effect by varying the apparatus variables.

KEKADA first chooses the magnification strategy and tries to magnify the phenomenon by varying the apparatus variables. However these attempts do not succeed in magnifying the effect.

Next KEKADA tries to assess the scope of the surprising phenomenon. It finds that other amino acids can also produce ammonia in kidney.

Study of the Deamination reaction

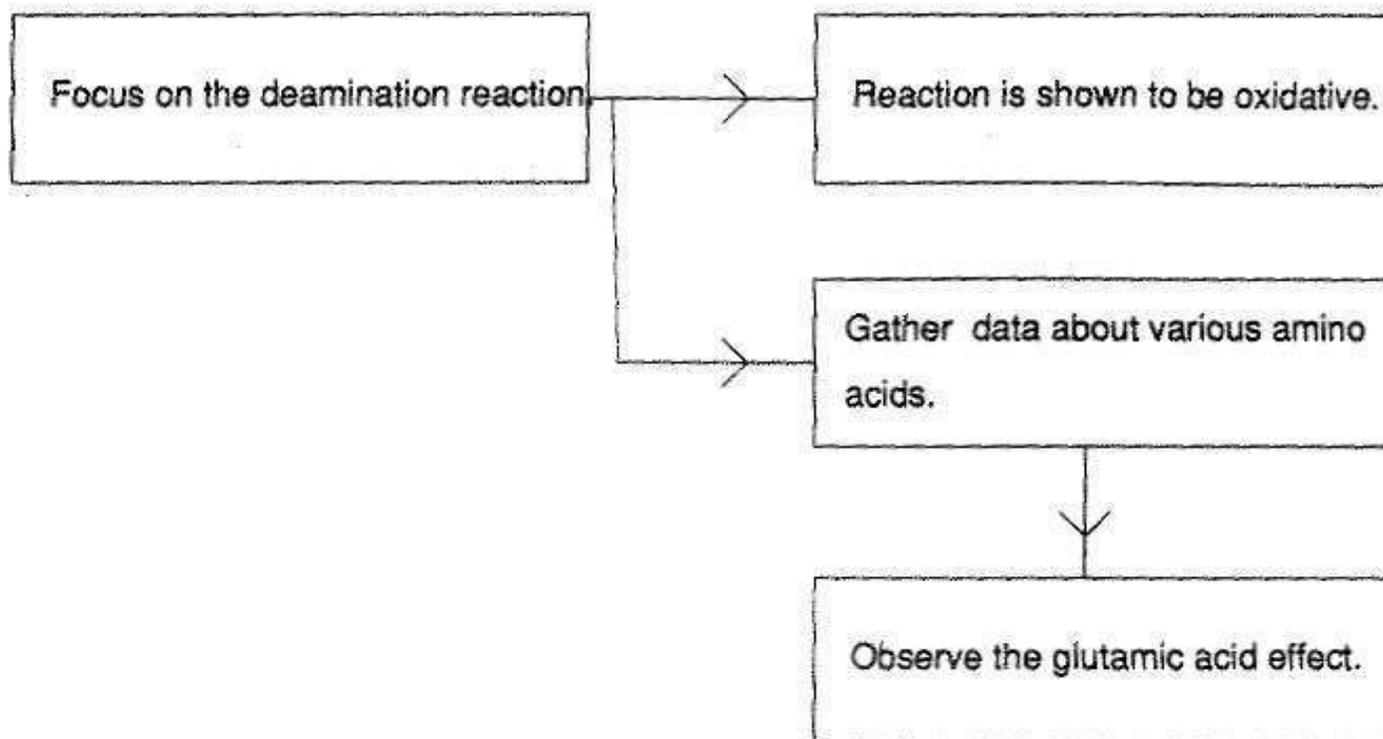


Figure 4-4: Study of deamination reaction

At this stage the program recognizes that the phenomenon under consideration is quite similar to deamination reaction and activates known hypotheses about the deamination reaction. These include

- 1) hydrolysis: $R-CH-COOH + H_2O = R-CHOH-COOH + NH_3$
- 2) reductive: $R-CH(NH_2)-COOH + H_2 = R-CH_2COOH + NH_3$
- 3) oxidative: $2 R-CH(NH_2)-COOH + O_2 = 2 RCOCOOH + 2 NH_3$

Thus generalizing the phenomenon allowed KEKADA to use previous knowledge about a particular class.

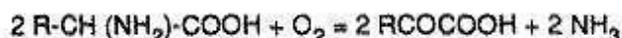
KEKADA also considers the following strategies.

1. Observed reaction might be in some way related to the urea synthesis reaction.
2. One could gather more data by carrying out experiments on various amino-acids.

3. The effect may be common to a even bigger class of amines or carboxylic acids.

4. One may try to magnify the effect by varying the apparatus variables.

KEKADA first decides to verify the oxidative hypothesis. To carry out an experiment, it needs to choose one of the amino acids. It decides to carry out the experiment on alanine for a number of reasons. Alanine is both reactive and cheap, and is thus particularly appropriate. Here we see one more advantage of the strategy of assessing the generality of a surprise. If the phenomenon turns out to be general, the experimenter has more choice in choosing a variable in the the phenomenon. To verify the oxidative reaction, KEKADA carries out experiments on alanine and on oxygen together. It is observed that oxygen is consumed at the q-rate of 5 and that ammonia and the keto-acid are produced at the q-rate of 10. As this is consistent with the chemistry of the oxidative reaction, KEKADA concludes that deamination must be occurring by an oxidative reaction:



Next KEKADA decides to gather more data on the deamination of amino-acids. A research strategy to produce data on an important problem is likely to produce information which is of interest to other scientists. If the scientist also looks for the unexpected, he may get a further lead to producing more interesting phenomena.

Thus KEKADA tries out experiments on various amino-acids. The last of the amino acids is glutamic-acid. It is found that the q-rate of oxygen consumption is 5, and the q-rate of production of keto-acid is 0. As the rate of production of keto-acid is lower than expected, this is noted as a surprise.

Study of Glutamic Acid Effect

Immediately KEKADA makes the glutamic acid reaction the focus of attention.

Presently it creates a number of hypotheses and strategies about the puzzling phenomenon.

1. There might be a missing inhibitor.
2. Either the glutamic acid or oxygen might be acting as a catalyst.
3. Examine if all the reactants are necessary for the reaction to occur.
4. Glutamic acid may be donating amino group to ammonia and COOH group to keto acid.

Glutamic acid may be donating carbonyl group to keto acid.

5. The phenomenon may be common to a class-- carboxylic acids, amino acids, amino acids, etc.

Study the effects of varying the concentrations of the reactants.

6. Magnify the effect by varying the apparatus variables.

First KEKADA attempts to magnify the phenomenon by varying parameters such as pH, place, and aerobicity. Attempts to magnify fail.

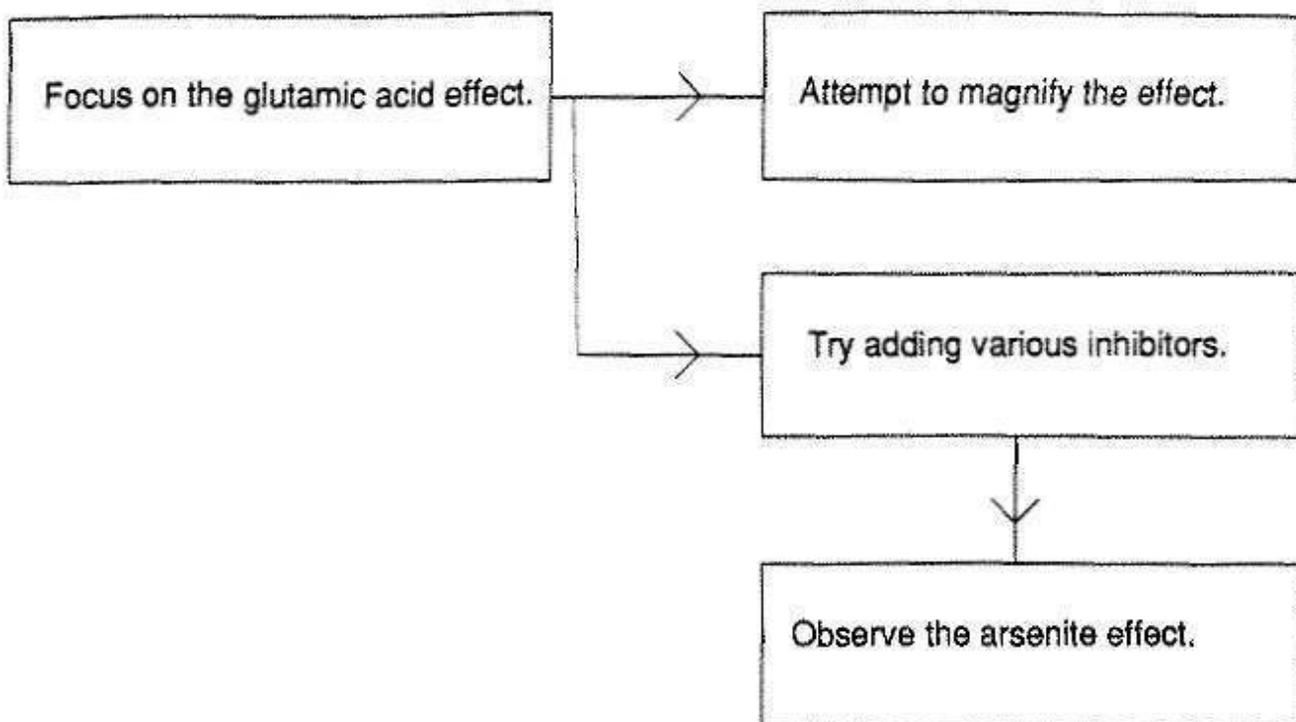


Figure 4-5: Study of glutamic acid effect

Next, it tries out a number of standard inhibitors, hoping that inhibitor will selectively block the reaction consuming the keto-acid. The initial inhibitors do not succeed in blocking the reaction. The last of the inhibitors is arsenite.

When a reaction is carried out with arsenite, and the rate of NH_3 increases as well as the keto-acid. No incremental increase in NH_3 production was expected by addition of arsenite. Thus the fact that addition of arsenite is capable of increasing the production of NH_3 is noted as a surprise, KEKADA focuses its attention on it.

Discovery of Glutamine Reaction

KEKADA generates following hypotheses and strategies about the surprising phenomenon.

1. Glutamic acid, arsenite, or oxygen may be a catalyst.
2. Try to find out if the phenomenon is limited only to arsenite, or if it is also exhibited by other metabolic-inhibitors.
3. Try varying the concentration of arsenite to see its effects.
4. Glutamic acid may be donating COOH or carbonyl group to keto-acid. Glutamic acid may be donating an amino-group to ammonia.
5. Deamination of glutamic acid in the absence of arsenite is a related reaction.

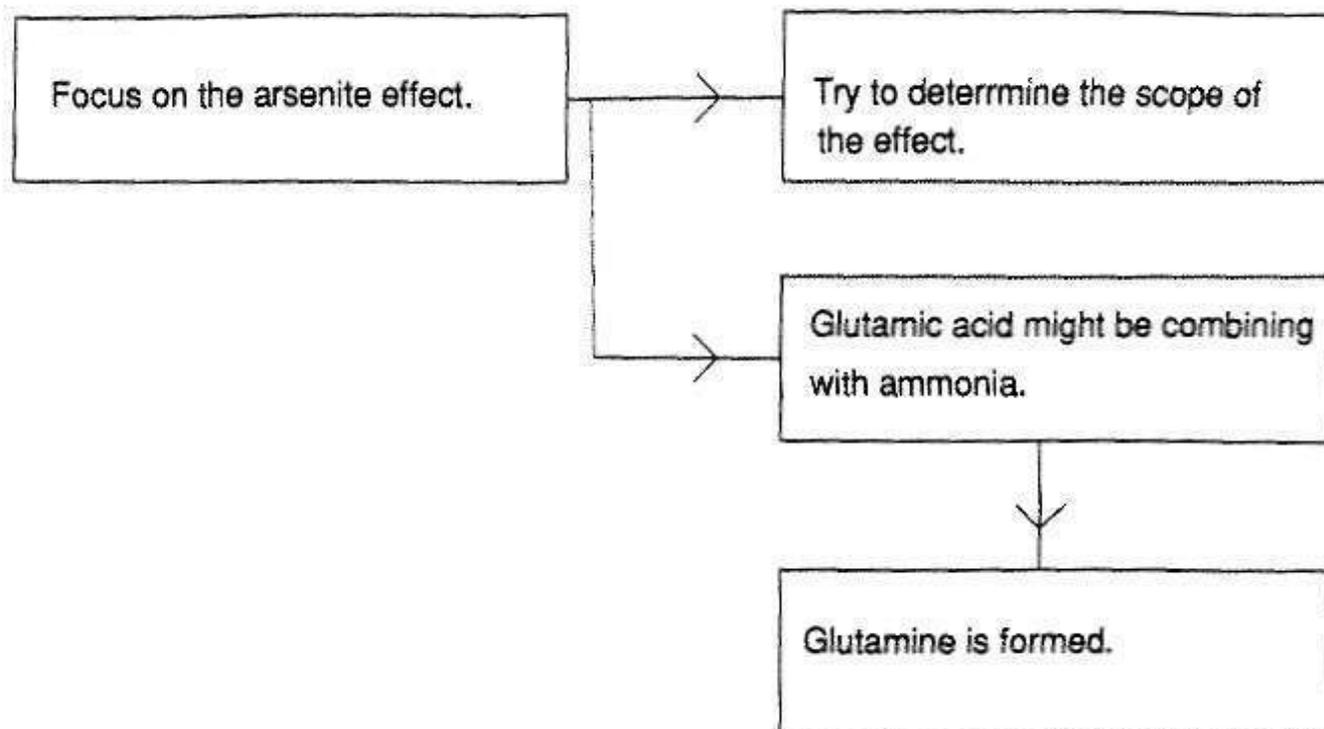


Figure 4-6: Discovery of Glutamine reaction

6. Examine if all the three reagents are required to exhibit the effect.

7. The phenomenon may be common to a class: carboxylic acids, amino acids, or amines.

Another strategy would be to vary the concentration of glutamic acid.

8. Ammonia may be reacting with one of the reactants in a side-reaction.

9. Magnify the effect by varying the apparatus variables.

First, KEKADA chooses the hypothesis that the phenomenon will have some scope over the class of carboxylic-acids.

However after the negative results for aspartic acid and other carboxylic acids, KEKADA reduces the confidence in the hypothesis that the phenomenon may also be exhibited by other carboxylic acids. Now KEKADA decides to consider the hypothesis that ammonia may be reacting with one of the other reactants in a side-reaction.

When KEKADA carries out the reaction with glutamic-acid and ammonia, glutamine is produced. At this stage KEKADA has discovered the important glutamine reaction.

4.2.5. Results produced by the program

This research produced a number of interesting results. It established that deamination of amino acids occurs by an oxidative reaction in kidney, and not in liver as has been previously assumed. It produced data on deamination rates of various amino acids. Furthermore it showed that glutamic acid combines with ammonia producing glutamine, a substance that was previously not known to play any role in metabolisms. Discovery of the glutamine reaction opened a whole set of new questions in metabolic biochemistry.

4.2.6. Analysis

4.2.6.1. Focusing on a surprise as a search-control strategy

Throughout the previous log, we find that decisions to focus on a surprise led to revealing a lot of interesting phenomenon. For example, consider the first experiment on glutamic acid which KEKADA finds surprising. Since KEKADA cannot explain an abnormal reaction of glutamic acid, the surprising reaction of glutamic acid indicated a dimension along which KEKADA lacked knowledge. As KEKADA attempted to understand the reaction better, it acquired some of the missing knowledge. Furthermore we saw that a number of hypothesis and strategy generators were effective in characterizing particular surprising phenomena.

4.2.6.2. Determination of scope

After observing that the ornithine-in-kidney effect is general to the class of amino-acids, KEKADA carries out experiments with alanine, a cheaper and more reactive agent. We also saw how prior knowledge about amino-acids helped in understanding the effect. Two of the many advantages of generalizing a phenomenon are: it improves our experimental control by giving us much wider choice, and by allowing us to use knowledge about a wider class, it increases the chances of solving the research problem.

4.2.6.3. Gather Data Strategy

We saw that while KEKADA gathered more data about the deamination reaction by running experiments on various amino acids, it produced important data on deamination of amino-acids and also came across the surprising glutamic acid effect. This shows the usefulness of the gather-data strategy. When one finds a novel and unusual phenomenon, even systematic collection of data about the phenomenon can be of great interest to the scientific community. Furthermore while collecting such data, surprises can turn up.

4.2.7. Domain-specific strategies

We also saw how domain-specific knowledge suggested the possibility of glutamic acid combining with ammonia.

4.2.8. Trace of KEKADA behavior

The system is focusing attention on the puzzling observation that that ornithine produces ammonia in kidney.

Heuristics	Results
HG11	The reaction could be an example of the deamination reaction. Three reactions have previously been hypothesized about deamination: an oxidation, an hydrolytic, and a reductive reaction.
HG4	Ornithine might be donating an amino group to ammonia;
HG1*	Creates various hypotheses including that " the effect may be common to a class of substances" (amino acids, amines, carboxylic acids). Consider the effects of varying the concentration of ornithine.
HG9	Magnifies the effect by varying the apparatus variables.
HSC1, DM1, EP9, E1(*)	Chooses the magnification hypothesis and carry out experiments varying the apparatus variables. Attempts to magnify the phenomenon fail.
{seq0}	
.....	
HSC1, DM1	Chooses the hypothesis that the effect may be common to some amino-acids.

