AD-A067 907 PITTSBURGH UNIV PA LEARNING RESEARCH AND DEVELOPMEN--ETC F/G 9/4
PLANNING NETS: A REPRESENTATION FOR FORMALIZING ANALOGIES AND S--ETC(U)
DEC 78 K VANLEHN, J S BROWN N00014-78-C-0022
UNCLASSIFIED TR-1 NL

| OF |
ADA
067907

END
DATE
FILMED
6-79
DDC

LEVEL II

AD A067907

DDC FILE COPY

PLANNING NETS: A REPRESENTATION FOR FORMALIZING ANALOGIES AND SEMANTIC MODELS OF PROCEDURAL SKILLS

Kurt VanLehn and John Seely Brown

DDC RECEIVED APR 25 1979 C

December 20, 1978

Technical Report No. 1

This research was sponsored by the Personnel and Training Research Programs, Psychological Sciences Division, Office of Naval Research, under Contract No. N00014-78-C-0022, Contract Authority Identification Number, NR 157-408.

This report is issued by the Learning Research and Development Center, supported in part as a research and development center by funds from the National Institute of Education (NIE), United States Department of Health, Education, and Welfare.

Reproduction in whole or part is permitted for any purpose of the United States Government.

Approved for public release; distribution unlimited.

79 04 23 006

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE

READ INSTRUCTIONS BEFORE COMPLETING FORM

1. REPORT NUMBER: Technical Report No. 1
2. GOVT ACCESSION NO.
3. RECIPIENT'S CATALOG NUMBER (9)
4. TITLE (and Subtitle) (6): PLANNING NETS: A REPRESENTATION FOR FORMALIZING ANALOGIES AND SEMANTIC MODELS OF PROCEDURAL SKILLS
5. TYPE OF REPORT & PERIOD COVERED: Technical Report
6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR(s) (10): Kurt VanLehn [redacted] John Seely Brown
8. CONTRACT OR GRANT NUMBER(s) (15): N00014-78-C-0022, MDA903-76-C-0108
9. PERFORMING ORGANIZATION NAME AND ADDRESS
University of Pittsburgh
Learning Research & Development Center
Pittsburgh, PA 15260
10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
Prog. Ele.: 61153N
Proj.: RR 042-06
Task Area: RR042-06-02
Work Unit: NR 157-408
11. CONTROLLING OFFICE NAME AND ADDRESS (11)
Personnel & Training Research Programs
Office of Naval Research (Code 458)
Arlington, VA 22217
12. REPORT DATE: 20 December 1978
13. NUMBER OF PAGES: 51
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) (12) 62p.
15. SECURITY CLASS. (of this report): Unclassified
15a. DECLASSIFICATION/DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT (of this Report)

Approved for public release; distribution unlimited.

(14) TR-1 (16) RR04206

17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)

(17) RR0420602

18. SUPPLEMENTARY NOTES: This research was also supported by the Advanced Research Projects Agency, Air Force Resources Lab, Army Research Institute for Behavioral and Social Sciences, and Navy Personnel Research and Development Center under Contract No. MDA903-76-C-0108; and by Xerox PARC in the Cognitive and Instructional Sciences Group, at Palo Alto, CA.

19. KEY WORDS (Continue on reverse side if necessary and identify by block number)

planning
constraints
semantics of procedures
skill, structure of

20. ABSTRACT (Continue on reverse side if necessary and identify by block number)

At some time in our lives, we have all been forced to learn the procedural skills which supposedly comprise mathematical literacy (e.g., place value addition) through the process of rote memorization, perhaps, enhanced by the use of "models" (e.g., the abacus). These models were intended to provide an intuitive basis for a given procedure. But, what really is a "model" of a procedural skill; how does it help in learning; how faithful can it be made to be; and, more generally, how can it help a procedure take on "meaning?"

DD FORM 1 JAN 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE
S/N 0102-LF-014-6601

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

388 006

Attempting to answer these questions led us to formalize the concept of an analogy between procedures based on a Sacerdoti-like representation called "planning nets." A planning net represents the synthesis of a given procedure from a set of "constraints" that define the properties of the arithmetic operation being implemented and the representation of the objects (numbers) being manipulated. An analogy between procedures is represented as a "maximal partial isomorphism" between the planning nets of the two procedures.

The planning net representation turns out to provide an elegant framework for defining the "teleologic semantics" of a procedure as well as for investigating how to construct a "natural sequence" of models (or microworlds) for a student to use in "inventing" his own procedure. Since both utilize the same framework, we have an extraordinarily powerful way to explain (or teach) the underlying teleology by showing how to relate it to a sequence of intuitively understood models.

This paper may be viewed at several levels: For the educational researcher it provides a framework for investigating the explanatory value of various manipulatory models for mathematical skills; for the cognitive scientist it provides a glimpse at a representation technique for formalizing procedural analogies and for representing the "deep structure" of a procedure; and for the AI person it provides some novel uses of planning nets.