EP1, DM2, ES*, E1

Decides to carry out an experiment on alanine in kidney. (Ammonia is produced with q-rate of 10)

HCM-5 Increases the confidence in the hypothesis that the phenomenon has some scope over the class of amino-acids.

Repeat [seq0] for various amino-acids.

HCM-7 Generalizes the description of the phenomenon to: amino acids produce surprisingly large amount of ammonia in kidney.

HG4 Amino-acids may be donating an amino group to ammonia and there may be some not yet known reactants involved in the reaction.

HG11 This might be the deamination reaction. Activates 3 hypotheses about it:

1)hydrolysis: $R-CH-COOH + H_2 O = R-CHOH - COOH + NH_3$

2)reductive: $R-CH(NH_2)-COOH + H_2 = R-CH_2 COOH + NH_3$

3)oxidative: $R-CH (NH_2)-COOH + O_2 = RCOCOOH + NH_3$

HG1 Tries to determine the effect of varying the concentration of amino-acid.

HG6 Urea synthesis is a related reaction.

HG12 Tries to gather more data by carrying out experiments on various amino-acids.

HG1 The effect may be common to a even bigger class of amines or carboxylic acids.

HG9 Decides to magnify the effect by varying the apparatus variables.

HSC1, DM1 Decides to verify the oxidative reaction.

EP2 Decides to carry out experiments on alanine and oxygen together.

E1, HCM-6 Alanine and oxygen together produce keto-acid with q-rate of 10 and ammonia is also produced at q-rate of 10. Oxygen is consumed at the q-rate of 5. Concludes that the oxidative hypothesis of deamination is correct.

HSC1, DM1 Decides to gather data on deamination of amino-acids.

EP10, DM2, ES*, E1

Tries out the reaction on various amino-acids. Last one being glutamic-acid (Q-rate of oxygen consumption 5, q-rate of production of keto-acid is 0)

PG1, PC1 Makes the glutamic acid reaction the focus of attention.

HG2 There might be a missing inhibitor.

- HG4 Either the glutamic acid or oxygen might be acting as a catalyst.
- HG7 Examines if all the reactants are necessary for the reaction to occur.
- HG4* Glutamic acid may be donating an amino group to ammonia and COOH group to keto acid. glutamic acid may be donating a carbonyl group to the keto acid.
- HG1* The phenomenon may be common to a class- carboxylic acids, amino acids, amino acids, etc. Study the effects of varying the concentration of the reactants.
- HG9 Magnify the effect by varying the apparatus variables.
- HSC1, DM1, E1(*)
Chooses the magnification hypothesis, and carry out a number of experiments by varying the pH, place, aerobicity. Attempts to magnify fail.
- [seq2]

- HSC1, DM1, EP6, DM2
Choose the missing inhibitor hypothesis.
- ES*, E1 No incremental action of HCN

- Repeats [seq2] for various inhibitors.
- Carries out the reaction with arsenite, and q-rate of NH_3 increases as well as the keto-acid. (No incremental increase in NH_3 production was expected by addition of arsenite.) $Q_{\text{NH}_3} = 10$; $Q_{\text{keto-acid}} = 10$; $Q_{\text{O}_2} = 5$.
- PG1, PC1 Notes this as a surprise.
- HG4* Glutamic acid, arsenite, or oxygen may be a catalyst.
- HG4 Try to find out if the phenomenon is limited only to arsenite, or it holds for larger number of members of the class of metabolic-inhibitors.
- HG4 Try varying the concentration of arsenite to see its effects.
- HG5* Glutamic acid may be donating COOH or the carbonyl group to keto-acid. Glutamic acid may be donating amino-group to ammonia.
- HG6 Formation of ammonia from glutamic acid (in the absence of arsenite) is a related reaction.
- HG7 The effect may be due to all the three reagents.
- HG1* The phenomenon may be common to a class- carboxylic acids, amino acids, amines...
Study the effect of varying concentration of glutamic acid.

- HG10 Ammonia may be reacting with one of the reagents to form a side-reaction.
- HG9 Magnify the effect by varying the apparatus variables.
- [seq3]
-
- HSC1, DM1 Chooses the hypothesis that the phenomenon will have some scope over the class of carboxylic-acids.
- EP1, E1, HCM-5 Scope experiment gives negative results for aspartic acid.
-
- HCM-3 After repeating [seq3] for various carboxylic acids, the hypothesis that the phenomenon may have scope over carboxylic acids is removed.
- HSC1, DM1 Chooses the hypothesis that ammonia may be reacting with one of the other reactants in side-reaction.
- EP11 Carries out an reaction with glutamic-acid and ammonia. This results in the production of glutamine.

4.2.9. Sensitivity analysis of the Problem Statement

Problem statement P2 include previous postulated hypotheses that amino acids might be producing ammonia in liver by an oxidative, hydrolytic or reductive reaction. It would not be historically plausible to include many more hypotheses in the problem statement.

We are also assuming that KEKADA needs to carry out an experiment only once to get a reliable result. In actual history, Krebs sometimes needed to repeat the experiment a number of times to get a reliable result. The reason was that washing procedures are not uniform and unknown amount of residual enzymes and other substances might be present. We assume that this variation is below the level of abstraction of the program.

4.2.10. Summary

The program begins with its attention on the puzzling phenomenon that ornithine produces urea in kidney. By the time the program run is interrupted, it has produced interesting data on amino-acid metabolisms and it has produced the important fact that glutamic acid combines with ammonia producing glutamine.

Its strategy to assess scope of an interesting phenomenon and its focusing attention on the

surprising phenomenon result in its attending to the part of the problem space containing interesting phenomena. Heuristic search using its domain-specific knowledge and its ability to carry out experiments allow it to understand further various surprising phenomena it encounters.

4.3. Behavior in the ether reactions

Williamson's discovery that alcohol contains two ethyl groups led to better understanding of etherification processes and thus resolved a long-standing problem about the structure of alcohol. The basic source of historical data used here is Leicester 1952.

4.3.1. Background of the discovery

In 1850 before Williamson started work on alcohols, there were a number of molecular structures proposed for common alcohol. Dumas had proposed that alcohol (C_2H_6O) had an etherin group. Laurent (1846) had proposed that alcohol had an ethyl group. It was against this background that Williamson started his work. He started carrying out experiments with the aim to produce alcohols. He decided to carry out an experiment with common alcohol, potash, and ethyl iodide. He expected that potash would replace hydrogen from common alcohol and furthermore ethyl group would replace K^+ . Thus he expected that a higher order alcohol would be formed. However to his surprise the substance formed was ether. Williamson carried out a fruitful series of experiments while studying this effect.

4.3.2. Initial state

We run KEKADA with the following background knowledge about chemistry: names of various substances, and their structural formulae. The fact that common alcohol together with potash and ethyl iodide produces common ether, is the puzzling phenomenon on which the system starts focusing attention.

4.3.3. Discovery of alcohol structure

An experiment is carried out with 3 reactants: common alcohol, K^+ , and ethyl iodide. This experiment produces ether.

Common alcohol + K^+ + $C_2H_5\cdot OH$ produces $C_4H_{10}O$

In response to the surprise the hypothesis and strategy proposers suggest the following hypotheses and strategies.

1. Ethyl iodide may be donating an ethyl group to ether.
2. Common alcohol may be donating an ethyl group to ether.
3. Etherification in the presence of sulphuric acid may be related to the observed phenomenon.
4. Study the effect of varying the concentration of the input reactants.
5. The phenomenon might be common to other alcohols, to ethyl halides, or to alkyl iodides.
6. Try to find if all the reactants are necessary to produce the phenomenon.
7. The effect should be magnified by varying apparatus variables.
8. Any of the input reactants could be acting as a catalyst.

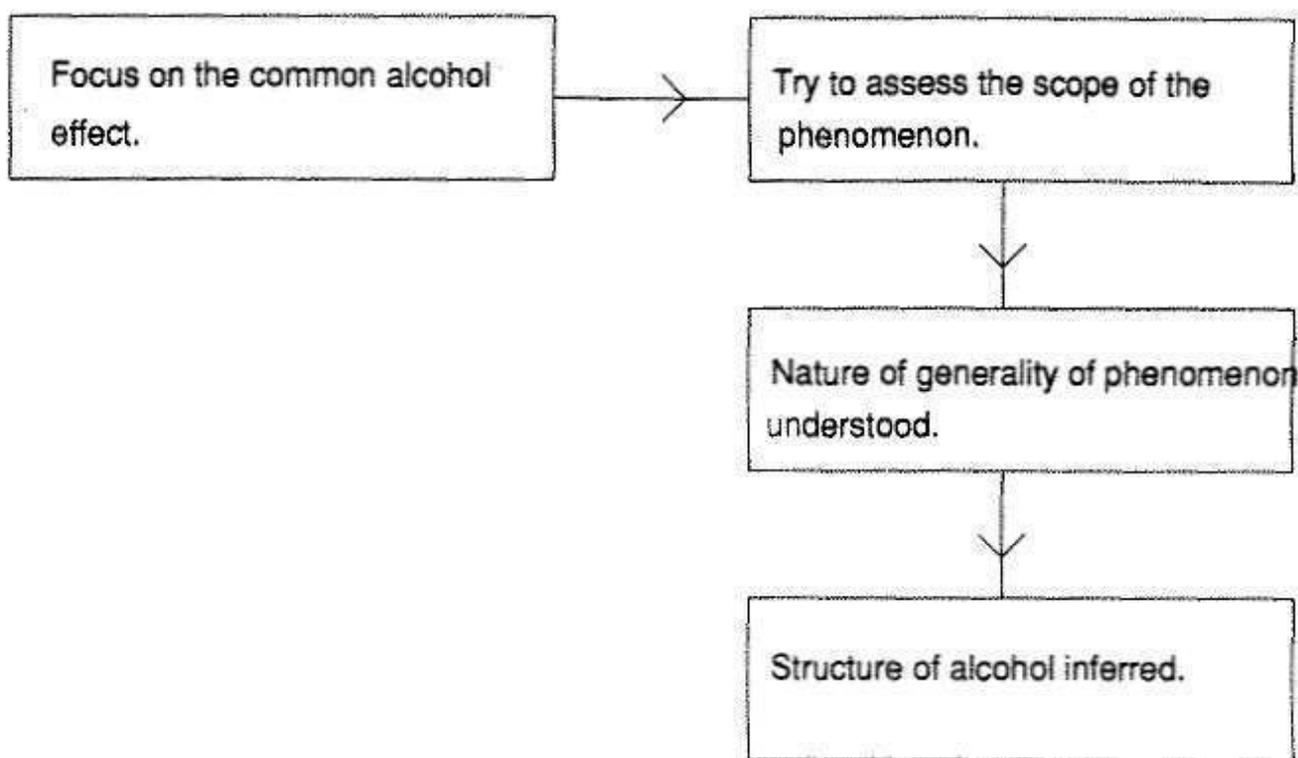
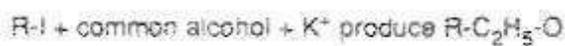


Figure 4-7: Discovery of Alcohol Structure

KEKADA next decides to assess the scope of the phenomenon over the class alkyl iodide. The phenomenon may not be limited to ethyl iodide, but may also be exhibited by other members of the class alkyl iodide. When experiments are carried out on other alkyl halides, it is found that in presence of potash and common alcohol they produce new ethers. In particular methyl iodide with common alcohol and potash produces methyl ethyl ether and in general

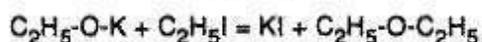


From the generalized description, it is obvious that common alcohol must contain the ethyl group.

Only the structural formula of alcohol containing the ethyl groups is consistent with the above

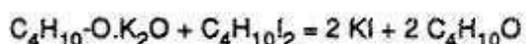
phenomenon. In this case the strategy of generalizing the phenomenon to an abstract level reduced the space of possible alcohol structures to one. Generalizing a surprising phenomenon would often eliminate many variables in a specific form of the surprise, and thus reduces the size of the relevant problem space. This can lead to understanding the effect better.

Let us compare KEKADA reasoning with the reasoning Williamson followed. Having faced with the same surprising phenomenon as KEKADA, he tried to see if the surprise can be explained by the two possible structural hypotheses about alcohol. If alcohol is $C_2H_5O H$ and the potassium compound is $C_2H_5 OK$, then the ethyl group from ethyl iodide may replace K to form ether. Thus



Another theory would explain the phenomenon by assuming that the potassium compound contains ether and potash, which separate during the action of ethyl iodide.

Thus



While Williamson considered the first explanation more convincing, he decided to carry out a *discriminatory* experiment by carrying out the reaction between common alcohol, potash, and methyl iodide. If the second explanation was right, one should obtain a mixture of common ether and an oxide of methyl. If the first explanation was right, one should obtain C_3H_8O . The experiment thus discriminated between the two explanations. KEKADA at present does not have the ability to carry out such *discriminatory* experiments.

4.3.4. Trace of program behavior

Heuristics	Results
	Experiment is carried out with 3 reactants common alcohol, K +, and ethyl iodide. Expectation set is that this should produce the higher homologue of alcohol. Instead, the output is ether. KEKADA is run with its focus of attention on this surprising reaction.
HG4	Ethyl iodide may be donating ethyl group to ether.
HG4	Common alcohol may be donating ethyl group to ether.
HG6	Etherification in the presence of sulphuric acid may be related.
HG1*	Study the effect of varying the concentration of the input reactants.
HG1*	The phenomenon would be common to other alcohols, to ethyl halides, to alkyl iodides.
HG7	Try to find if all the reactants are necessary to produce the phenomenon.

- HG9 *The effect should be magnified by varying apparatus variables.*
- HG4 *Any of the input reactants could be acting as a catalyst.*
-
- [begin seq1]
- HSC1.DM1 *Evaluate the alternatives. Decide to assess the scope of the phenomenon over the class alkyl iodide. (Phenomenon may not be limited to ethyl iodide, but may be exhibited by other members of the class alkyl iodide)*
- EP1, E1 *Chooses methyl iodide. The experiment on common alcohol, K⁺ and methyl iodide produces ethyl methyl ether.*
- HCM-5 *Increase the confidence in the scope-hypothesis.*
- [End seq1]
-
- Repeat [seq1] for various alkyl iodides.
- HCM-7 *Generalize the phenomenon to the following description: R-I + C₂H₅-OH + K⁺ produce R-C₂H₅-O*
- HG8 *Conclude that the correct structural formula of alcohol contain C₂H₅, and is C₂H₅OH*

4.3.5. Results produced by KEKADA

In the research described in the previous section, KEKADA establishes the generality of the etherification reaction. Furthermore it concludes that alcohol contains an ethyl group and has the formula C₂H₅OH.

4.3.6. Sensitivity analysis of the problem statement

Our problem statement which is based on Williamson's own account states two hypotheses about the alcohol structure. However there was at least one more hypothesis by Berzelius that ether and alcohol were oxides of different radicals, to which Williamson does not refer in his paper. If our problem statement were to have this additional hypothesis, KEKADA would still be productive in producing new facts and but would not be able to reduce the number of possible hypotheses about the structure of alcohol to one.

4.3.7. Summary

While studying a puzzling phenomenon the program infers the correct structural formula of alcohol. The strategy of assessing the scope of a phenomenon and generalizing the description of the phenomenon essentially reduces the space of possible of alcohol structures to one.

4.4. Behavior in magneto-electricity

Faraday's discovery of the induction of electricity from magnetism is a classic discovery in the history of science. In its practical form, it led to large-scale production of electricity by mechanical means. Faraday's diaries (Martin, 1932) and Ryan Tweney's analysis (1985) of these diaries are our basic sources of information. We used Faraday's correspondence (Williams, 1971) and Williams' analysis (1965) as supplementary sources.

4.4.1. Background of the discovery

In 1820, Oersted reported that a wire carrying electric current can induce magnetism. This caused an immediate sensation. It appeared to support the hypothesis that a single force exhibits itself in the form of electrical and magnetic forces. If this hypothesis was true, it should be possible to induce electricity from magnetism in a way symmetric to Oersted's experiment. Before 1830, different researchers had tried to demonstrate such induction of electricity and failed. However, the research of electricity and magnetism produced a number of new phenomena and theories in the 1820s. One of the interesting phenomena was reported by Francois Arago in France. He reported that the rotation of a copper disc in the presence of earth's magnetic field made it magnetic. This was a surprise, as simple motion seemed to be inducing magnetism in a non-magnetic substance. Further work on Arago's effect demonstrates that magnetism in motion induces magnetism in copper and other metals previously not known to be magnetic.

In United States, Joseph Henry reported on how powerful magnetic power could be generated from electro-magnets. Faraday's friend Gerritt Moll informed Faraday of Henry's experiments on June 7th 1831. It was against this background that Faraday resumed his experimentation on August 29, 1831 and carried out fruitful investigations on magneto-electricity.

4.4.2. Initial Working Memory

KEKADA begins with certain background knowledge about magnets and about electricity. This knowledge includes knowledge about instruments and about various available methods of measurement of electric current. The program also knows about the effect reported by Arago. The system begins by focusing on the following phenomenon: On switching on electric current in a coil around a cylinder, a temporary current is produced in another nearby coil.

4.4.3. Overview of KEKADA's research on magneto-electricity

KEKADA's research on magneto-electricity can be divided in three stages: observation of make-break effect, characterization of make-break effect, and observation of induction from magnetic motion.

Observation of Make-Break effect: KEKADA explores if the transient current also reappears when the electric current is switched off from the coil; and finds that it does.

Characterization of Make-Break Effect: Further KEKADA attempts to understand the phenomenon better by making changes in the apparatus. It also observes that making or breaking of magnetic circuit can also produce current in the adjoining coil.

Observation of induction from magnetic motion: KEKADA also entertains the possibility that phenomena it is observing might be related to Arago's effect. One of the experiments it tries out is moving a magnet towards a coil. In this experiment, it observes that motion of the magnet alone is capable of producing electrical current in a coil.

4.4.3.1. Observation of Make-Break Effect

KEKADA begins with its focus of attention on a puzzling phenomenon. In the apparatus shown in figure 4-8, on closing the electric circuit in coil A around the cylinder, current is produced in the coil B. Let me refer to this as the *make* effect.

Hypotheses and strategy generators in KEKADA suggest the following hypotheses and strategies about this surprising phenomenon:

1. Try to magnify the effect by varying the associated apparatus variables. By using slightly different apparatus or methods of measurement, it might be possible to make the effect clearly visible.
2. Try to see if the effect is also evident for various values of steady-state electric current and also on breaking the electric circuit on side A. Effect may not be restricted to only switching on of the electric current. It may also be valid for steady state values of the current and it may perhaps be evident when the circuit is switched off.

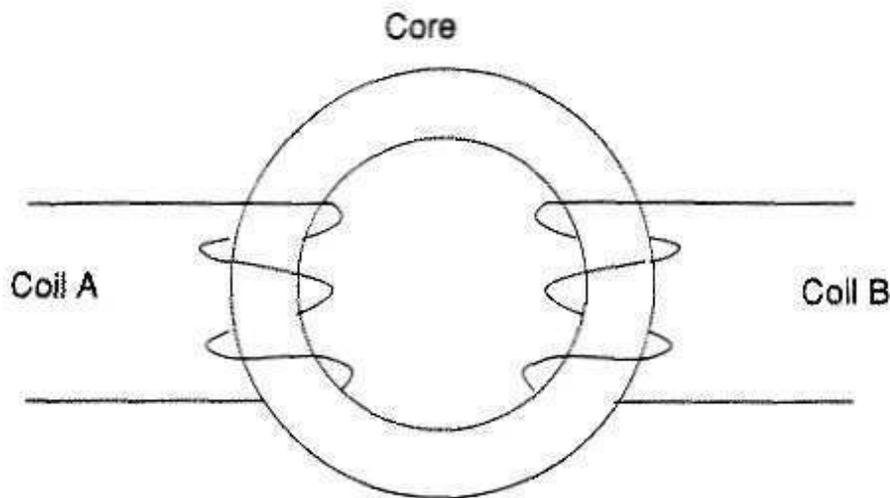


Figure 4-8: Study of induction of electricity

3. Effect may be common to other known forms of electricity. For example, would static electricity also display such an effect?

4. The surprising phenomenon might somehow be related to the Arago's effect. However as the current phenomenon and the Arago's effect have only one variable magnetism in common, KEKADA gives a low priority to attending to this hypothesis.

5. Another possible strategy to understand the surprising phenomenon better is to determine if both the independent variables are required in the phenomenon.

KEKADA generates the different strategies with fixed order priorities without using any detailed knowledge about the experimental setup. While carrying out the experiment in which the make-effect was observed, any physicist of that time would have noted that it is possible to make an additional observation about the behavior on breaking of electrical circuit in the coil A without carrying out an additional experiment. The detailed knowledge about the setup would make that an obvious choice. In actual historical fact, Faraday made the observations of an adjoining coil when a circuit is switched on and off in a single experiment. But KEKADA operates at a higher level of abstraction and lacks the knowledge to carry out this reasoning. Thus KEKADA would ask for separate experiments to make these observations. As we mentioned previously, we allow the user to reset KEKADA's priorities so that KEKADA's behavior matches better with that of a particular scientist. The user sets a high priority for the strategy of assessing the scope of the phenomenon. Thus KEKADA decides to assess the scope of the phenomenon.

Thus KEKADA tries to find out if the effect could also be produced by the breaking of the electric

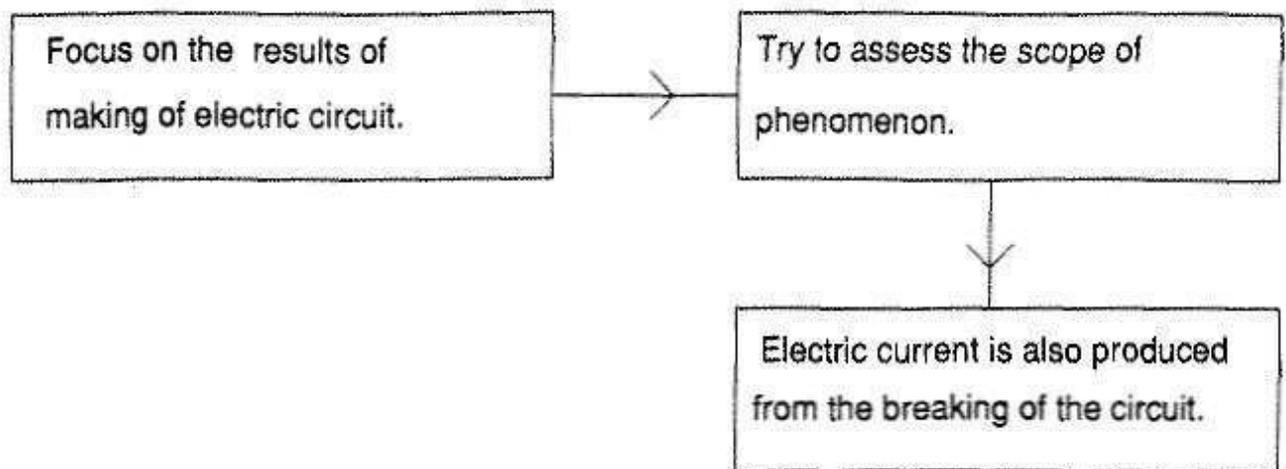


Figure 4-9: Study of induction of electricity

circuit in the coil A. When the experiment is carried out, it is observed that the electric current is produced in the second coil on the breaking of the electric current in the first coil.

It is then concluded that the effect is produced either by making or breaking of the electric circuit. The make effect provides further evidence to support some of earlier conclusions we made about the surprising phenomena. A surprising phenomenon indicates a dimension along which the program lacks knowledge. The surprise has a wider scope than originally observed. The phenomenon actually would be exhibited by any change in the current value. Unfortunately the problem space of the program does not have the concept of the derivative of the current. Thus it only considers the change from off state of circuit to on state of the circuit, and vice versa. Its attempt to determine the scope of the make effect immediately reveals this fact.

4.4.3.2. Characterization of Make-Break Effect

Next the program attempts to magnify the Make-Break effect by changing the method of measurement, by increasing the length of the side A, and by also trying to making other small changes in the apparatus. On increasing the length of side A, it is found that more electric current is produced in the second coil. The subsequent experiments are carried out on a longer coil. The rationale of the strategy of manipulating apparatus is that it increases our chances of making the crucial observations in the further experimentation on the puzzling phenomenon. This would help us understand a phenomenon better.

One strategy used by some experimental scientists who deal with complex phenomena involving a large number of variables is factor analysis. This would simplify the description of it, thereby clarifying our understanding of the phenomenon. As we identify that some variables are relevant, and others irrelevant, we reduce the sizes of relevant problem spaces.

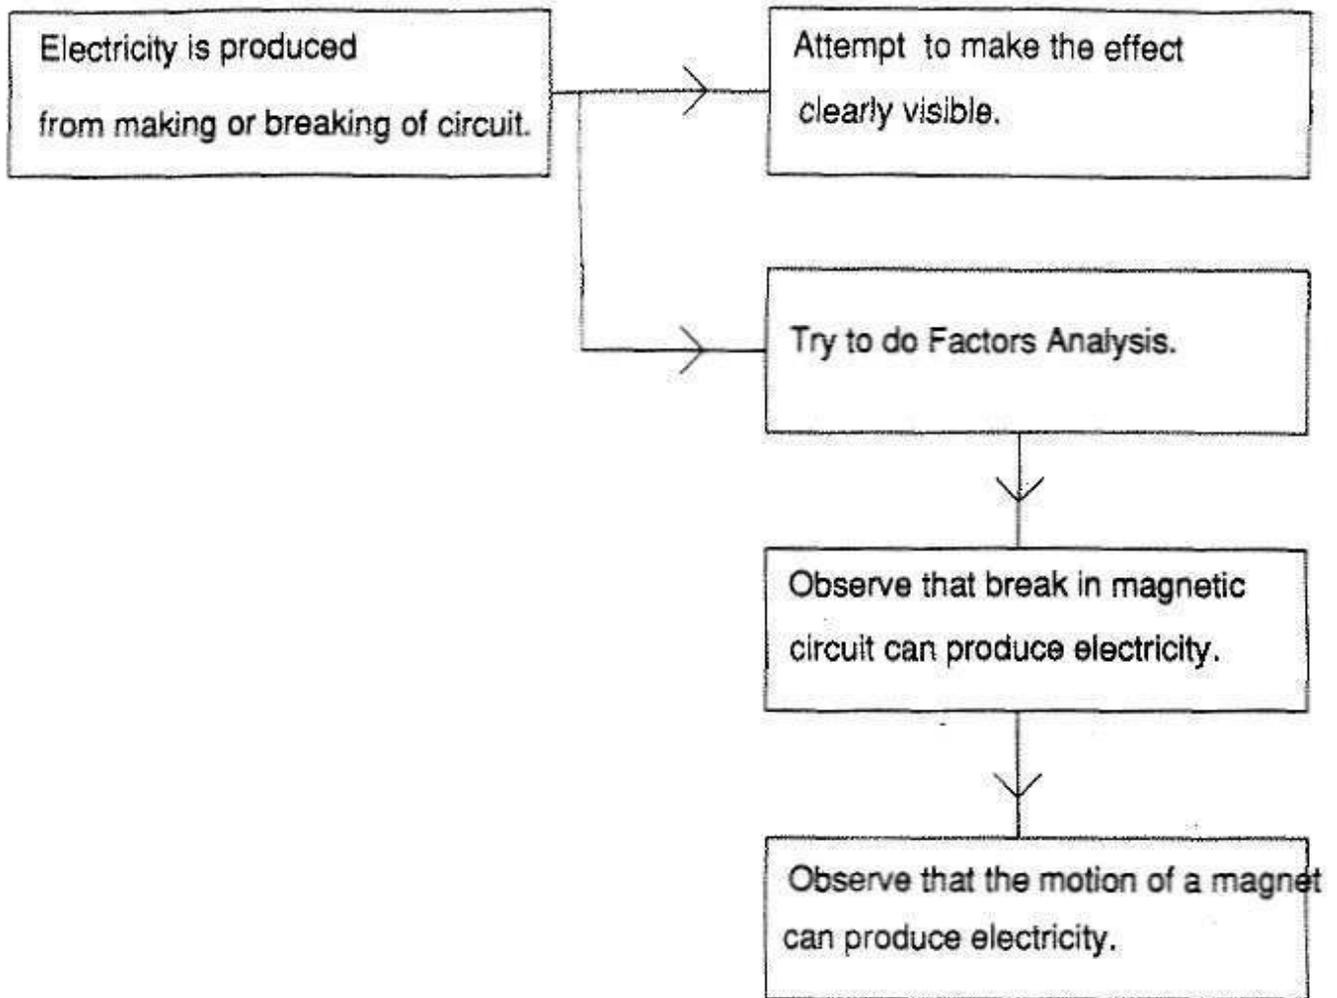


Figure 4-10: Discovery of induction from motion of a magnet

For the phenomenon under consideration, there are only two independent variables electric current and magnetism so that a simple scheme suffices. Two experiments are carried out - one involving making an electric circuit but without any core that would induce magnetism; another in which a magnetic circuit is broken but without any electricity in the coil A. KEKADA, which operates at a higher level of abstraction, does not design the apparatus in which a magnetic circuit is broken. On carrying out the experiment on the magnetic circuit, it is observed that making the magnetic circuit can by itself induce current in the adjoining coil. Immediately KEKADA decides to focus its attention on this surprising phenomenon.

4.4.3.3. Observation of induction from Magnetic Motion.

A number of hypotheses and strategies are generated about this surprising phenomenon.

1. That the phenomenon might be related to Arago's effect.
2. The effect may not be limited to making a magnetic circuit but may also be produced by other arrangements of a magnetic circuit.

3. Try to magnify the effect by making small changes in the apparatus associated with phenomenon.

KEKADA first tries to magnify the effect by making small changes in the apparatus associated with phenomenon. These attempts do not succeed.

Next KEKADA tries to examine if this effect is in some way related to Arago's effect. One comment is due here about KEKADA's access to knowledge about Arago's effect. We have allowed KEKADA to have access to Arago's effect but not many other phenomena reported in the 1820s. We haven't explored the issue of whether the 'relate similar phenomenon' heuristic in its present form will be able to generate good hypotheses if it is given access to the larger set of phenomena reported at this time.

On attending Arago's effect, KEKADA creates two hypotheses 1. That making or breaking of a magnetic circuit might be producing the electric current in the adjoining coil. 2. That the motion of magnet might be producing magnetism in non-magnetic materials.

To test the second hypothesis, it carries out an experiment on a magnet in motion. It is found that a magnet in motion can produce electricity in the second coil. At this stage the program KEKADA has discovered the induction of electricity from the motion of a magnet.

4.4.4. Analysis

4.4.4.1. Role of focusing on a surprise

In KEKADA's behavior on magneto-electricity, we saw that surprises provided pointers to the regions of experiment space containing more interesting phenomena. Thus KEKADA was able to use the strategy of focusing on a surprise effectively.

4.4.4.2. Magnification Strategy

The first strategy KEKADA employs in the face of a surprise is to attempt to magnify the phenomenon. This increases the observability of the phenomenon. Systematic manipulation of a surprising phenomenon increases the chances that a phenomenon will be observed at a later stage.

4.4.4.3. Factor Analysis

We saw that in the initial description of the phenomenon, both the magnetic field and the electric current were potential causative agents. By separating the two factors, KEKADA could infer that it is possible to produce an electric spike in a coil when the magnetic circuit is broken. Factor analysis thus simplified the description of the observed phenomenon.

4.4.4.4. Relating similar phenomena

We saw that KEKADA conjectured that there might be some relation between the observed surprising phenomenon and Arago's effect and the experiments revealed such a relation. The reason for this lies in the uniformity of nature. If there are two very similar surprising phenomena, then there is some chance that this is not a mere coincidence and that there is a common hidden mechanism.

4.4.5. Trace of the program behavior

The system starts with its focus of attention on the following phenomenon: On starting on electric current in a coil around a cylinder, a current is produced in another coil.

Heuristics	Results
HG1*	Create the hypothesis that the effect may be evident for other values of the quantity of volta-electric current. It might also be evident on breaking the electric circuit, and for steady state values of electric current. Create hypotheses that the effect may be common to make-break class of values of quantity of volta-electricity. Create the hypothesis that the effect may be common to other known forms of electricity. Thus it may be common to static electricity, Ampere's electric currents inside magnets.
HG9	Hypothesize that the cause of the hypothesis will become evident if the effect is magnified by changing the apparatus variables.
HG6	Hypothesize that the given effect may be related to the Arago's effect as both these have the independent variable "magnetism."

- HG7 The effect may be due to both the independent variables.
- HSC1,DM1 Chooses the hypothesis that the phenomenon may also be valid for breaking of electric current.
- EP2,DM1 Carries out an experiment observing what happens when electric current in the circuit is broken.
- E1 Electric current is produced in the second coil, too.
- HCM-5,HCM-7 Decides to generalize the surprising phenomenon to a new description that either making or breaking of electric current would produce electric current in the second coil under the experimental conditions.
- HG1* Create the hypothesis that the effect may be evident for other values of the quantity of volta-electric current. It might also be evident for the steady state values of electric current. Create the hypothesis that the effect may be common to other known forms of electricity.
- HG9 Hypothesize that the cause of the hypothesis will become evident if the effect is magnified by changing the apparatus variables.
- HG6 Hypothesize that the given effect may be related to the Arago's effect as both these have the independent variable "magnetism."
- HG7 The effect may be due to both the independent variables.
- HG12 Studies the phenomenon in more detail by gathering data on results of experiments for various values of 'make-break' class.
- HSC1,DM1,EP9, E1, HCM-9 *
- Chooses the hypothesis that the effect is because of invisibility of certain observations and needs to be magnified.
- Tries out the various experiments to magnify the effect. Tries out other methods of measurement. Changes the length of side A , changes instrument-type and material of coil on side A and change other apparatus variables. It is found that use of a powerful galvanometer makes the effects more easily visible. Henceforth, a powerful galvanometer will be used as a means of measurement.
- HSC1,DM1 Chooses the strategy to do factor analysis.
- EP3 Decides to carry out experiments on volta-electricity and on magnetism separately and on both together.