ACCESSION for
NTIS White Section
BDC Buff Section
UNANNOUNCED
JUSTIFICATION
BY
DISTRIBUTION/AVAILABILITY CODES
Dist. SPECIAL
A

Planning Nets:
a representation for formalizing analogies and semantic models of procedural skills

Kurt VanLehn* and John Seely Brown*
December 20, 1978

Abstract

At some time in our lives, we have all been forced to learn the procedural skills which supposedly comprise mathematical literacy (e.g., place value addition) through the process of rote memorization, perhaps, enhanced by the use of "models" (e.g., the abacus). These models were intended to provide an intuitive basis for a given procedure. But, what really is a "model" of a procedural skill; how does it help in learning; how faithful can it be made to be; and, more generally, how can it help a procedure take on "meaning?"

Attempting to answer these questions led us to formalize the concept of an analogy between procedures based on a Sacerdoti-like representation called *planning nets*. A planning net represents the synthesis of a given procedure from a set of *constraints* that define the properties of the arithemetic operation being implemented and the representation of the objects (numbers) being manipulated. An analogy between procedures is represented as a *maximal partial isomorphism* between the planning nets of the two procedures.

The planning net representation turns out to provide an elegant framework for defining the *teleologic semantics* of a procedure as well as for investigating how to construct a *natural sequence* of models (or microworlds) for a student to use in "inventing" his own procedure. Since both utilize the same framework, we have an extraordinarily powerful way to explain (or teach) the underlying teleology by showing how to relate it to a sequence of intuitively understood models.

This paper may be viewed at several levels: For the educational researcher it provides a framework for investigating the explanatory value of various manipulatory models for mathematical skills; for the cognitive scientist it provides a glimpse at a representation technique for formalizing procedural analogies and for representing the "deep structure" of a procedure; and for the AI person it provides some novel uses of planning nets.

To appear in R. E. Snow, P. A. Frederico, & W. E. Montague (Eds.), *Aptitude learning and instruction: Cognitive process analyses*. Hillsdale, N.J.: Lawrence Erlbaum Associates, 1979, and as an Office of Naval Research report from Learning Research and Developement Center, University of Pittsburgh, Pittsburgh, Pa.

*Current address: Xerox PARC, 3333 Coyote Hill Road, Palo Alto, California.

79 04 23 006

Table of Contents

1. Introduction	**3**
1.1 Analogy between procedures	3
1.2 Organizational Overview	5
2. A General Theory of Analogy	**7**
2.1 Mapping between "deep structures"	7
2.2 Basic definitions	8
3. Finding the Right Representation for Analogies	**12**
3.1 Traces	12
3.2 Flow charts	12
3.3 Procedural nets	13
3.4 Planning knowledge seems necessary	14
3.5 Constraints and planning heuristics	16
3.6 Planning a base-1 subtraction procedure	18
3.7 Planning nets	24
3.8 Planning net *mp*-morphisms formalize procedural analogies	26
3.9 Difference generators are used to predict closeness	28
3.10 Discussion	30
4. Analogies and Teleologic Semantics in Educational Research	**32**
4.1 Local explanations: manifestation and motivation	32
4.2 Principles for sequencing local explanations	35
4.3 A space of *mp*-morphisms	39
4.4 Using the *mp*-morphism space in discovery learning curriculums	39
5. Conclusions	**43**
References	**44**
Notes and Appendix	**45**

This research was supported in part by the Advanced Research Projects Agency, Air Force Human Resources Laboratory, Army Research Institute for Behavioral and Social Sciences, and Navy Personnel Research and Development Center under contract number MDA903-76-C-0008, and by the Office of Naval Research under contract number N-0014-78-C-0022.

We would like to thank James Greeno and Al Stevens for various insightful conversations on the nature of procedural skills in elementary mathematics, and Doug Lenat, Austin Henderson, and Richard Burton for their thorough criticisms of early drafts of this paper.

1. Introduction

At some time in our lives, we have all been forced to learn the procedural skills which supposedly comprise mathematical literacy (e.g., place value addition) through the process of rote memorization, perhaps, enhanced by the use of "models" (e.g., the abacus). These models were intended to provide an intuitive basis for a given procedure. But, what really is a "model" of a procedural skill; how does it help in learning; how faithful can it be made to be; and, more generally, how can it help a procedure take on "meaning?"

This paper is directed at understanding how procedures can take on "meaning". It is intended to provide a small step in that direction by discussing a particular kind of "semantics" for procedural skills, which we call *teleologic* semantics, in the context of the unambiguous and totally specifiable procedural skills of elementary mathematics.

The teleologic semantics of a procedure is knowledge about the purposes of each of its parts and how they fit together. Such knowlege is the province of true masters of the procedure. Its value is extolled by the proverb, "To really understand something, one must build it." Teleologic semantics is the meaning possessed by one who knows not only the surface structure of a procedure, but the details of its design.

This paper has two arguments. First we motivate the particular representation that we use for teleologic semantics, which we call *planning nets*, by showing how it can capture analogies between procedures as seen by an expert at those procedures. Secondly, we show that teleologic semantics, as formalized by planning nets, is useful by describing several potential applications in the field of education. In particular, some consideration is given to how teleologic semantics can be explained, and how it provides a useful framework for developing "optimal" sequences of "model" procedures (or microworlds) for guided discovery learning.

1.1 Analogy Between Procedures

Before we delve into a technical discussion of procedural analogies, let us consider a simple example of an analogy between the procedure for adding two multi-digit numbers and a "model" procedure for addition which manipulates physical objects that represent numbers. The model procedure is a physical procedure in that it manipulates physical objects that stand for numbers. Before we can describe the procedure, we will briefly describe the objects that it manipulates, namely Dienes Blocks.

Planning Nets

The Dienes Blocks representation of numbers

Dienes Blocks provide an explicit representation of base 10 numbers -- namely a set of *unit* blocks for representing the units; a set of *long* blocks consisting of ten unit blocks molded into a long stick for representing the tens; a set of *flat* blocks consisting of ten long blocks laid next to each other, thus forming a 10 x 10 square for representing the hundreds; and finally a set of *cubes* in the form of 10 x 10 x 10 units for representing the thousands. A number (of 4 or less digits) can be physically represented by selecting the number of unit blocks to correspond to the units digit, the number of long blocks to correspond to the tens digit and so on. Hence a particular multi-digit number is represented by piles of units, longs, flats, and cubes. Here, for example, is 123 represented in Dienes Blocks:

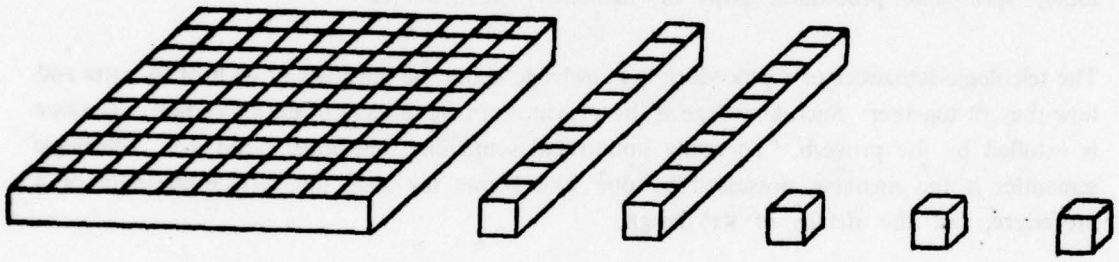

The base-10 nature of the symbolic place-value scheme for representing numbers is then made explicit since one can see the direct translation of a number represented as piles of Dienes Blocks into a base-1 system (i.e. the total number of units comprising all the blocks in all the piles).

Dienes Block Addition

Addition of two multi-digit numbers represented as concrete Dienes Blocks involves forming set unions, and "trading". The units pile for each of the two numbers is first unioned together. This corresponds to adding the units column. Next, the resulting set is examined. If it contains more than ten unit blocks, then ten blocks are removed from this set and traded for a long block (consisting of ten units) which is then placed in a pile of long blocks of the top number. This corresponds to carrying from the units to the tens column in standard addition. The procedure now repeats, unioning the longs piles, then the flats, etc.

A theory of analogy between procedures, applied to this case, should be able to capture not only the fact that Dienes Block addition and standard additon produce the same answers given the same inputs, but that their *internal structure* corresponds as well. Set unions match with column

sums, trading matches carrying, and so on.

Two-pass addition illustrates differences in closeness

We were recently struck by the way Dienes Blocks were being used in a school. In particular, the Dienes Blocks procedure being taught was not as described above but instead had the students combining all the piles of blocks together and then returning to the units pile and trading up and so on. Thus, in standard multi-digit addition, a carry is (potentially) performed after each column operation, whereas in this version of Dienes Block addition the "trading" (or carrying) operation was being deferred until all the columns have been initially processed. One intuitively feels that this second, two-pass procedure is not as closely analogous to standard addition as the previous, one-pass Dienes Block procedure.

A theory of analogy should have some formal measure that can predict how close an analogy is. The theory below has such a formal mechanism, called a *closeness metric*. The degree of correlation between the predictions of the closeness metric and subject's intuitive judgements of closeness is one verification condition for the theory. (See note 1 for caveats on this claim.)

Why arithmetic?

The examples in the paper are all drawn from the computational procedures of arithmetic even though the techniques we have developed have wide applicability. We limited our examples to arithmetic for several reasons. Everyone knows how to add and subtract, so lack of familiarity with the example domain will not hinder comprehension of these admittedly rather abstract formalisms. Arithmetic is a highly evolved, complex system of procedures. It has iteration, recursion, tables of facts, and of course a rather non-trivial data representation, namely place-value numbers. Lastly, arithmetic is taught in school. This means our theories are more likely to accrue the benefits of thoughtful, experience-based criticism from those with a sincere interest in putting the theories to work.