ES*.E1*	Carries out the experiments. Making of magnetic circuit alone produces electric current in the adjoining coil.
PG1.PC1	The phenomenon of production of electricity from making of magnetic circuit becomes the focus of attention.
HG6	Hypothesizes that the given effect may be related to the Arago's effect as both these have the independent variable "magnetism."
HG1*	Creates the hypothesis that the effect may be evident on other values of quantity of magnetism. Thus it may also be evident on breaking of magnetic circuit, and for other values of steady state current.
HG9	Hypothesizes that the cause of the hypothesis will become evident if the effect is magnified by changing the apparatus variables.
HSC1.DM1.EP9, E1	Chooses to magnify the effect. Tries out various experiments to magnify the effect by varying the apparatus variables. Nothing succeeds.
HSC1.DM1	Chooses the hypothesis that Arago's effect is related to given effect.
EP8	Creates the hypothesis that the two effects belong to common class (Effect1: Making or breaking of magnetism produces volta-electric current, Effect2: Motion of the magnet produces magnetism in non-magnetic materials.)
HSC1.DM1	Chooses the above-mentioned class-hypothesis.
EP1, E1	Carries out an experiment with the magnet in motion. It produces electricity in the second coil.

4.4.6. Results produced by KEKADA

During the research described in the previous subsection, KEKADA produced a number of interesting and important phenomena. The most important result was that electricity can be produced from the motion of a magnet towards a coil.

4.4.7. Sensitivity analysis of the problem statement

Knowledge in the problem statement P3 has a transition class with the values 'switching the circuit on' and 'switching the circuit off.' Tweney has argued that historical evidence indicates that Faraday's experience in the months before August 1831 made him prone to observing transient effects. This provides some plausibility of having such a class, but no proof.

Secondly our problem statement allows KEKADA access to the Arago's effect, but not to any other effects. We made some comments on this issue earlier in this chapter.

Thirdly, the design of actual experimental setup did involve some creative thinking in some of the experiments in Faraday's case. By choosing a suitable level of abstraction, we have thus simplified the problem.

4.4.8. Summary

The program starts with basic knowledge about various pieces of apparatus needed for electricity and magnetism experiments and with the knowledge of a phenomenon that switching on the electrical current in a coil surrounding a core can induce current in another coil. KEKADA further characterizes the effect. In particular, it shows that the motion of a magnet can induce electric current in a coil. Its goal of focusing on surprising phenomenon and trying to characterize it allows KEKADA to produce novel phenomena. In characterizing a specific phenomenon, it successfully employs a number of strategies: factor-analysis, magnification strategy, and trying to find a relation to similar facts.

4.5. Behavior on KEKADA surprise

4.5.1. Introduction

In the methodology section we noted that AI has often been characterized as an empirical inquiry. (Newell & Simon 1976, Lenat & Feigenbaum 1987) We will quote the *Empirical Inquiry Hypothesis* as has been articulated by Lenat and Feigenbaum:

"Intelligence is still so poorly understood that Nature still holds most of the important surprises in store for us. So the most profitable way to investigate AI is to embody our hypothesis in programs, and gather data by running the programs. The surprises usually suggest revisions that start the cycle over again. Progress depends on these experiments being able to falsify our hypotheses; i. e., these programs must be capable of behavior not expected by the experimenter." [Lenat & Feigenbaum, 1987]

They argue that the history of AI research shows that much of the progress has been brought about by empirical inquiries.

A number of our colleagues in CMU have been surprised by the existence of the program KEKADA [Kulkarni&Simon 1988]. That KEKADA.1, which has only about 64 heuristics, can still account for the behavior of a great scientist appeared to be unusual. One aspect of the surprise can be formulated thus : When KEKADA is run on the urea-synthesis example, the number of state transitions of the program that match with the behavior of Krebs is larger than expected. For the sake of clarity, we will refer to the program that produces this surprise as X. We will now explore what actions KEKADA strategies would suggest if faced with the surprise produced by program X.

4.5.2. Initial state

We run the program with the following background knowledge: Names of components of X by Kulkarni and Simon (1988), Names of various computers, and domains ordered by cost and availability. The run begins with the system's focus of attention on the following phenomenon: When X is run on the urea synthesis problem, the number of state transitions of the program that match with the behavior of the scientist is much larger than expected.

4.5.3. KEKADA's research on cognitive science

In response to this surprise, KEKADA generates a number of hypotheses and strategies it could use to deal with this puzzling phenomenon.

1. One possible strategy is to make the phenomenon easily visible by using better computers or architecture. One may thus use faster computers so that it is possible to experiment with the program with ease, or one may use an easily available computer such as Macintosh or IBM-PC so that the program can be easily distributed and studied by other researchers. One may try to use architecture such as GRAPES or SOAR and to see if the effects are clearer. In the current run, we have not provided KEKADA with the domain-specific knowledge necessary to reason deeply about the possible benefits of switching to an architecture with particular features. There is another possible way to magnify the observed effect that KEKADA does not suggest. One might examine larger numbers of state-transitions. Thus one might want to see if we can extend the program so that it accounts for Krebs' day-to-day work over his entire lifetime. Some of the effects that are not easily visible in the current run would be visible in such a run, should it become possible to produce it.

2. As the program X has a number of subprocesses, a possible strategy is to divide and conquer. Try to focus on one of the subprocesses and try to understand that subprocess better. One might thus want to focus only on the issue of choosing a problem or proposing an experiment and study that issue in much more detail.

3. Another possible strategy is to vary the amount of knowledge KEKADA has. How much better performance will KEKADA display if it is provided with a larger data-base?

4. Similar phenomena might exist even for task types other than post-surprise experimentation. Can we produce similar phenomena for programs making mathematical discoveries, or for discoveries involving a change in the representation?

5. Similar surprising phenomena might exist in other scientific disciplines. The phenomenon observed might not be restricted to biochemistry. Thus KEKADA when provided with appropriate domain knowledge might be able to produce behavior resembling behaving functions of scientists in other domains.

6. All three independent variables: the general-component, the domain-specific knowledge, and the way the initial state is encoded may be responsible for the behavior.

As one sees, all of these constitute reasonable strategies to deal with the observed surprising phenomenon.

It is routine in AI to see if lessons learnt from an AI program running in a specific domain are valid in other domains. KEKADA decides to assess the scope of the surprising phenomenon. The advantage of assessing the scope would be that if the surprising phenomenon is found to have more scope, then we will know that the features of a specific domain have not caused the surprising phenomenon and it would be easier to reason about the puzzle.

KEKADA has a list of domains ordered by availability and cost of implementation. It chooses domains from this list one by one. It asks the user to see if KEKADA's general component when added with appropriate domain-knowledge can produce the behavior resembling that of Faraday on magneto-electricity, that of Williamson on alcohol structure, and that of Krebs on glutamine-synthesis. When the user confirms that they do, KEKADA infers that the general components of program X when supplemented by appropriate domain-specific knowledge would exhibit behavior resembling that of a scientist doing post-surprise experimentation in any scientific discipline.

In reality, however, on the basis of our general knowledge of how various sciences operate, we would realize that this is an over-generalization. It would be safer just to say that the phenomenon is valid

for other domains like the ones we looked at. Here we have one more illustration of the importance of the scope strategy. A generalized phenomenon is of more importance to the scientific community than a specific one. If a scientist starts to study experimentation in a particular domain, his work can benefit from conclusions about a phenomenon that is valid for a broad variety of domains. But it might be difficult for him to relate his work to a result that is specific only to, say, biochemistry.

KEKADA at this stage again generates a number of hypotheses and strategies about this generalized phenomenon:

1. All three independent variables: the general strategies in KEKADA, appropriate domain-knowledge for a given domain, appropriate initial encoding for the initial state may be responsible for the behavior. KEKADA does not have the ability to do a causal analysis of the program like that done in the chapter 5 of this thesis. Chapter 5 is another example of how it is much easier to make a causal analysis if we can show that a phenomenon is valid for a large number of cases, and not limited to a specific case.

2. Make the phenomenon easily visible by using better computers or architecture.

3. One possible strategy is to vary the amount of knowledge KEKADA has.

4. Similar phenomena might exist even for task-types other than post-surprise-experimentation.

5. Divide and conquer might be a possible strategy to understand the phenomenon better.

Here we see that the generators in KEKADA are able to suggest reasonable hypotheses and strategies at an abstract level. An experiment in cognitive science research can take months to carry out. The thought involved in carrying out such an experiment would not show up in our simulation. Researchers might spend a lot of time thinking about the global goal of the experiment before actually undertaking the experiment. They might thus bring to bear significant amount of knowledge while attacking the problem.

The next subsection gives the actual trace of the program behavior.

4.5.4. Trace of program behavior

Heuristics	Results
HG3	Divide and conquer is a possible strategy to understand the phenomenon better.
HG1	Another possible strategy is to vary the amount of knowledge KEKADA has.
HG1	Similar phenomena might exist even for task-types other than post-surprise-experimentation.
HG1	Similar phenomena might exist in other scientific disciplines.

- HG7 All three independent variables may be responsible for the behavior.
- HG9 Make the phenomenon easily visible by using better computers or architectures.
- [Begin seq1]
-
- HSC1, DM1, EP1, DM2, E1
- Decides to assess the scope of the surprising phenomenon. Repeats the surprising phenomenon on Krebs' later work leading to glutamine synthesis discovery. Result is that the number of transitions matching Krebs' behavior remain high.
- HCM5 Increases the confidence in the hypothesis that the phenomenon may be valid for other scientific disciplines.
- {End seq1}
-
- Repeat [seq1] for a number of disciplines.
- HCM-7 Concludes that the originally observed phenomenon is more general and that for any scientific discipline, the above-mentioned phenomenon will be observed.
- HG9 Makes the phenomenon easily visible by using better computers or architectures.
- HG1 One possible strategy is to vary the amount of knowledge KEKADA is provided with.
- HG1 Similar phenomena might exist even for task-types other than post-surprise-experimentation.
- HG3 Divide and conquer might be a possible strategy to understand the phenomenon better.
- HG7 All three independent variables may be causally responsible for the behavior.

4.5.5. Results produced by KEKADA

In the research described in the previous section, KEKADA tried to determine the scope of a surprising phenomenon and furthermore concluded that the phenomenon has wider scope and is not restricted to a particular domain. The program further suggested a number of hypotheses and strategies which could be used to extend this research.

4.5.6. Summary

In this section we saw a specific example in which the strategy to focus on a surprise and to assess the scope of a phenomenon were effective strategies. Furthermore we saw that the hypothesis and strategy generators of KEKADA were capable of generating reasonable strategies in the specific instance.

4.6. Other Cases from History of Science

In this section I will describe some more cases from the history of science. These cases are typical cases, describing behavior of scientists in the face of a surprise. They would serve to demonstrate use of KEKADA heuristics by scientists in many different domains.

The account of Priestley's work is based on the historical study by James Conant (1957). The remaining accounts are based on the descriptions in Asimov's Guide to Science (1972).

4.6.0.1. Priestley's work on oxygen

Priestley's discovery of Oxygen was "the central event in the overthrow of the phlogiston theory." [Conant, 1957] Before Priestley started his work on red oxide of mercury, Bayen had reported erroneously that red oxide when heated produces carbon dioxide. Thus when Priestley carried out the experiment heating the red oxide, he was surprised to find that the gas evolved was not carbon dioxide. He further found that a candle burned with vigorous flame in this gas. He therefore erroneously concluded that the gas was laughing gas. He proceeded to check if two other substances - red precipitate and red lead also would produce the same gas when heated. Both these substances when heated produced a gas in which a candle burned vigorously. Priestley further tried to test the gas by other methods. Unlike the laughing gas, this gas burned vigorously even after shaking in the water. Priestley then tried the nitrous air test in which a gas is mixed with nitrous air and a candle is burned in the mixture. The nitrous air test showed that the gas was more wholesome than the common air. This new gas was oxygen. Thus on his research on oxygen, Priestley effectively used three of the KEKADA strategies- focusing on surprises, magnification, and determination of scope.

4.6.0.2. Discovery of X-rays

Roentgen's discovery of X-rays was an immediate sensation, as scientists had not previously known any radiation with powerful penetrating capability. On November 5, 1895, Roentgen stumbled upon a surprising phenomenon. He had a cathode-ray tube in a closed cardboard box. He found that the tube emitted radiation that made a luminescent material outside the box glow. Roentgen was surprised to find that there existed radiation which could penetrate through a box. Next he applied the magnification strategy. He tried to see if the radiation could even penetrate the walls of the room, and it did. The amazing penetrating power of these rays immediately made clear the importance of the phenomenon.

4.6.0.3. Discovery of radioactivity

After Roentgen reported his discovery, other scientists tried to assess the scope of the phenomenon by asking: do radiations other than cathode rays also contain X-rays? Henri Becquerel was studying to see if the radiations of the potassium uranyl sulphate contained X-rays.

This could be tested by exposing the sulphate to sunlight while the compound is put on a photographic plate wrapped in black paper. If the fluorescent radiation contained X-rays, it would darken the plate. However Becquerel had to postpone the experiment because of lack of sunlight in cloudy weather. After several days, he developed the photographic plates, even though he had not carried out the experiment. To his surprise, the plate was darkened by exposure to some radiation. Thus the emission of the radiation did not depend on the potassium uranyl sulphate being exposed to sunlight. When he further tried to assess the scope of the phenomenon, he found that even other uranyl compounds exhibit radioactivity.

4.6.0.4. Discovery of radioactive chains

Crookes was surprised to find that a substance he had isolated from uranium was more radioactive than uranium. He referred to this substance as Uranium X. Becquerel further characterized the phenomenon and showed that Uranium continually generates Uranium X. Next the scientists attempted to assess the scope of this phenomenon by trying to see if other radioactive substances exhibited the same property. Rutherford found thorium generated thorium X. Furthermore soon the scientists came up with a series of new substances and named these substances radium A, radium B, etc. It became evident that radioactive substances are transformed into new substances in radioactive chains. Here we see another example of how original surprising phenomenon indicated a dimension along which scientists lacked knowledge and how the scope-strategy can be effectively used to characterize a surprising phenomenon.

4.6.0.5. Discovery of Photo-electric Effect

Lenard won a Nobel Prize in 1905 for his discovery of the photo-electric effect. Later it was reported that light with increased intensity did not emit electrons with more energy, but light with different wavelength could. Einstein showed that Planck's quantum theory explained this surprising phenomenon. KEKADA clearly lacks the ability to create an elaborate explanation involving a deep theory.

Chapter 5

Analysis

That "creativity" is beyond analysis is a romantic illusion we must now outgrow.
Peter Medawar(1969)

In this chapter, we will analyse KEKADA and its performance.

In the first section, we produce statistics showing that KEKADA heuristics as they are encoded are not specific for the problems on which they are run, but are general in applicability. In the second section, we will discuss the empirical evidence showing that most of the KEKADA heuristics have in fact been used in more than one problem. The third section notes the variety of domains and problems on which KEKADA was shown to be useful. In the fourth section, we will analyse the reasons why KEKADA heuristics were effective on the problems in the last chapter and we discuss the role each aspect of KEKADA played in its success. The fifth section views KEKADA as a parsimonious model producing the behaving function of Hans Krebs. Together this evidence suggests that KEKADA strategies would be capable of producing intelligent experimental program on problems in a wide range of domains.

In another section, we will attempt to characterize the class of problems on which KEKADA would be useful. We will also give reasons why exact characterization of such a class may not be possible.

5.1. Generality of the heuristics

As figure 5-1 indicates, 28 of the 43 heuristics in the program are domain-independent. The remaining 15 are specific to a domain such as biochemistry, but none are specific to a particular problem such as urea synthesis. However we allow user-defined ordering of strategies in the program. The effect of the ordering is shown in the figure 5-3.

Category of Heuristics	General	No	Domain-specific	No
Expectation-setters	ES1-5	5		
Hypothesis Generators	HG1,3,6,7,9,11,12	7	HG2,4,5,8,10	5
Problem Generators	PG1	1		
Problem Choosers	PC1	1		
Experiment Proposers	EP1,3,7,8,9,10	6	EP2,5,6,11	4
Hypothesis or Strategy Choosers	HSC1	1		
Experimenters	E1	1		
Hypothesis or Confidence Modifiers	HCM3,5,7,9	4	HCM1,2,4,6,8,10	6
Decision Makers	DM1,2	2		
Total		28		15

Figure 5-1: Generality of heuristics

5.2. Applicability of the heuristics in multiple tasks

As shown in the figure 5-2, 31 of the heuristics were actually used in more than one task.

5.3. Empirical evidence of the effectiveness of KEKADA

The program has been run in 4 different domains and was able to produce interesting research results in all these domains. Further, we also looked at a random sampling of some cases from the history of science and found that these strategies were widely used.

Number of heuristics used in more than one problem = 31

Number of heuristics used in only one problem = 12

Number of Problems	The heuristics	No
5	EP1,HCM5,HG1,HG7,HG9,HSC1,DM1,E1	8
4	HCM7	1
3	EP2,EP9,HG6,DM2,PG1,PC1 ES1,ES2,ES3,ES5	10
2	EP3,EP6,HCM3,HCM10,HG2,HG3 HG4,HG5,HG11,HG12,ES4,HCM2	12
1	EP5,EP7,EP8,EP10,EP11,HCM1,HCM4 HCM6,HCM8,HCM9,HG8,HG10	12

Figure 5-2: Applicability of heuristics in multiple tasks

5.4. Analysis of the program performance

In this section, we will discuss the role various KEKADA heuristics play in the performance of the program.

5.4.1. Role of dual search of KEKADA

We saw throughout the last section that the central source of gaining new knowledge was experiments on nature. At times experiments were only confirmatory, they allowed KEKADA to change confidence in an existing hypothesis. Thus KEKADA's experiment with very low concentration of ornithine verified the catalytic hypothesis it was entertaining. At other times experiments gave new information about what the outputs of the experiment are. Of special importance, were the experiments that produced surprising results, thus giving a lead to the regularities nature held. The experiments were used for gathering new information about the external environment and for modifying confidences in the existing information.