1.2 Organizational Overview

The paper is divided into three parts. The first part (section 2) is an exposition of some of the basic concepts of formal theories of analogy. We assume that an analogy can be represented as a *mapping* between a *deep structure* representation of each procedure which is expressed as a *maximal partial isomorphism* between the two deep structures. Thus, after an analogy has been comprehended, we would expect to find cognitive structures that could be modeled by three components: two of which represent the abstraction or deep structure of the two procedures and the third which represents the structure perserving map (i.e., analogy) between these two structures.

Planning Nets

The second part of the paper (section 3) motivates the planning net representation of teleologic semantics by using it as the deep structure component of a theory of analogies between procedures. Part three of the paper (section 4) is an examination of some of the applications of this theory to education. In particular, we discuss a paradigm for explaining the teleologic semantics which involves using a sequence of analogies such that each analogy illustrates exactly one concept underlying the synthesis of the given "target" procedure (e.g., place-value subtraction). This paradigm is then augmented with a set of "naturalness" principles for structuring a sequence of analogies thereby addressing the problem of how to design an optimal sequence of "micro-worlds" or models for enchancing discovery learning.

We caution the reader that our style of arguing with examples has led to the incorporation of a great deal of detail into the subsequent pages. However, if Artificial Intelligence has contributed anything to cognitive psychology, it is an appreciation that ignoring trivial detail often leads to overlooking non-trivial problems.

2. A General Theory of Analogy

This section presents a theory of analogy that is so general that it is almost vacuous. It appears that virtually any theory of analogy, including the theory of procedural analogies that is presented below, can be recast as a special case of this general theory. Thus, this general theory is apparently immune to refutation. Nonetheless, it allows discussion of some concepts common to all analogies, such as "closeness", before becoming immersed in the details of procedures and their representations.

2.1 Mapping between "deep structures."

We view an analogy as a comparison of two "things" that can be broken down into three parts: (a) an analysis of the first thing into some abstract description (or deep structure), (b) an analysis of the second thing into another abstract description, and (c) a mapping between the two descriptions. This tripartite breakdown is the foundation of the general theory of analogy. Exactly this breakdown is also found in Tversky's work on similarity, a domain which illustrates the general theory more clearly because of the simpler "deep structures" that are used (Tversky 1977).

Much research on similarity has used pairs of geometric figures or letters. A typical task is to rate the similarity "o" to "c". Tversky's analysis of this task is to assume a feature space, describe each figure as a set of features, then predict the similarity judgements with some "metric" on the overlap of the feature set of "o" with the feature set of "c". The correlation of the judgements with the predictions serves as a verification condition on the feature space and the metric. Often, the features are not very abstract -- "o" might be mapped into the description {curved, circular, closed} while "c" would become {curved, circular, open}.

Much of the research on analogy has studied a task one often finds on intelligence tests, namely, to fill in X in a statement of the form "A is to B as C is to X." Most commonly, the four elements A, B, C and X are either words or geometric figures. A simple example of a word analogy problem is "Red is to Stop as Green is to (a. Go, b. Halt)." Superficially, this appears to be a different sort of task than the similarity task since there are four things rather than two. But the two tasks become very much the same when one considers the analogy task to be a comparison of *relationships* rather than directly apprehendable things. This is a widely held view of analogy. Indeed, the instructions to one analogy test, as quoted by Evans (1968, pg 272) read "Find the rule by which Figure A has been changed to make Figure B. Apply the rule to Figure C. Select the resulting figure from Figures 1 to 5."

Actually, these instructions represent just one strategy for answering analogy problems. Evans ANALOGY program, for example, used a different strategy, whereby it extracted an AB rule, then

found five rules for pairs C1, C2, C3, C4 and C5, then finally chose one rule of the five as being the most similar to the AB rule. The existance of many different strategies for solving analogy problems also obscures the parallels of this task to the similarity and metaphor tasks. And yet, when one is done finding the analogy, one possesses the same three maps: an abstraction from AB, an abstraction from CX where X is the chosen answer, and the partial match (or mapping) between the two resulting abstract descriptions.

In short, if one ignores the strategic differences between solving an analogy and evaluating a similarity, and one puts relationships on an equal footing with letters and geometric figures, then there is very little difference between the analogy task and the similarity task. After either task is completed, the cognitive structures can be modeled by three components: the two abstract descriptions and the mapping (in the form of a match) between them.

2.2 Basic definitions

In this subsection, several basic concepts will be discussed. They all follow rather immediately from the view of analogy described above. As above, they will be motivated and illustrated with examples from Tversky's theory of similarity.

Intersection and difference sets

A good way to summarize the outcome of the matching map is in terms of one *intersection set* and two *difference sets*. As an example, take the similarity task mentioned above, to evaluate the similarity of "o" and "c". Their descriptions, let's say, are the feature sets {round, curved, closed} and {round, curved, open}, respectively. Call these sets A and B, the abstract descriptions of "o" and "c". Then, the intersection and difference sets are

$$A \cap B = \{ \text{round, curved} \}$$

$$A - B = \{ \text{closed} \}$$

$$B - A = \{ \text{open} \}.$$

This is not particularly startling, to be sure. Indeed, there are stereotypical ways of referring to these sets in English similies: "A is like B in that $A \cap B$," or "A is like B except that A-B instead of B-A."

Maximal partial graph morphisms generalize the notion of "match"

With more complex languages than feature spaces for expressing abstract descriptions, one must of course give a new definition of "match." For example, consider the analogy (from Sternberg 1977) "Washington is to 1 as Lincoln is to 5." Suppose semantic nets are the representation

language. The abstract description of the relationship Washington:1 is a certain chain of semantic links from the node "Washington" to the node "1". The description of Lincoln:5 is a different chain. However, when one finally finds the correct way to view the two relationships (which is rather non-trivial for this example), then the two chains end up bearing the same sequence of link names, namely: Last-name, image-of, portrait-on, dollar-amount. That is, "Washington" is the *last name* of the man NODE31, the *image of* NODE31 appears in the picture NODE7, NODE7 is the *portrait on* the kind of dollar bill NODE68, and the *dollar amount* of NODE68 is "1". The chain for the Lincoln:5 relationship is a completely distinct chain, but it has exactly the same sequence of link labels. In this sense, the analogy is perfect.

To make these two chains match, the definition would have to be sensitive to (a) the order of the links, and (b) the labels on the links. A definition in terms of intersection of sets of links would be inappropriate because none of the links are identical, and because such a definition would ignore the topology of the descriptions. A definition of "match" that is appropriate for semantic nets (or any other representation with the topology of a labeled directed graph, including planning nets), can be defined in terms of a graph isomorphism:

> *Adjacency:* Two links of a graph are adjacent if they are incident with a common node.
>
> *Isomorphism:* An isomorphism of labeled directed graphs is a 1-1 correspondence on the links that preserves the adjacency, direction and label of the links.

The "match" of the two semantic net chains X and Y can now be defined to be the *maximal* graph isomorphism from a subgraph (subsequence) of X to the subgraph of Y. By "maximal", we mean the isomorphism that pairs the largest number of links correctly. Unfortunately, use of maximality precludes any mathematical guarantee of the uniqueness of the resulting isomorphism. However, in practice we have yet to be plagued by a non-unique maximal isomorphism.

Note that we have defined "match" as a map which is an isomorphism between *sub*graphs of the two deep structures. The map between deep structures is not necessarily total (i.e., onto) in either direction (we are in the process of investigating a revision of this aspect of the definition as well as the interesting situation where it is many-to-one and hence would have the properties of a homomorphism). In other words, the analogy is a mapping which is a *maximal partial graph isomorphism.* However, we will abbreviate our terminology somewhat and say that the analogy from A to B is formalized by the mp-morphism from A to B (i.e., we will speak of the analogy as being this structure preserving map.)

To replace the terms "intersection set" and "difference set", we will simply use "intersection subgraph" and "difference subgraph". There are, of course, two difference subgraphs for a mp-morphism, namely the residue portions of each of the deep structures being compared.

Throughout this report, we will continue to use the symbology of sets for these concepts, even though the designated entities are not sets, but subgraphs.

Closeness metrics

Both the similarity task and the analogy task involve the ranking of the match between two things, or rather between their abstract descriptions. The subject is asked to rank the degree of similarity or choose the closest analogy. We assert that both kinds of judgements can be modeled by a function over the intersection set (or subgraph) and two difference sets (or subgraphs). In similarity research, this three argument function is often called a "similarity metric" even though there are cases when the function is not a proper mathematical metric (see Tversky 1977). With the same sloppiness, we will call the function that ranks the closeness of analogies a *closeness metric.*

These metrics can be rather complex. Certain features might be more salient than others, and one might model this difference by giving the former more weight in a summation over the various sets. These metrics might even be asymmetric (see note 2) which means they are not proper "metrics" in the strict mathematical sense. In short, determining the intersection set and the two difference sets is not the end of the story for predicting similarity judgements; the metric can play a decisive role.

Monotonicity, etc.

We take the position that a precise statement of the closeness metric for procedural analogies can only be determined from detailed empirical studies. However, Tversky has shown that if certain formal conditions on the metric can be guaranteed, such as its monotonicity over subsumption of the intersection and difference sets, then the metric can have a simple, linear form (Tversky 1977). (One of us--VanLehn--has investigated some of the conditions for procedural analogies, and will discuss them in a later report.)

Individual differences and learning

We have been speaking of *the* abstract description (or deep structure) of a thing as if this object is the same for all people. In some tasks, such as assessing the similarity of letters, it seems reasonable for literate individuals to have roughly the same representation language and the same abstraction functions for extracting descriptions from the letters. But this assumption is rather implausible in many other cases. In these cases, individual differences in conceptions of the things being compared is likely to influence judgements of the closeness of analogy. This would make verification of a theory significantly more difficult.