However the experiment spaces were large in size and random experiments in these spaces would not have produced interesting and useful information. Therefore the program always started with a hypothesis or a strategy generated by its hypothesis and strategy generators and used experiment-proposers to generate an experiment. We saw that such a scheme produced experiments with an acceptable rate of generating good information. The program also needed knowledge to interpret the results of the experiments. Thus a specific piece of knowledge allowed it to conclude from an experiment that ornithine is a catalyst. Thus *the strategy of carrying out a dual space search was at the heart of KEKADA's ability to make new discoveries, that is generating useful new information.*

5.4.2. Focusing on surprise as a search-control strategy

In all the twelve cases discussed in the previous chapter, surprises are pointers to parts of the problem spaces where a number of phenomena (interesting with respect to initial knowledge) exist. In all these cases, the decision to focus attention on such surprising phenomena enabled the program to keep its attention on "interesting" phenomena. In KEKADA's behavior, *the surprises play the crucial role in discovery of novel phenomena.*

Beyond this empirical evidence, the analytic reason for the surprise-focusing strategy is clear: when a phenomenon in the external environment is surprising with respect to the internal knowledge, it indicates a dimension along which the system is missing significant knowledge. Understanding a surprise may remove a deficiency in internal knowledge. Due to the uniformity of nature, similar interesting phenomena may exist.

Once a program detects a surprising phenomenon and puts it on the agenda, KEKADA has a heuristic that always prefers the most recent surprising phenomenon over any other tasks. Though this scheme works fine in all the cases in which the program has been run, analytic evaluation would indicate the scheme is not fully satisfactory. KEKADA does not have the fundamental ability to choose between two different puzzling problems. KEKADA.1, the previous version of KEKADA (Kulkarni & Simon, 1988) had the following criteria:

[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.

These criteria are the kind of criterion we would use to choose among the problems in science. Essentially we want to choose problems that are both important and solvable within the state of the art. KEKADA does not have the ability to evaluate the importance of a new problem. One possible criterion is to see if the surprise is related to previous supergoals of the system.

We leave this as an open question which the program does not address. If the goal of science is to produce and understand "novel and useful" phenomenon, our program has a notion of "novelty", but lacks the ability to evaluate the relative importance of focusing on different tasks on the agenda.

One may want to ask why KEKADA's simple scheme works. One would notice that KEKADA's experiments are related to its general objectives, and thus the surprises will actually turn out to be relevant to the general area of interest of the system.

As Pasteur said, surprises occur to the prepared mind. KEKADA prepared its mind by setting expectations with every experiment it carried out. Thus we saw how expectation-setters along with the surprise-detector heuristic were needed for implementation.

5.4.3. Role of KEKADA control structure

As we noted in chapter 3, *KEKADA's control structure simply combines the ideas of its dual search and surprise-focusing strategy without allowing the two to interfere with each other.* We saw in the previous section that it does serve that purpose in all scenarios.

5.4.4. Strategies used to understand the surprising phenomenon better

We saw that having focused on a surprising phenomenon, the program employs the following strategies.

1. Magnify the phenomenon by varying apparatus variables.
2. Divide and conquer: the surprising effect may depend on one of the subprocesses.
3. Assess the scope of the surprising phenomenon.
4. Determine if all the independent entities are necessary to produce the surprising phenomenon.
5. Try to find any relation to closely related phenomena and study any hypotheses suggested domain-specific heuristics.

6. Use a number of domain-specific strategies.

Below we discuss the role these strategies play in the behavior of the program.

5.4.5. Magnification Strategy

Empirical evidence of usefulness of the strategy: In all the cases we saw, the attempts to magnify the phenomenon were limited to a few experiments, and thus a reasonable cost. In the magneto-electricity scenario, some of the effects may not have been visible if the program had *not* attempted to first magnify the effect and then tried other manipulations. In the case-study of Rutherford's work on alpha particle bombardment, his decision to use a powerful device made visible what was the first man-made transmutation. When Roentgen observed the surprising radiation that penetrated the card-board, he checked if the radiation could penetrate the thicker wall, too; and it did. The surprise in the magnified form- that the radiation could penetrate a thick wall- was a more novel and important phenomenon.

Analytic evidence of the utility of strategy: The advantage of magnification is clear: it increases our chances of making the crucial observations in the further experimentation on the puzzling phenomenon, which would help us clarify the surprise. For example, at times, the surprise may occur because that the observations are spurious or small in magnitude and are not visible. Secondly if a surprise is magnified, it is likely to be more "interesting" for the same reasons discussed above about why surprises are interesting. A magnified phenomenon may also be more useful, in some cases.

In the program's runs, the cost of this strategy was limited due to the fact program only knows a limited number of relevant variables and sets of values associated with them. However the program does not have the notion of 'cost,' and adding it might be one way of improving the performance of the program. Second, the program uses a relatively simple scheme in its attempts to magnify the phenomenon. Another way to extend the program would be to allow it to attempt magnification by using more elaborate general problem solving processes. This might allow it to make the kinds of changes in the apparatus which Anderson carried out when he observed a particle behaving unexpectedly.

Concomitant variation of an independent variable with a dependent variable can be also be used to infer a causal relation between them. KEKADA at present lacks such an ability. But that would constitute a useful extension.

5.4.6. Determination of scope

The ornithine effect in kidney and the surprises in the AI case and in the ether case were more general than the originally observed phenomena. The history of science provides numerous examples where initially observed specific surprising phenomena are found to be more general on closer examination. (Some of these were cited in the special section on examples from the history of science in the previous chapter.) In all these cases *the surprising phenomenon is interesting, hence showing that it is valid for a general class furthers its importance*. As we saw in the ether case, there are fewer hypotheses that can fit a generalized description than a particular one, thus we have a simplified problem which can be solved more easily. In the case of the study of the ornithine effect in kidney, after observing that the effect is general to the class of amino acids, experiments were then carried out with alanine, a cheaper and more reactive agent. Besides, prior knowledge about amino acids helped in understanding the effect. Thus two more advantages of generalizing a phenomenon are that *it improves our experimental control by giving us much wider choice, and allows us to use knowledge about a wider class*. As we are likely to have more knowledge about a wider class, it increases the chances of solving the problem at hand.

Another point to be noted is that we are not seeking precise quantifications while determining the scope of the phenomenon. Only rough quantifications suffice to give us the above-mentioned advantages. We know from BACON that a greedy regularity seeking heuristic is useful in the global goal of seeking regularities. Our wish to generalize is also partly a reflection of the regularity seeking character of nature.

Now we need to assess the cost of the strategy. In the cases we looked at, this is a limited cost strategy. We discontinue using the strategy after the cost exceeds a certain value. As in many cases, more specific understanding of a puzzle is likely to involve a long-term effort, searching a lot of very specific explanations (the space of specific hypotheses about the causality may be large), a limited amount of exploration right in the beginning is a highly cost-effective strategy.

Concomitant variation of an independent variable with a dependent variable can also be used to infer a causal relation between them. KEKADA at present lacks such an ability.

5.4.7. Divide and conquer

Given a certain state of sophistication of instruments, certain hypotheses can be tested by a single experiment, whereas other complex hypotheses can only be tested by a large number of experiments. If a complex hypothesis has many subcomponent hypotheses that can be tested individually by a single experiment, then it would be efficient to consider one subcomponent at a time, carry out an experiment and use the information gained from this experiment in further choice of hypotheses.

Furthermore, generators that create a simple component of a more complex hypothesis can be potentially used to generate a larger number of hypotheses, and would give a system required parsimony. Thus a parsimonious system would only possess generators that can create simple hypotheses, and complex hypotheses would have to be generated by combining simple hypotheses. Therefore it is much easier for such a system to generate good simple hypotheses about subcomponents, than generate good complex hypotheses.

For these two reasons, an agent might prefer to divide a complex phenomenon into simpler components and focus on a component at a time. This will enable it to generate good hypotheses and strategies and it will also be able to test each hypothesis with a single experiment.

The program's strategy of dividing processes into subprocesses is a specific implementation of our intuitive notion of the divide-and-conquer strategy. Even the factor-analysis strategy used by the program can be viewed as a specific implementation of the divide-and-conquer strategy. These implementations together however do not completely capture our intuitive notion of the divide-and-conquer strategy.

Reducing the complexity of a phenomenon by focusing on a smaller part allows an agent to have good generators of hypotheses about a part of the phenomenon and also conveniently test a hypothesis. KEKADA used this strategy only in the urea synthesis scenario and it did not pay-off. We have discussed the analytic reasons why such a strategy is nevertheless known to be useful.

5.4.8. Factor analysis

We saw that this strategy was effective in the case of magneto-electricity. In the make-effect both magnetic field or electric current individually or together could have caused the observed induction of electric current. By focusing on each factor separately, KEKADA simplified the phenomenon, thus clearly establishing that a magnetic field alone is capable of inducing electric current.

In general, the factor analysis strategy allows us to remove irrelevant variables. Thus it simplifies the problem by reducing the degrees of freedom, and thus the size of the relevant problem space.

5.4.9. Gather Data Strategy

We saw that while KEKADA gathered more data about the deamination reaction by running experiments on various amino acids, it produced important data on deamination of amino acids and also came across the surprising glutamic acid effect. This shows the usefulness of the gather-data strategy. When one finds a novel and unusual phenomenon, even systematic collection of data about the phenomenon can be of great interest to the scientific community. Furthermore while collecting such data, surprises can turn up.

5.4.10. Related phenomenon

KEKADA also has a strategy which recommends trying to find relations between the surprising phenomenon and a similar phenomenon.

We saw the utility of the strategy in the magneto-electricity case. When KEKADA was studying the phenomenon in which breaking a magnetic circuit induces current in the adjoining coil, it tried to relate it to Arago's effect. This led the system to see if motion of the magnet can produce electricity in the coil and it does. KEKADA generates this strategy with one of the lowest preference.

As KEKADA is supplied with a limited amount of knowledge, one might question whether, in case KEKADA had general knowledge about lots of phenomena, the similarity heuristics might detect such a large number of apparently related phenomena as to create a search problem. We leave this as an open question.

5.4.11. Domain-specific strategies as efficient and practical instantiations of General strategies

The domain-specific strategies are efficient and practical instantiations of more general principles. We would like to illustrate this with the implementation of Mill's difference principle.

Mill's Difference Principle tells us that if under conditions C1, phenomenon P occurs and under C2 it does not occur, then the differences between C1 and C2 are causally responsible for the phenomenon P. But when we apply this principle to a phenomenon in which little urea is formed in tissue slice and it is known that an amino acid would form urea under the same conditions in the human body, difference principles could potentially generate hundreds of organic compounds and other conditions that exist in the body but are not present in tissue culture. These hypotheses would create a search problem and could

not be easily tested by an experiment. Our domain-specific implementation would suggest only testing stimulators. Thus it suggests a few substances which can be experimentally tested.

In a task-environment where our access to instruments is limited, in some cases knowledge can't be replaced by more search, because hypotheses generated in a larger search space can't be tested with the available instruments. Alternatively if one has access to all the instruments that are used by the scientific community and the skills to use these instruments, knowledge would only trade against much larger amount of search.

Thus by using the general principles in domain-specific forms, we reduced the search problem and generated some plausible hypotheses which could be experimentally tested.

5.4.12. The effect of the intervention of the user

As we mentioned earlier, we have allowed the user to reset the priorities among different strategies. In some cases, this was done to get a better fit with the behavior of the scientist. It also resulted in reducing the run-time of the program. However the system does not go astray in the absence of this user-interference. It just needs to carry some additional experiments before making the discoveries.

Figure 5-3 shows the number of experiments KEKADA needs to carry out in the absence of the user intervention. It shows that the number of experiments would increase only by a factor of about two in the absence of this intervention. (In the cases where KEKADA generates a set of strategies all with equal priority, the table refers to the worst case scenario.)

5.4.12.1. Urea synthesis problem

To illustrate how the user-intervention is used, let us examine a particular example.

When faced with the Ornithine effect, KEKADA produces a number of strategies and hypotheses, with the priorities indicated in the brackets.

```

Try to magnify the effect (1)
Divide and conquer (3)
Assess the Scope over the class of amino acids(4)
Assess the Scope over the class of amines(4)
Assess the Scope over the class of carboxylic acids(4)
Vary the concentration of Ornithine and of ammonia(4)
Factor analysis(5),
Ornithine may be acting as a catalyst(6),
Ammonia may be acting as a catalyst(5),
Hypotheses about what reactant might be donating what group(6),
The effect might be related to the arginine reaction(6) .

```

User orders the strategies as below:

Problem	Experiments with user-defined priorities	Experiments without user-defined priorities (worst case)
Urea Synthesis	48	82
Amino-acid metabolisms	34	71
Magneto-electricity	37	54
Alcohol Structure	3	12

Figure 5-3: Number of Experiments carried out

Try to magnify the effect (1)
 Assess the Scope over the class of amino acids (4)
 Assess the Scope over the class of amines (4)
 Assess the Scope over the class of carboxylic acids (4)
 Ornithine may be acting as a catalyst (6).

In the absence of the user's intervention, the program would have to choose randomly between the strategies with equal priorities. In the worst case scenario, KEKADA carries out 22 more experiments in the absence of user-defined priorities as compared to the run with such priorities. In this particular case, Krebs did carry out experiments to test such hypotheses as: Ornithine might be donating an amino group to urea. Therefore the program will produce a better fit with Krebs' behavior in the absence of these particular user-defined priorities. However in almost all of our other cases, we have used user-defined priorities to improve the fit with the behavior of the scientist.

5.4.13. The role of the missing processes

In this previous chapter, we noted a few discrepancies between the behavior of KEKADA and that of a scientist working on the similar problem, and we found that the scientists were employing some additional processes. Below we will list the various processes we observed and discuss if these processes could be used to improve the performance of the program.

Processes may do work that is not directly relevant for the task solution, but might be important on

some other problems. For example, consider the experiments Krebs carried out on liver tissue slices of well-fed and ill-fed rats. KEKADA does not have the knowledge about the relation of experiments in tissue slices to reactions in the human body, and thus does not carry out such experiments. However there could be other problems where possessing such knowledge is necessary to make progress.

In some cases, scientists employ additional processes to use new results in the context of other related problems. In the previous chapter, we discussed Priestley's use of KEKADA strategies. While Priestley was carrying out these experiments, he was also looking at the results and using STAHL-like heuristics to make inferences about components of chemical reactants. KEKADA works on only one problem at a time and limits the kinds of inferences it does from a new experiment.

As we saw in all our examples, KEKADA works at a higher level of abstraction, and the scientists do need certain processes to be able to plan the details of an experiment. Processes working at a lower level of abstraction can occasionally affect the designs at a higher level of abstraction, as we saw in the case of Faraday's observations on switching off an electric circuit. As KEKADA works only at a higher level of abstraction, it is simpler to design, but is unable to reason at a lower level of abstraction and use this reasoning in the context of the problem at a higher level abstraction.

Lastly, while working on one particular problem, productions whose condition sides are not met, would not play any role. Thus our chemistry-specific heuristic do no work in the run on magneto-electricity and do not play any role. Thus a large number of processes a scientist possesses may not play any role in solving a specific problem.

One would therefore expect that our program, lacking these processes, would be able to make discoveries on a range of problems smaller than the range of problems a typical researcher is capable of. When we tried to set if KEKADA can run on an arbitrary problem faced by Faraday, we found that there is direct empirical evidence indicating so.

5.4.14. Role of the Architecture

Processes were represented as productions in OPS5. As many of KEKADA strategies take the form of conditions and actions, that was particularly appropriate. However there are some KEKADA strategies which reason about the member class relation. An architecture like GRAPES which has schemas and rules would be better suited to these.

5.5. KEKADA as a model of Hans Krebs

KEKADA can also be viewed as a parsimonious model producing the behavior of Hans Krebs. We include in the appendix a paper that analyses the validity of KEKADA as a model of the heuristics Hans Krebs used in his discovery of the Ornithine cycle. (Kulkarni & Simon 1988) This provides additional evidence that KEKADA strategies should be applicable to a larger class of problems than the five tasks in the last chapter.

5.6. The class of problems over which KEKADA is applicable

Even in the case of a human researcher, we can not characterize precisely the class of open problems in a research area s/he will be able to solve. We might say that the area of specialization of a researcher is graph theory. That would mean that he is capable of solving some problems in that area. We can only characterize a wider class of problems of which he has the ability to solve some. Thus experimental physicist might be able to make some discoveries in experimental physics after trying his hand on a larger number of problems. In the case of KEKADA, this class can be characterized as the class of problems in empirical sciences in which a surprise is involved. However the size of the class of problems it would be able solve would be much smaller than the size of the class of problems a typical human researcher can solve.

5.7. Summary

In this chapter, we discussed the empirical and analytical evidence that shows that KEKADA strategies are capable of carrying out an intelligent experimental program on a wide range of problems. KEKADA strategies are not specific for the particular problems on which they were run. Furthermore most of these have been useful on more than one problem in the runs in the previous chapter. We also analyzed the reasons why KEKADA strategies are effective in producing interesting results.

Chapter 6

Conclusion

Whence is it that nature does nothing in vain; and whence arises all that order and beauty which we see in the world?

Sir Issac Newton(1931).