Individual differences affect analogy, but analogy also affects the individual differences. That is, one can learn from analogies. More specifically, when an individual understands an analogy, he or she may become aware of descriptive features that they were not previously aware of. So, a complete theory of analogy must allow for an evolution of an individual's conception of the things being compared over the course of testing.

In this research, we will ignore these difficult methodological problems by assuming that the subjects who are judging the closeness of the analogies are *experts*. That is, they all have a complete representation of the things being compared, and hence can be assumed to have roughly the same representations, and secondly, that they already know all there is to know about the things being compared, and therefore learn very little over the course of the testing.

3. Finding the Right Representation for Procedural Analogies

In this section, several candidate representations for procedures will be examined as a basis for a theory that predicts the closeness of procedural analogies (note 3). Possible representations range from a very superficial one, namely a simple chronological list of actions, on up to a very abstract representation that involves goals, constraints and other planning knowledge, namely planning nets. Our research has shown that planning nets are the only serious contender, so the discussion of the others will be quite brief. However, the more superficial representations are mentioned in this section for a reason, namely, to show how a human (or machine) can construct a very abstract representation of a procedure by ascending through several levels of representation. We do not claim that the structure of this section models the abstraction process that a person executes when assimilating a procedural analogy, but it does provide an indication of the complexity that such a process would have to have.

3.1 Traces

The *trace* of a procedure is simply a chronological list of the actions it performed during one particular execution. This representation of a procedure can be constructed directly from observation of the execution of the procedure (although, there are the usual problems in choosing the "grain size" of primitives -- see note 4). However, traces are a highly inappropriate representation for procedures, as the following example indicates.

Consider an analogy between Dienes Block addition and written addition. These two traces would probably have few, if any, action labels that match. The action "write "4"" would have to be matched against a group of four actions labeled "place one block in the pile", whereas the action "write "8"" would have to be matched against a group of eight block placing actions. Such sophisticated matching could not be represented by a mp-morphism. Indeed, the match seems to require the *concept* of "write *n*" and the *concept* of "repeat single block placement *n* times." These are abstractions over action sequences, and so should be part of the representation rather than the matching mechanism. Incorporating such concepts into the representation lifts us to the next level of abstraction.

3.2 Flow charts

By generalizing over a large collection of traces, one could derive a notion of the observed procedure that could be represented with a programming language, such as flow charts. Granted, this generalization would be non-trivial: repetitious sequences of actions would become loops, objects that are manipulated similarly become the contents of variables, etc. Nonetheless, constructing a program from examples is well within human ability.

However, flow charts would also be a poor representation for analogy. Consider a simple subtraction procedure for numbers represented as base-1 blocks as illustrated by the flow chart below. The primitive terms used in this flow chart are as follows. LH stands for someones's left hand. TOP and BOT stand for placemats on the TABLE. The BOT set of base-1 blocks is subtracted from the TOP set of blocks by pairing off a block from each, using the primitive actions PICK/FROM and PUT/ONTO, and tossing them onto the table. When the bottom "number" is "zero" (i.e., empty), whatever is left in the top "number" is the answer. However, notice that by merely shuffling the order of the steps somewhat and using two hands instead of one, a new procedure can be constructed that is intuitively very similar to the old procedure, and yet its flow chart (see below) shares virtually no isomorphic subgraph with the old procedure's flow chart. Since the intersection graph is so small relative to the difference subgraphs, a reasonable closeness metric would have to report that the two procedures are not very close -- a false prediction. So for this and other reasons, flow charts also seem to be a poor representation or level of abstraction for procedural analogies.

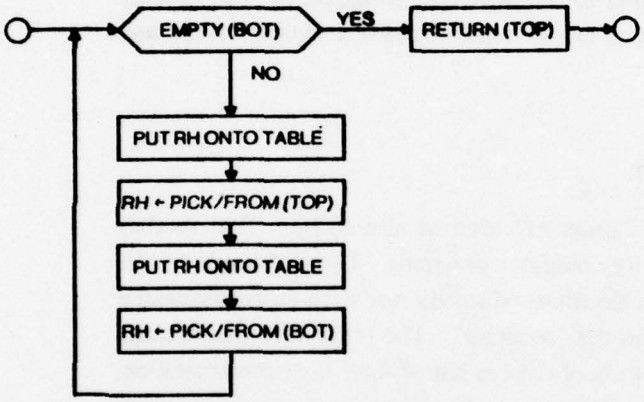

FLOW CHART FOR A BASE-1 BLOCK SUBTRACTION PROCEDURE USING ONE'S RIGHT HAND

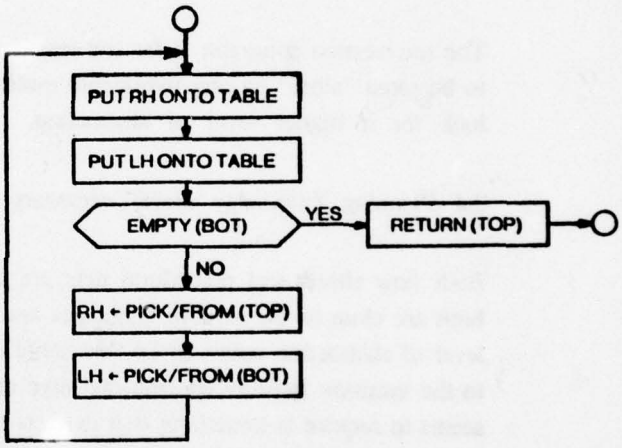

FLOW CHART FOR A BASE-1 BLOCK SUBTRACTION PROCEDURE USING TWO HANDS

3.3 Procedural nets

On the basis of the example above, it might appear that flow charts are too committed to a set order of performing steps, since the two base-1 flow charts have the same steps, but order them slightly differently. Also, these charts lack the typical hierarchy of subprocedures that is used in computer programs to modularize and organize the procedure. This suggests using a structure that emphasizes the subprocedure hierarchy, and deemphasizes the temporal sequence of subprocedures.

Just such a structure has been developed for modeling children's bugs in arithmetic procedures -- namely BUGGY's *procedural net* representation (Brown and Burton 1978). Although we will not pause here to explain this representation, a procedural net for a very familar procedure, namely standard subtraction, is included as figure 1. However, procedural nets also fail as a basis for a theory of analogy, as illustrated in the following example.

Consider two Dienes Block subtraction procedures: (A) In "big-pile" Dienes block subtraction, a number is represented by one big pile of Dienes Blocks. (B) In "sorted" Dienes Block subtraction, all the blocks are kept sorted into little piles according to their shape. Intuitively, these two procedures are quite closely analogous. But when the procedural nets are formed, and the matching is done, we find the following statistics:

A∩B contains 6 nodes.
A-B contains 10 nodes.
B-A contains 16 nodes.

The intersection subgraph is far too small compared to the difference subgraph for this analogy to be rated "close" by any reasonable metric. So again, we must abandon a representation, and look for a higher level of abstraction.

3.4 Planning knowledge seems necessary

Both flow charts and procedural nets are at the "program" level of abstraction. That is, they both are close to the sorts of languages one sees for computer programs. The problem with this level of abstraction seems to be that some design decisions which do not seem so consequential to the intuition have an enormously large effect on the "program". The framework that analogy seems to require is something that extracts these sorts of choices out of their final manifestation, makes them explicit, and relates them in a reasonable way to other, more important elements of the design. In short, what seems necessary *is a representation of the design process* behind a procedure -- this allows one to say which choices are important, and which are relatively minor. The process of creating a procedure from a set of constraints is traditionally called "planning" by the Artificial Intelligence community. So, the abstract representation that analogy seems to require appears to involve planning knowledge and planning inferencing.

Planning knowledge includes not only the functional decomposition of the surface structure of the procedure but also the reasoning that was used to transform the goals and constraints which define the *intent* of the procedure into its actual surface structure. The formalism that we use to represent this knowledge we will call *planning nets*. These planning nets are an extension of Sacerdoti's pioneering work on representing procedural knowledge for robotics (Sacerdoti 1977).