6.1. Summary of the Ideas in the Chapters

In Chapter 2 we characterized the nature of the problem of scientific research and defined a set of problems. We also described the methodology used to study this problem. Chapter 3 described the program KEKADA. KEKADA carries out a dual space search . We discussed how KEKADA focuses on surprises and tries to characterize them. Chapter 4 described the behavior of the program when it is placed in different task-situations. We ran the program successfully in the domains of biochemistry, magneto-electricity, chemistry and cognitive science. Chapter 5 analyzed the behavior of the program. We discussed the empirical and analytical evidence indicating the effectiveness of KEKADA strategies in a wide variety of domains. This chapter will summarize the conclusions drawn from the work .

6.2. Central Thesis of this research

The thesis of this research is that KEKADA is capable of carrying out intelligent experimental programs on problems similar to those faced by a number of experimental scientists. KEKADA has a set of experimentation strategies, that were detected from the traces of the behaviors of scientists. KEKADA strategies include : focusing on a surprising phenomenon, characterizing the surprising phenomenon by general strategies such as magnification, applying divide-and-conquer, determining the scope of phenomenon, factor-analysis, relating to similar phenomena, and domain-specific strategies and hypotheses. The domain-specific heuristics in KEKADA are efficient and practical instantiations of general strategies such as - controlled experimentation, determination of complexity of a process, testing of a causal chain, componential analysis, differencing and divide-and-conquer. However KEKADA lacks

many abilities a human researcher possesses and thus it would not be a good experimental scientist in its present form.

6.3. Contributions of the work

This work advances the state of the art in Artificial Intelligence by producing a set of computational strategies for experimental research applicable in a wide variety of domains. The KEKADA work also shows how a detailed historical account of scientists work could be examined to extract domain-independent and domain-specific computer-implementable strategies that scientist used in his research.

6.4. Relation to other work

In this subsection, we will discuss the work in discovery and experimentation which is related to my research work.

Dendral and Meta-Dendral: Dendral showed that one way to deal effectively with large search spaces, is to use knowledge to identify constraints, and then apply the generate-test method. Thus Dendral suggests one way to deal with problem of creating good generators. While DENDRAL can only use information about structures from mass spectroscopy, KEKADA's study of the chemistry of alcohol shows that it is possible to get more knowledge by carrying out reactions and this knowledge may serendipitously lead to clues about the structure of some molecules.

BACON: Langley, Simon, Bradshaw and Zytkow(1987) designed a series of programs, including BACON, GLAUBER, DALTON and STAHL, that used a small number of simple and general heuristics to induce a large number of theories from data. A natural question which BACON programs raised was where the data came from. The present research is aimed at addressing that issue and is a natural extension of the BACON research. To illustrate the point clearly, consider the phenomenon that red oxide of mercury when heated produces oxygen. This reaction has been used as a data point by STAHL(Zytkow&Simon1986), STAHLP(Rose&Langley1988) and REVOLVER(Rose&Langley1988) programs. We earlier discussed how Priestley used KEKADA strategies in producing this data point.

Other programs such as ABACUS (Falkenhainer&Michaelski, 1986),and COPER (Kokar, 1986) have further addressed some of the issues in the induction of theories from data.

Schank's work: Roger Schank (1982) has argued for a central role for expectation failures in learning. His work has largely focused on reminders in the face of expectation failures. KEKADA work

focuses on experimentation strategies which could be used in the face of an expectation failure and thus is more in spirit of the work of Carbonell (e.g. Gil&Carbonell 1987). A number of KEKADA strategies (e.g. determination of scope) would benefit from reminders of similar experiences. Our mechanism for retrieval of similar phenomena in the 'relating similar phenomena heuristic' is relatively simple and would benefit from a more sophisticated model of memory. In the cases in Schank's book, one does not have the freedom to carry out experiments on external environment and therefore KEKADA strategies could not be used. On the other hand, the explanations in his book would not be applicable to many of the surprises KEKADA encounters. With the state of the art in the 1930s, the only explanation of the ornithine effect was the detailed mechanism by which urea is formed.

MOLGEN: In 1979, Friedland designed MOLGEN which used skeletal plan refinement to design experiments in molecular biology, seeking to model the human experiment design process. His work shows that the design of experiments involves a gradual refining of more abstract plans. KEKADA is capable of designing experiments at a certain level of abstraction, leaving out the details. The design of the left-out details would need further use of (mainly domain-specific) heuristic processes.

While MOLGEN can generate a plan for a specific experiment, KEKADA can interpret the results of the experiments, using them to modify the hypotheses and in turn again carry out further experimentation. Friedland and Kedes (1986) (Also, see Karp 1986) have further proposed to build MOLGEN2 using interviews of Yanofsky in molecular genetics.

As KEKADA works on a domain with a shallow theory, the credit-assignment problem is not a crucial one. Accumulation of empirical data in the form of programs working on different domains with different characteristics would eventually enable us to understand which mechanisms are effective in which kind of domains.

EDISON: EDISON (Dyer,Flowers&Hodges1986) uses processes of mutation, generalization, and analogical reasoning, and organized, indexed episodic memory to invent novel new devices. As the KEKADA work in electro-magnetism produces phenomena which were directly used as new devices, we could make two points. KEKADA strategies which focus on a surprise and attempt to characterize it are useful in the invention of new devices. However KEKADA also used the the heuristic of relating similar phenomena in one of its discoveries in magneto-electricity. In that case, it related an existing phenomenon with a previously known phenomenon, that reported by Arago. We also noted there that as we had supplied KEKADA with Arago's effect but haven't investigated what would happen if KEKADA knew a lot more phenomena. The question of memory retrieval is likely to become important in such a case. While

KEKADA implements the strategy of relating to a similar phenomenon in a specific form, this strategy in its general form includes analogical reasoning. However analogical reasoning can be a good generator of new experiments which we have not explored in KEKADA.

PRODIGY: While KEKADA uses experimentation only to acquire factual knowledge, experimentation could also be used to refine partially specified operators, and to acquire control knowledge. PRODIGY is a general-purpose planner (Minton&Carbonell, 1987). Gil and Carbonell (1987) describe a number of techniques PRODIGY uses to refine partially specified operators. Another well-known program which uses experimentation to refine its operators is LEX (Mitchell,1983).

AM: AM is a program that can rediscover many concepts in mathematics (Lenat, 1976). It uses an agenda-structure in focussing attention on *interesting* concepts. While the task of AM is not isomorphic to that of KEKADA, there are some similarities in the two programs. For example, while AM uses a number of heuristics to evaluate interestingness of a concept, one of the heuristics recommends increasing the interestingness in an unusual relation between two concept, very similar to KEKADA. AM has also been viewed as carrying out a dual-space search (Barr & Feigenbaum, 1981). The experiments of AM, like the KEKADA experiments in AI, are not on an external environment but on thought material.

KEKADA work provides extensive evidence for the utility of a surprise-focussing strategy to many domains that carry out experiments in an external environment. Thus there might be a number of mechanisms that can be used both to carry out experiments on the external environment and to carry out information processing experiments.

Meta-DENDRAL used the version-space method to infer rules from mass-spectroscopy data. (Buchanan&Feigenbaum, 1978; Mitchell, 1978) Another program, Eurisko(Lenat,1983) can also discover new heuristics as well as make changes in representation. KEKADA does not have the ability to discover new productions or make changes in its representations. Thus these other programs show one natural way to extend KEKADA's abilities.

ADEPT and PHINEAS: Falkenhainer and Rajamoney report on a system that can generate a theory by using analogical reasoning and then carry out experiments to verify that theory. (Falkenhainer&Rajamoney1988) This shows that there are probably many more strategies to generate experiments than those in KEKADA.

BIGTRAK Work: Klahr, Shrager and Dunbar (Shrager&Klahr 1986, Klahr&Dunbar1988) have studied children's discovery of the functionality of a new key in the toy BIGTRAK. They found that children carry out dual-space search. This comes as a further evidence of the utility of dual-space search.

Work in history of science: KEKADA work shows how a detailed historical account of scientists work could be examined to extract domain-independent and domain-specific computer-implementable strategies that scientist used in his research. This shows that the accounts in history of science are informal and incomplete descriptions of information processing involved in science. Computer programs require precision in specifying how information processing transformations are occurring. KEKADA research extended earlier work on STAHL and Thagard's program and others in showing how historical accounts could be used effectively in creating computer programs.

6.5. Directions for future research

There are a number of possible ways in which this research can be extended. Subsections below describe these in detail.

6.5.1. Comparison of plan-tweaking with experimentation

Schank has noted that subjects sometimes create an explanation by plan-tweaking when faced with a surprise. It remains an open issue to understand when an agent should resort to KEKADA experimentation and when it should generate an explanation by the plan-tweaking.

6.5.2. Problem-choosers

We might want to extend KEKADA by replacing the present scheme of choosing a problem by a more elaborate problem-chooser as we discussed in chapter 5.

6.5.3. Creation of discovery assistants

One possible extension to this work will be to create automated research assistants which can be distributed to a research community.

6.5.4. Representation Change

As we saw in this thesis, new phenomena often emerge as surprises in the context of study of some other phenomena. The older representations may not be most efficient to reason about the newly discovered phenomena. Thus when Krebs discovered the ornithine cycle, he continued to represent it in terms of the basic representation of reactions. It was only later when the importance of the peculiar & new

cyclic phenomena was realized, a study was done of the cyclic nature of the phenomena and new representation which made certain features easily accessible emerged. Change of representation in discovery is in general one of the open areas.

6.5.5. Integrated models

A possible next step would be towards integration into a discovery system that has the various experimentation strategies of KEKADA, BACON heuristics for creating new laws and new terms, and ability to reason about causality. It may also have an EURISKO-like ability to make changes in representation. Further extension would allow it to discover new heuristics.

Appendix A

Glossary

Alanine: $\text{CH}_3\text{CH}(\text{NH}_2)\text{COOH}$, 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.

Cysteine: This amino acid has chemical formula $\text{CH}_2(\text{SH})\text{CH}(\text{NH}_2)\text{COOH}$

Cadaverine: $\text{H}_2\text{N}(\text{CH}_2)_5\text{NH}_2$

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 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 $\text{H}_2\text{N}(\text{CH}_2)_4\text{CH}(\text{NH}_2)\text{COOH}$.

Q-rate: mm^3 per hour per mg tissue.

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 average 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.

Appendix B

KEKADA Heuristics in An Ordered Form

B.1. Expectation-Setters

[ES1] Expected lower-bound is set to be the lowest of the 'lower-bounds' of matching expectation-summaries. if the attribute is numeric.

[ES2] Expected upper-bound is set to be the highest of the 'higher-bounds' of matching expectation-summaries. if the attribute is numeric.

[ES3] For an attribute with symbolic value, expected value is set to be the the expected-value associated with the matching summary element.

[ES4] If the previous expectations were not based on experimental evidence and the attribute is numeric, then 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. The expected-value with a symbolic attribute is that observed in the experiment.

[ES5] if the previous expectations were based on at least one experiment, and the present experiment violates the bounds or value in the expectation-summary, update it.

B.2. 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.

B.3. Problem-Choosers

At present we have a computationally simple scheme for choosing a new problem. KEKADA focuses on any surprise it encounters.

[PC1] 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.

B.4. Hypothesis or Strategy Choosers

[HSC-1] Evaluate the alternative hypotheses or strategies and choose one or more of them for consideration using the decision-making rules.

B.5. Experimenters

In the current system, there are no experimentation heuristics.

[E1] The outcomes of experiments are stored in certain working memory elements and are directly copied from them.

DM1 production implements this strategy.

[DM1] The order in which the hypotheses or strategies are considered is based on the user-specified priorities, which closely resemble the program-defined priorities.

Furthermore at times KEKADA needs to choose among alternative members of a class. e.g. it may to have to choose among a number of amino acids. The DM2 heuristic implements this.

[DM2] The substances are stored in the form of an ordered list. (It is assumed that this list has been ordered by cost and availability criteria.)

[HG1] If an independent variable <at> associated with an independent entity of a surprising phenomenon has value <v> and <cl> is a class of values of <at> containing <v>, then consider the strategy to determine the scope of the phenomenon over the range of values specified by <cl>.

[HG2] If substances previously known to influence the phenomenon were absent from the surprising phenomenon, then hypothesize that the absence of such an activator/inhibitor is the causal factor behind the surprise. (Set priority equal to be 2.)

[HG3] If there is a hypothesis that a phenomenon has subprocesses and the phenomenon is noted as surprising, hypothesize that the surprising result depends on one of the subprocesses (divide-and-conquer strategy).

[HG4] If a chemical reaction produces some output (with q-rate or quantity above a minimum threshold), create hypotheses asserting which reactant donates which group to the output substance and if there is more than one reactant, then a reactant may be a catalyst. If the reaction has only one known input and only one output, then guess that there must be other auxiliary inputs among substances around or other auxiliary unknown outputs. Note it is possible that more than one structural formula has been hypothesized about a given input or output. (Set priority to be 6.)

[HG5] If the given phenomenon is a chemical reaction and a one-step reaction from inputs to outputs of a reaction is found not to be possible with only two inputs , then create the hypothesis that an intermediate exists. (Set priority to be 6.)

[HG6] If the goal is to study a puzzling phenomenon and another phenomenon and the surprising phenomenon contain two common dependent or independent entities , then create a hypothesis that the other phenomenon may be related to the surprising phenomenon.

[HG7] If the phenomenon has two or more independent entities, then consider the strategy of deciding whether all entities are necessary to produce the phenomenon.

[HG8] If the surprising phenomenon is a chemical reaction and the input reactant <i1> is hypothesized to have the formula <ii1> or <ii2>, the group <gr> is present in the formula <ii1> but not <ii2>, none of the other input reactants contains the group <gr>, and one of the output reactants contain the formula <gr>, conclude that the reactant <i1> has the formula <ii1> and not <ii2>.

[HG9] If a phenomenon is found to be surprising, then a possible strategy is to attempt to magnify the effect (or make it conveniently visible) by changing the apparatus variables associated with the phenomenon.

[HG10] If the phenomenon under study is a chemical reaction and the incremental q-rate of an output rises unexpectedly on adding an inhibitor, then conclude that this reactant is being consumed in a chemical reaction. (Thus it may be either degrading or reacting with one of the other reactants.) (Set priority to be 6.)

[HG11] Activate any previously known hypotheses about the surprising phenomenon.

[HG12] If one of the variables of the surprising phenomenon is a class, then a possible strategy is to gather more data by carrying out experiments on the various members of this class.

B.6. Experiment-Proposers

[EP1] If the preferred strategy is to assess the scope of a surprising phenomenon over a class of values of an attribute of an independent entity, then use the decision-makers to choose a value A in that class, and decide to study the phenomenon with A as the value of the variable.

[EP2] If there is a hypothesis that the chemical reaction under study contains 2 subreactions PR1 and PR2 one followed by the other, and I1 is the set of inputs to PR1, and I2 the set of inputs to PR2, then study these two reactions, measuring the rates of formation of the outputs.

[EP3] If the preferred hypothesis is that the phenomenon has A and B as 2 independent entities, carry out experiments on A and B in combination and on A and B separately. An independent variable associated with A may be dependent on some variables associated with B and vice versa. In such a case, the experimenter should add additional apparatus to give causal support to the variable.

[EP4] If the phenomenon under consideration is a chemical reaction and the preferred hypothesis is that <in> donates the group <gr> to the <out>, then carry out the reaction making a special effort to measure the rate of consumption of <in> and the rate of formation of <out>.

[EP5] If the chosen hypothesis is that the reactant A in an experiment is a catalyst, then carry out the experiment over long periods but with very low initial quantity and concentration of A. Measure final quantities of all outputs.

[EP6] If the chosen hypothesis is that the reason for a surprising outcome is in the absence of some entity, choose one of the entities that earlier experiments seem to have associated with the given class of processes and study the effects of adding this entity to the independent entities associated with the surprising phenomenon.

[EP7] If the chosen strategy is divide-and-conquer, carry out each of the subprocesses of the phenomenon under various conditions.

[EP8] If the preferred hypothesis is to study the relation of another phenomenon to a surprising phenomenon, then create the following hypotheses and add them to the hypothesis set:

(a) If the surprising phenomenon and the related phenomenon have a common dependent entity <u> and a common independent entity <vna1>; and <vna1> has an attribute <at>; and the surprising phenomenon and the related phenomenon have values of <at>, of <va> and <va2> respectively, a possible strategy is to try to find out if there is a set of values of that attribute containing both <va> and <va2> which will exhibit the same phenomenon.

(b) if the surprising phenomenon and the related phenomenon have a common dependent entity $\langle u \rangle$ and the surprising phenomenon has an independent entity $\langle va \rangle$ and the related phenomenon has an independent entity $\langle va2 \rangle$, a possible strategy is to try to find a set of entities containing both $\langle va \rangle$ and $\langle va2 \rangle$ that will exhibit the same phenomenon.

(c) If the phenomenon is a chemical reaction, the surprising reaction and the related reaction have a common output $\langle u \rangle$ and the related phenomenon has an input $\langle i \rangle$, then create the hypothesis that $\langle i \rangle$ is an intermediate in the surprising reaction.

[EP9] If the chosen strategy is to magnify a certain effect, then carry out experiments varying the value of each apparatus variable over the set of values associated with this variable. Note that a variable may be associated with the method of the measurement of a dependent variable.

[EP10] If the preferred strategy is to gather data about a phenomenon and an attribute of an independent entity is a class, then use the decision-makers to choose a value A in that class, and decide to study the phenomenon with A as the value of the attribute.

[EP11] If the hypothesis under consideration is that reactant $\langle r1 \rangle$ may be involved in an unknown reaction and there exists a hypothesis that $\langle r1 \rangle$ and $\langle v \rangle$ react together and $\langle v \rangle$ is a reactant in the surprising phenomenon, carry out a reaction with $\langle r1 \rangle$ and $\langle v \rangle$ as the reactants.

B.7. Hypothesis and Confidence Modifiers

[HCM-1] If the domain is chemistry And the goal of two of the experiments currently carried out is to study the hypothesis that B is an intermediate in the reaction from ACL to C And these two experiments measure the rates of formation of C from A and from B, And A is the member of the class ACL, modify the implied-success or implied-failure slot in the confidence about the above hypothesis depending on whether there is faster formation from B or from A.

[HCM-2] If the goal of the experiment is to study the hypothesis that the cause of the surprising phenomenon lies in the absence of an independent entity And in the experiment which was just carried out, the entity currently guessed to be missing did not have any effect on the phenomenon; increase "failed-effort" slot in the confidence that an independent entity is missing, by 1.

[HCM-3] If the amount of effort spent on an existential hypothesis reaches a specified high value (which we have assigned to be 3), make the hypothesis inactive.

[HCM-4] If the domain is chemistry ; And the goal of three of the experiments currently carried out is to study the hypothesis that ACL and B react together to form C; And these three experiments measure

the rates of formation of C from A, and from B, and from A and B together; And A is a member of the class ACL; modify the implied-success or implied-failure slot in the confidence about the above hypothesis depending on which of the following two is greater: the rate formation from A and B together, or the sum of the rates from A and from B. (Allow for an error tolerance factor.)

[HCM-5] If the goal of the experiment currently carried out was to assess the scope of a surprising phenomenon; then check the similarity between the resulting phenomenon and the surprising phenomenon and accordingly modify the confidences in the hypothesis guessing the scope of the phenomenon over a class.

[HCM-6] if the goal of the experiment last carried out was to verify whether a hypothesized process description is the correct process description of a surprising phenomenon, And the results of the experiment confirm this,

conclude that this hypothesized process description is the correct description of the surprising phenomenon.

[HCM-7] If there is a hypothesis that the surprising phenomenon may have scope over a class <c> And the success-slot in the confidence of this hypothesis exceeds the threshold value for generalization, generalize the surprising phenomenon and transfer the control to the hypothesis-generators.

[HCM-8] If in a chemical reaction, a small amount of an input can produce large amounts of the output, conclude it acts as a catalyst and there exists an intermediate in the catalytic reaction, and apply HCM-10.

[HCM-9] If the goal of the experiment was to magnify the effect or make it more visible, and it is observed that by changing the value of some apparatus variable the phenomenon is magnified, decide to carry further study of the surprising phenomenon with the new value of the apparatus variable.

[HCM-10] If the preferred strategy is to verify the existence of an intermediate in an experiment , 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) that has been evaluated most likely to be an intermediate in the reaction. If in one of the sub-reactions there is more than one input, conclude that there is yet another intermediate. Ask user to carry out literature survey and recursively apply HCM-10.

Appendix C

A detailed trace of KEKADA

In this appendix, we will give a more detailed description of how the working memory of KEKADA changes when it is run. The description below assumes that the reader is familiar with OPS5.

C.1. The generation of hypotheses and strategies in the face of Ornithine Effect

Below I reproduce a part of the trace of KEKADA from the subsection 4.1.8. I will describe how the working memory in KEKADA changes in the later sections.

Heuristics	Results
HG4*	Possibility that ornithine or ammonia is a catalyst.
HG5	Intermediate exists.
HG7	Creates a clue that mixed action of both the inputs.
HG4*	Hypotheses about who donates what to the reaction.
HG6*	Possibility of relation to similar reactions, for example, to arginine reaction
HG3	Problem may be in one of the sub-reactions of the process
HG1*	Possibility that the phenomenon may be common to a class of substances: namely amino-acids, amines, carboxylic-acids.
HG9	Magnify the effect by varying the apparatus variables.

C.2. Pieces of knowledge in the working memory

I will first describe some of the pieces of knowledge in the working memory, which are relevant to the current run of the program.

; Active group of productions is hypothesis-generators.

9484: (context ^name hypothesis-generation ^active yes)

; Ornithine and ammonia were observed to produce large amount of urea in

; a surprising phenomenon.

9294: (surprise ^name g00064 ^domain chemistry)

9363: (var-of-surprise ^father g00064 ^type independent ^attribute name ^value ornithine ^no 1)

9369: (var-of-surprise ^father g00064 ^type independent ^attribute name ^value ammonia ^no 2)

9361: (var-of-surprise ^father g00064 ^type dependent ^attribute name ^value urea ^no 1 ^expected-value urea)

9362: (var-of-surprise ^parent-var-type dependent ^father g00064 ^type dependent ^attribute q-rate ^value 9 ^no 1 ^expected-lower-bound 2 ^expected-upper-bound 6)

;Structural information about various substances

324: (has-formula ^substance ammonia ^formula ammonia)

279: (formula-substance ^substance-name ammonia ^level 2 ^atom amino ^number 1)

320: (has-formula ^substance ornithine ^formula ornithine)

284: (formula-substance ^substance-name ornithine ^level 2 ^atom amino ^number 2)

285: (formula-substance ^substance-name ornithine ^level 2 ^atom cooh ^number 1).

325: (has-formula ^substance urea ^formula urea)

272: (formula-substance ^substance-name urea ^level 2 ^atom amino ^number 2)

273: (formula-substance ^substance-name urea ^level 2 ^atom carbonyl ^number 1)

319: (partial-match ^group1 amino ^group2 amino)

317: (partial-match ^group1 cooh ^group2 carbonyl)

;Description of a process hh5

82: (process ^name hh5)

83: (var-of-process ^father hh5 ^type independent ^attribute name ^value amino-acid ^no 1)

84: (var-of-process ^father hh5 ^type dependent ^attribute name ^value urea ^no 1)

89: (hyp-process ^name h6 ^part-of hh5 ^priority 4)

91: (hyp-process ^name hhh5 ^part-of h6 ^next-part hh6)

;h4 is the arginine reaction

74: (process ^name h4)

75: (var-of-process ^type independent ^no 1 ^attribute name ^value arginine ^father h4)

76: (var-of-process ^father h4 ^type dependent ^attribute name ^value urea ^no 1)

77: (var-of-process ^father h4 ^type dependent ^attribute name ^value ornithine ^no 2)

;domain-defined taxonomies

- 49: (member ^name ornithine ^attribute name ^order-no 7 ^class carb-acid)
 48: (member ^name ornithine ^attribute name ^order-no 7 ^class amine)
 30: (member ^name ornithine ^attribute name ^order-no 3 ^class amino-acid)

C.3. Changes in the working memory

Below I will describe an OPS5 trace which shows the kinds of changes occur in the working memory. For the sake of brevity, some of the productions are in (watch 1) and the others are in (watch 2).

5737. hypothesis-generation:catalyst:HG4% 9484 9362 9362 9363 9369 9294

=>wm: 9485: (description ^name g00079 ^father g00064 ^type catalyst ^reactant1 ornithine ^priority 6)

=>wm: 9486: (confidence ^name g00079 ^success 0 ^implied-success 0 ^implied-failure 0 ^fail 0 ^failed-effort 0)

;ammonia might be acting as a catalyst

5738. hypothesis-generation:catalyst:HG4% 9484 9362 9362 9369 9363 9294

5739. hypothesis-generation:intermediate-exists:HG5 9484 9294 9363 9369 72 73 71

5740. hypothesis-generation:mixed-action-of-vars:HG7 9484 9294 9363 9369

=>wm: 9491: (description ^name g00082 ^father g00064 ^type mixed-action ^priority 5)

=>wm: 9492: (confidence ^name g00082 ^success 0 ^implied-success 0 ^implied-failure 0 ^fail 0 ^failed-effort 0)

;ammonia might be donating an amino group to urea.

5741. hypothesis-generation:donates-group:HG4% 9484 9362 9362 9294 9369 9361 324 325 279 272 319

=>wm: 9493: (description ^name g00083 ^father g00064 ^type donates-group ^reactant1 ammonia ^reactant2 urea ^group amino ^priority 6)

=>wm: 9494: (confidence ^name g00083 ^success 0 ^implied-success 0 ^implied-failure 0 ^fail 0 ^failed-effort 0)

;Ornithine might be donating an amino group to urea

5742. hypothesis-generation:donates-group:HG4% 9484 9362 9362 9294 9363 9361 320 325 284 272 319

;Ornithine might be donating carbonyl group to urea

5743. hypothesis-generation:donates-group:HG4% 9484 9362 9362 9294 9363 9361 320 325 285 273 317

; the process I0 might be related to the surprising phenomenon

5744. hypothesis-generation:note-related-fact2:HG6% 9484 9294 9361 9363 256 262 259

; Problem may be in one of the sub-reactions.

5745. hypothesis-generation:defect-in-subpart:HG3 9484 9294 9363 9361 30 82 83 84 89 91

=>wm: 9501: (description ^name g00087 ^father g00064 ^type defect-in-subpart ^priority 3 ^process-name h6)

=>wm: 9502: (confidence ^name g00087 ^success 0 ^implied-success 0 ^implied-failure 0 ^fail 0 ^failed-effort 0)

;the process h4 (arginine reaction) might be related to the

; currently found surprising phenomenon

5746. hypothesis-generation:note-related-fact2:HG6% 9484 9294 9361 9363 74 76 77

=>wm: 9503: (description ^name g00088 ^father g00064 ^type related-fact ^fact-id h4 ^priority 6)

=>wm: 9504: (confidence ^name g00088 ^success 0 ^implied-success 0 ^implied-failure 0 ^fail 0 ^failed-effort 0)

;the phenomenon may not be specific to Ornithine, but may be common to ;carboxylic-acids.

5747. hypothesis-generation:scope-rule1:HG1 9484 9294 9363 49

=>wm: 9505: (description ^name g00089 ^father g00064 ^type scope-over-class ^class-name carb-acid ^member-value ornithine ^member-attribute name ^member-no 1 ^priority 4)

=>wm: 9506: (confidence ^name g00089 ^success 0 ^implied-success 0 ^implied-failure 0 ^fail 0 ^failed-effort 0)

;the phenomenon may not be specific to Ornithine, but may be common to amines.

5748. hypothesis-generation:scope-rule1:HG1 9484 9294 9363 48

;the phenomenon may not be specific to Ornithine, but may be common to amino ;acids.

5749. hypothesis-generation:scope-rule1:HG1 9484 9294 9363 30

=>wm: 9509: (description ^name g00091 ^father g00064 ^type scope-over-class ^class-name amino-acid ^member-value ornithine ^member-attribute name ^member-no 1 ^priority 4)

=>wm: 9510: (confidence ^name g00091 ^success 0 ^implied-success 0 ^implied-failure 0 ^fail 0 ^failed-effort 0)

; Attempt to magnify the effect.

5750. hypothesis-generation:magnify:HG9 9484 9294

=>wm: 9511: (description ^name g00092 ^father g00064 ^type magnify-the-effect ^priority 1)

=>wm: 9512: (confidence ^name g00092 ^success 0 ^implied-success 0 ^implied-failure 0 ^fail 0 ^failed-effort 0)

C.4. A firing of a production

How did HG1 get fired in 5749th cycle?

```
(p hypothesis-generation:scope-rule1:HG1
  (context ^name hypothesis-generation ^active yes)
  (surprise ^name <surpr>)
  (var-of-surprise ^type independent ^attribute <at> ^value <val>
    ^no <n> ^parent-var-type <pvt>)
  (member ^name <val> ^class <c1>))
-->
(bind <desc1>)
(make description ^father <surpr> ^name <desc1> ^type scope-over-class
  ^member-attribute <at> ^member-value <val> ^member-no <n>
  ^member-pvt <pvt>
  ^class-name <c1> ^priority 4)
(make confidence ^name <desc1> ^success 0 ^implied-success 0
  ^implied-failure 0 ^fail 0 ^failed-effort 0))
```

Right hand side matches against 9484 9294 9363 30.

9484: (context ^name hypothesis-generation ^active yes)

9294: (surprise ^name g00064 ^domain chemistry)

9363: (var-of-surprise ^father g00064 ^type independent ^attribute name ^value ornithine ^no 1)

30: (member ^name ornithine ^attribute name ^order-no 3 ^class amino-acid)

Rule gets selected by the conflict resolution scheme, and produces the right hand side.

```
=>wm: 9509: (description ^name g00091 ^father g00064 ^type scope-over-class ^class-name amino-acid
  ^member-value ornithine ^member-attribute name ^member-no 1 ^priority 4)
```

```
=>wm: 9510: (confidence ^name g00091 ^success 0 ^implied-success 0 ^implied-failure 0 ^fail 0 ^failed-
  effort 0)
```

APPENDIX D

KEKADA.1 Simulation of the Discovery of Urea Cycle

The Processes of Scientific Discovery: The Strategy of Experimentation

DEEPAK KULKARNI
HERBERT A. SIMON

Carnegie-Mellon University

Hans Krebs' discovery, in 1932, of the urea cycle was a major event in biochemistry. This article 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. Furthermore, 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 article 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 article include the work of Buchanan and others on Meta-DENDRAL (Buchanan & Feigen-

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 Schaffner, 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.

baum, 1978), of Lenat on AM (Davis & Lenat, 1980), of Friedland (1979) on MOLOEN 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 article 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 Kurt 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" (Fru-ton, 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

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.

PROCESSES OF SCIENTIFIC DISCOVERY

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 recollections were recorded. Kekulé 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 be placed 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 genetics.

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 scientists 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 leads 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, personal communication). 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, personal communication). 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 impor-

tance 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 Löffler that “urea formation in the surviving liver is bound up with the integrity of the cell structure” (Löffler, 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

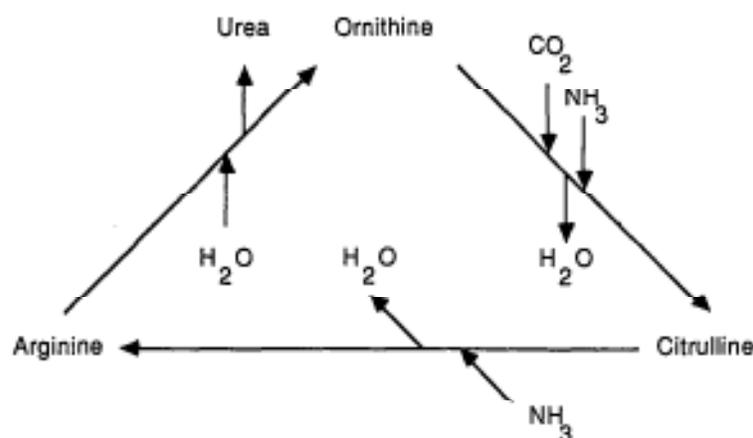


Figure 1. The Ornithine cycle

PROCESSES OF SCIENTIFIC DISCOVERY

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 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

PROCESSES OF SCIENTIFIC DISCOVERY

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, and so forth.

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 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 reactor some amino acids may produce urea" or "ornithine and ammonia produce urea"), (2) its description is given in terms of compound groups. (Example: "NH₂COOH 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

PROCESSES OF SCIENTIFIC DISCOVERY

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 search the instance space fall into two categories:

1. **Experiment-proposers**, which propose experiments based on existing hypotheses.
2. **Experimenters**, which carry out experiments.

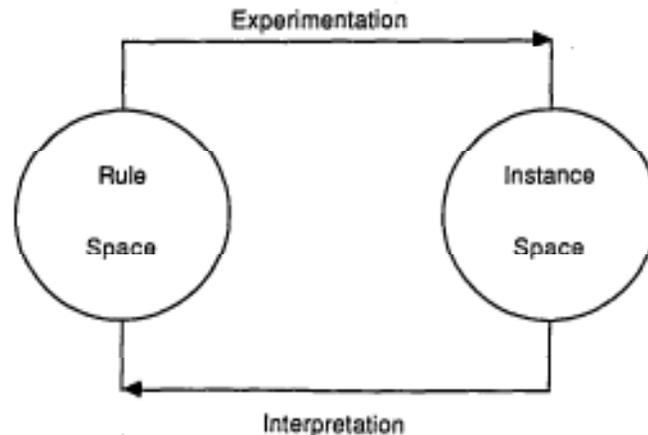


Figure 2. Two-space Model of Learning

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 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.

2.5.1. Interaction of Heuristics. We now can 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 types 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 problems to the agenda and preempt the system to focus on a different problem.

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, Farrell, Kant, & 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 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.

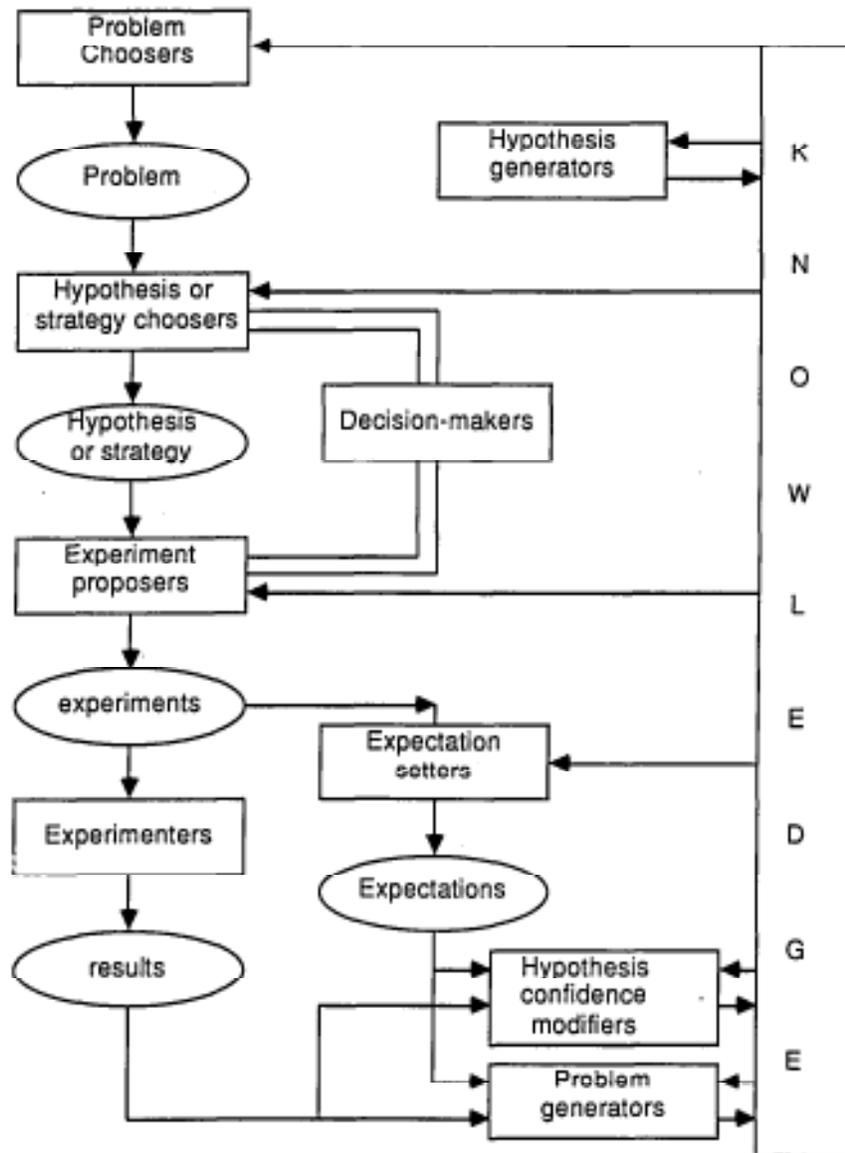


Figure 3. Interaction of heuristics

- [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.

```

; 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. OPSS code for the surprise-detector heuristic

[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 predefined 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 hypothesis

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

PROCESSES OF SCIENTIFIC DISCOVERY

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.

2.14. 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 coefficient 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).

2.15. 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; or 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.

2.16. 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.

2.17. 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 colleagues. 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.

2.17.1. 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, and so forth, 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 supple-

mentary 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.

2.17.2. *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.

PROCESSES OF SCIENTIFIC DISCOVERY

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, <i>decides that intermediate exists.</i>
HG2	<i>Some essential substance is missing from the tissue slice preparation.</i>
HG3	<i>The reason for surprise may be one of the subreactions.</i>
HG1*	<i>The phenomenon may be common to some or all elements of a class.</i>
[seq0]	

[Begin seq0]	
HSC1	Evaluates the alternatives.
DM4,9*	Decides to consider the hypothesis that an absence 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]	

KULKARNI AND SIMON

[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,CF-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]

PROCESSES OF SCIENTIFIC DISCOVERY

[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

Set 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 NH₂ group in urea is ammonia.

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

HM6 Concludes that the source of amino group is NH₂.

HSC1 Evaluates the various alternatives.

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

EP6, 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	<i>Creates hypothesis that there exists intermediate in the reaction.</i>
HM2,B1-11*	<i>Creates candidates for intermediate. Balances the reactions. Counts the number of inputs. Evaluates the intermediates. Chooses arginine.</i>
HG5	<i>Creates a hypothesis that there exists intermediate in the reaction.</i>
(User, when asked to carry out a survey, creates elements corresponding to citrulline and other substances.)	
HM2,B1-11*	<i>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.</i>

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 interpret 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 led 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 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

PROCESSES OF SCIENTIFIC DISCOVERY

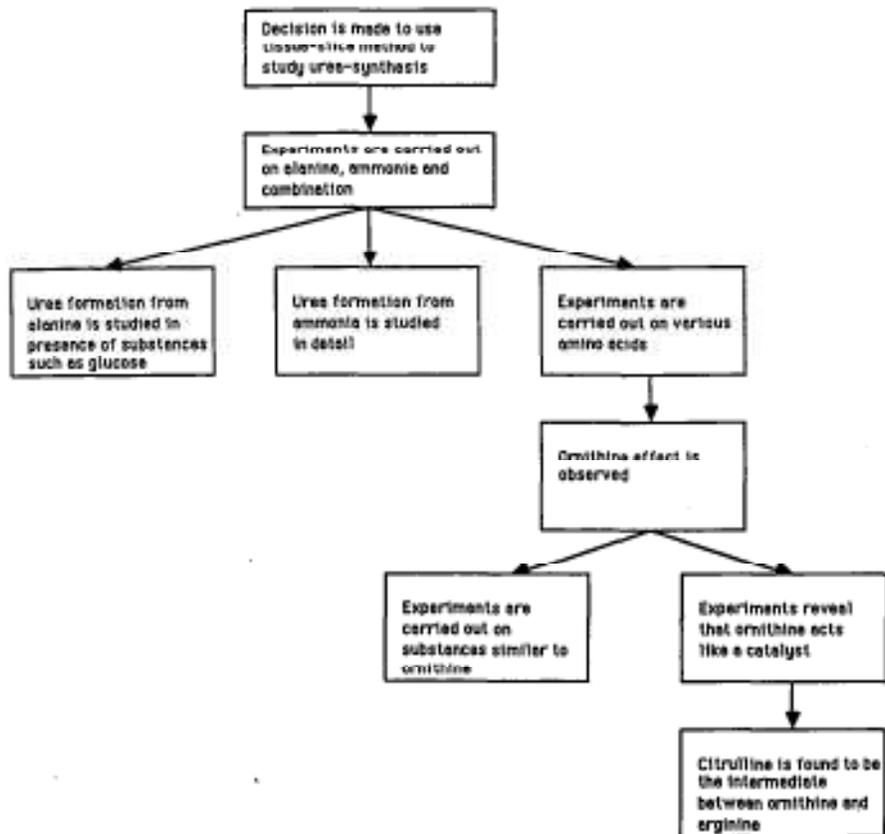


Figure 5. Progress of KEKADA in the discovery

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

PROCESSES OF SCIENTIFIC DISCOVERY

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 tried 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, three 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.

PROCESSES OF SCIENTIFIC DISCOVERY

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

KULKARNI AND SIMON

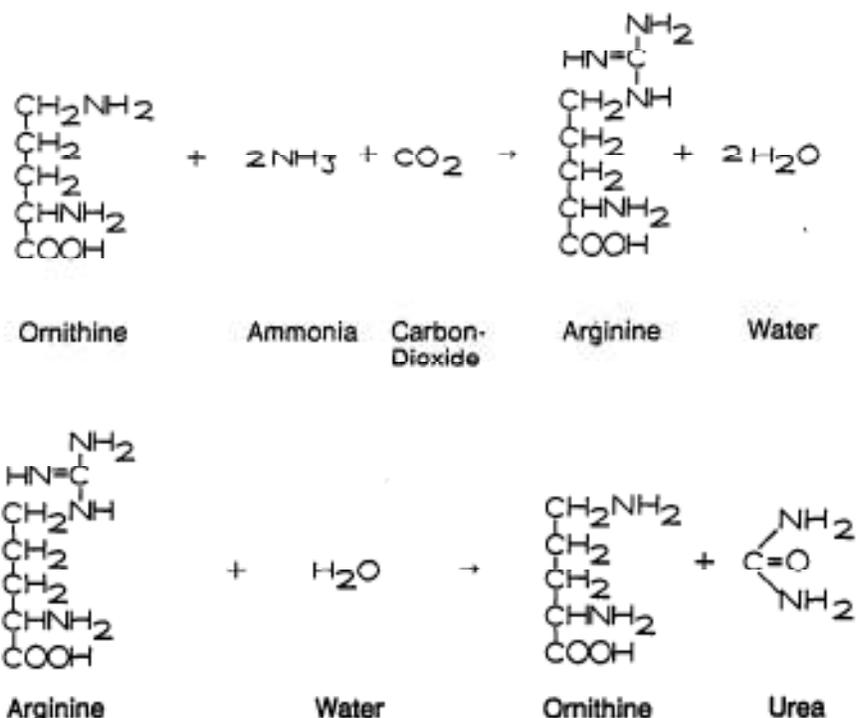


Figure 6. Ornithine as catalyst

carries out an experiment to compare the rate of production of urea from ornithine and from arginine.

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 shifts² 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 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 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, 1947; Fruton, 1972; Holmes, 1986, personal communication; Luck, 1932) that they were in common use in the study of metabolic reactions in biochemistry in the early 20th century, before 1931 and for some years later. Thus, they constituted accepted domain-specific strategies which a newcomer like Krebs was likely to know after a brief introduction to 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 the log in 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).

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

- (4) Thanks to Holmes (1986, personal communication), 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 study.

These considerations show that although KEKADA was handcrafted 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 Hans 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 heuristics 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," that is, 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 additional tests.

■ Original Submission Date: February 5, 1987; Revised July 31, 1987;
Accepted August 15, 1987.

REFERENCES

- Anzai, Y., & Simon, H.A. (1979). The theory of learning by doing. *Psychological Review*, 86, 124-140.
- Baldwin, E. (1947). *Dynamic Aspects of Biochemistry*. New York: Macmillan.
- Brownston, L., Farrell, K., Kant, E., & Martin, N. (1985). *Programming expert systems in OPS5: an introduction to rule-based programming*. Reading, MA: Addison-Wesley.
- Buchanan, B.G., & Feigenbaum, E.A. (1978). Dendral and Meta-Dendral: Their application Dimension. *Artificial Intelligence*, 11, 5-24.
- Davis, R., & Lenat, D. (1980). *Knowledge-based systems in Artificial Intelligence*. New York: McGraw Hill.
- Friedland, P. (1979, August). Knowledge based experiment design in molecular genetics. *Proceedings of the Sixth International Joint Conference on Artificial Intelligence* (pp. 285-287). Tokyo.
- Frustron, 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, 216-225.
- Langley, P., Simon, H.A., Bradshaw, G.L., & 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.
- Löffler, 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.

PROCESSES OF SCIENTIFIC DISCOVERY

- Luck, J.M. (1932). *Annual review of biochemistry*. Stanford, CA: Stanford University Press.
- Shrago, J., & Klahr, D. (1986). Instructionless learning about a complex device: The paradigm and observations. *International Journal of Man-Machine Studies*, 25, 153-189.
- Simon, H.A., & Lea, G. (1974). Problem solving and rule induction: A unified view. In L. Gregg (Ed.), *Knowledge and Cognition*, Erlbaum, Hillsdale, N.J.

I. GLOSSARY

Alanine: $\text{CH}_3\text{CH}(\text{NH}_2)\text{COOH}$, 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 $\text{CH}_2(\text{SH})\text{CH}(\text{NH}_2)\text{COOH}$

Cadaverine: $\text{H}_2\text{N}(\text{CH}_2)_5\text{NH}_2$

Guanidino: The Guanidino group is characterized by $(\text{NH}_2-\text{C}(\text{NH})-\text{NH}-)$. 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 $\text{H}_2\text{N}(\text{CH}_2)_4\text{CH}(\text{NH}_2)\text{COOH}$.

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 average 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.

References

- Anzal, Y. and Simon, H.A. (1979). The theory of learning by doing. *Psychological Review*, 86, 124-140.
- Asimov, I. (1972). *Asimov's Guide To Science* New York:Basic Books.
- Baldwin, E. (1947). *Dynamic Aspects of Biochemistry*. New York:The Macmillan Company.
- Bobrow, D.G. and Hayes, P.J. (1985). Artificial Intelligence- where are we? *Artificial Intelligence* 25, 375-415
- 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
- Carbonell, J.G. and Gil, Y. (1987). Learning by Experimentation. *Proceedings of Forth Machine Learning Workshop*, Irvine, CA.
- Conant J. (Ed.) (1957). *Harvard Case Histories in Experimental Science*, Harvard Univ Press.
- Davis, R. and Lenat, D. (1980). *Knowledge-based systems in Artificial Intelligence*. New York: McGraw Hill.
- Dyer, M., Flowers, M. and Hodges, J. (1986). EDISON: An Engineering Design Invention System Operating Naively. UCLA Technical Report CSD-860087.
- Einstein, A. (1950). *Out of My Later Years*. London: Thames and Hudson.
- Falkenhainer, B. and Rajamoney, S. (1988). The Interdependencies of Theory Formation, Revision and Experimentation. *Proceedings of the Fifth International Conference on Machine Learning*, 353-366
- Forsyth, R. and Rada R. (1986). *Machine Learning: applications in expert systems and information retrieval*. Ellis Horwood Limited. Chichester, England.
- Friedland, P. (1979). Knowledge based experiment design in molecular genetics. *Proceedings of the Sixth International Joint Conference on Artificial Intelligence*, 285-287.
- Futon, J.S. (1972). *Molecules and life: Historical essays on the interplay of chemistry and biology*. New York:Wiley-Interscience.
- Glere, R.N. (1988). *Explaining Science A cognitive approach*. University of Chicago Press.
- Gruber, E. (1974). Darwin on man: A psychological study of scientific creativity. New York: E.P. Dutton & Co.
- 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*.
- Klahr, D. and Dunbar, K. (1988). Dual space search during scientific reasoning. *Cognitive Science*, 12, 1-48.
- Kulkarni, D. and Simon, H.A. (1988). The Processes of Scientific Discovery: The Strategy of Experimentation. *Cognitive Science*, 12,139-175
- Kuhn, T.S. (1970). *The Structure of Scientific Revolutions*, pp62. University of Chicago Press.
- 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
- Leicester, H.M. (1952). *A source book in chemistry*. New York: McGraw-Hill.
- Lenat, D.B. (1983). EURISKO: A program that learns new heuristics and domain concepts. *Artificial Intelligence*, 21, 61-98.
- Lenat, D.B. and Feigenbaum, E.A. (1987). On Thresholds of Knowledge. *Tenth International Conference on Artificial Intelligence* 1173-1182 Milan, Italy: Morgan-Kaufmann.

- Lenat, D.B., Hayes-Roth F., and Klahr P.,(1983).*Cognitive Economy in a Fluid Task Environment*, 133-146 In Proc. of Int. Machine Learning Workshop, June 22-24, 1983 Allerton House, Monticello, Illinois.
- 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.
- MacGibbon,(1973). *Albert Einstein:creator and rebel* London: Hart-Davis.
- Martin, T. (1932). (Ed.) *Faraday's Diary*. London: Bell.
- Medawar, P. (1969). *Induction and Intuition in scientific thought*. American Philosophical Society, Philadelphia.
- Minton, S.N. and Carbonell, J.G. (1987). *Strategies for Learning Search Control Rules: An Explanation-Based Approach*. *Proceedings of IJCAI-87*, Milan, Italy.
- Mitchell, T.M., Utgoff, P. E.and Banerjee, R.B. (1983).*Learning by Experimentation: Acquiring and Refining Problem-Solving Heuristics*. In R.S. Michalski, J.G. Carbonell, and T.M. Mitchell (Eds.) *Machine Learning, An Artificial Intelligence Approach*, Palo Alto, CA: Tioga Press.
- Newton, I. (1931). *Opticks*, pp.1704. London, 1931.
- Poole, R. (1988). *Was Newton Wrong?* , *Science*, Vol 241, pp 789-790.
- Quinlan, J.R. (1983). *Inductive Inference as a tool for the construction of high-performance programs*. In R.S. Michalski, T.M. Mitchell, and J. Carbonell (Eds.), *Machine Learning*. Palo Alto, Calif.: Tiago.
- Rose, D. and Langley, P. (1986). *Chemical Discovery as Belief Revision*. *Machine Learning*, 1, 423-451.
- Rose, D. and Langley, P. (1988). *A Hill-Climbing Approach to Machine Discovery*, In *Proceedings of the Fifth International Conference on Machine Learning*. San Mateo, CA: Morgan Kaufmann.
- 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. (1981). *Sciences of the Artificial*. Cambridge MS. MIT Press.
- 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.
- Schank, R. (1982). *Dynamic Memory*. Cambridge University Press.
- Stranathan, J.D. (1942). *The Particles of Modern Physics* The Blakiston Company, Philadelphia.
- Tweney, R.D. (1987). *A Framework For The Cognitive Psychology Of Science* In B.Cholson (Eds.) *Psychology of Science and Metascience*, Cambridge: Cambridge University Press.
- Tweney, R.D.(1985). *Faraday's Discovery of Induction: A Cognitive Approach* In David Gooding, et al(Eds.) *Faraday Rediscovered: Essays on the Life and Work of Michael Faraday, 1791-1867*. New York: Stockton Press.
- Williams, L.P.(1965). *Michael Faraday, a biography*. New York Basic Books.
- Williams, L.P. (1971). (ed.) *The Selected Correspondence of Michael Faraday*, Cambridge:Cambridge University Press.
- Zytkow, J.M. and Simon, H.A. (1986). *A theory of historical discovery: The construction of componential models*. *Machine Learning*, 1, 107-136.