Planning Nets

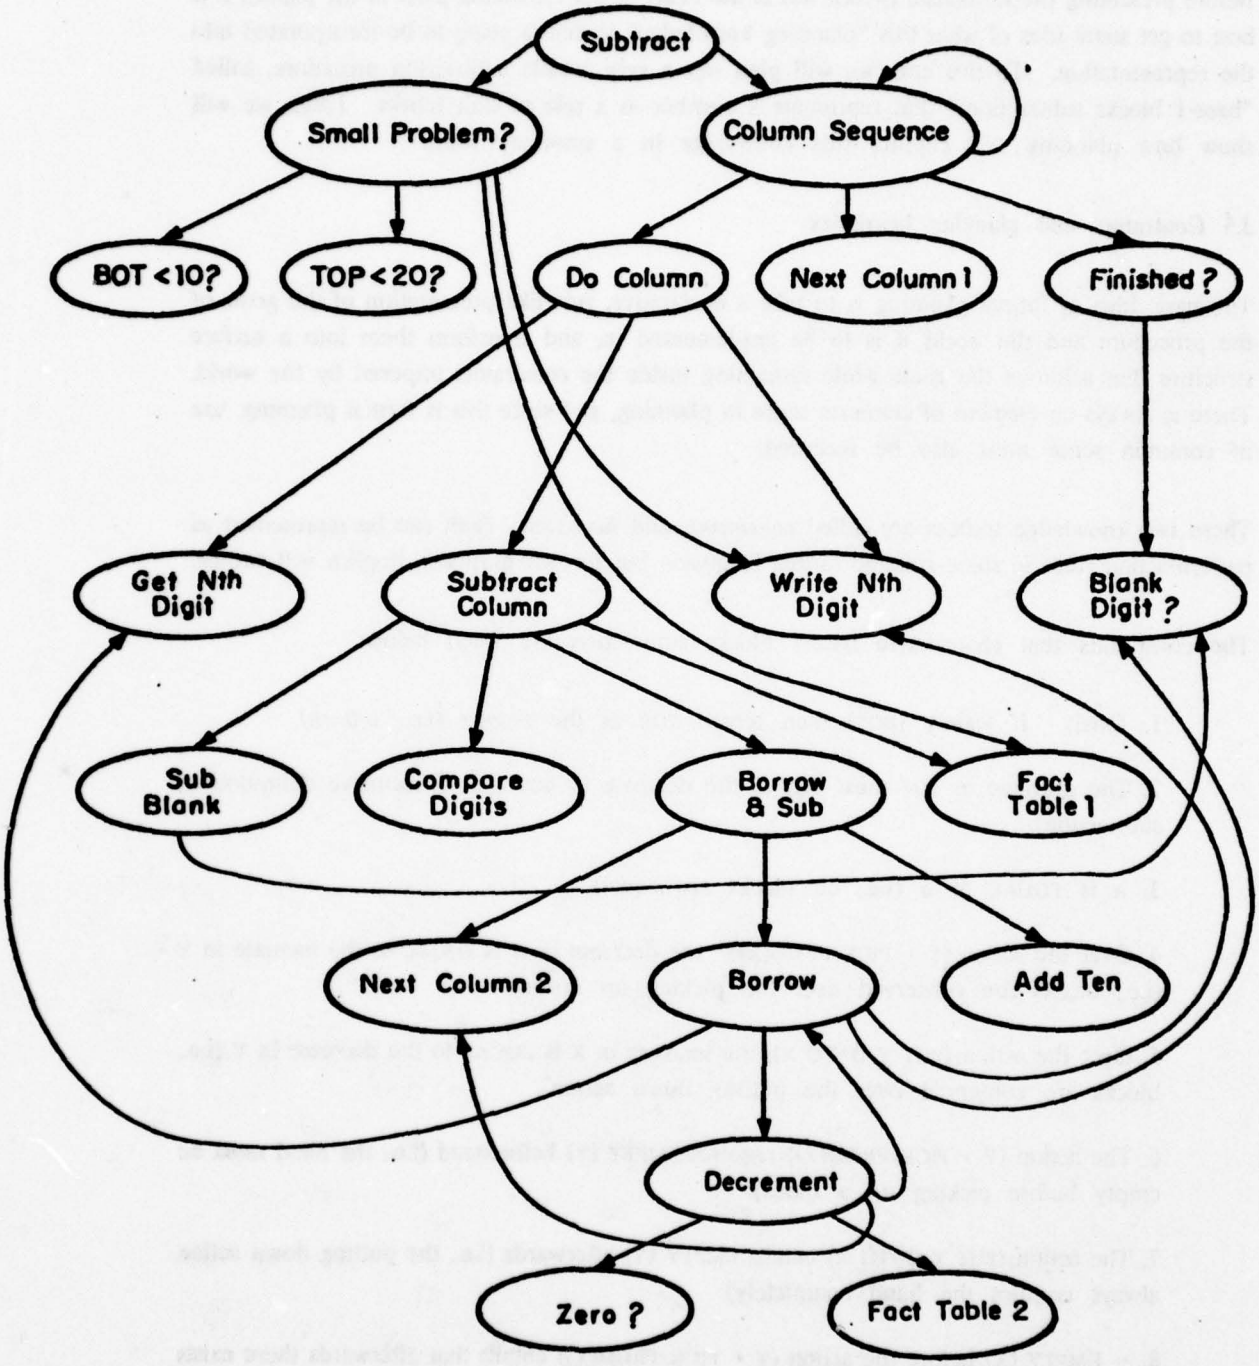

Figure 1

Before presenting the formalism (which lies at the heart of the remaining parts of the paper), it is best to get some idea of what this "planning knowledge" is that is going to be incorporated into the representation. To this end, we will plan out a very simple subtraction procedure, called "base-1 blocks subtraction," that represents a number as a pile of unit blocks. Later, we will show how planning nets capture this knowledge in a summary form.

3.5 Contraints and planning heuristics

The basic idea of formal planning is to take a declarative, rule-like presentation of the goals of the procedure and the world it is to be implemented in, and transform them into a surface structure that achieves the goals while remaining inside the constraints imposed by the world. There is always an element of common sense in planning, and since this is formal planning, use of common sense must also be recorded.

These two knowledge sources are called *constraints* and *heuristics*. Both can be represented as pattern-action rules in some suitable formal language, but for our purposes, English will suffice.

The constraints that characterize base-1 blocks subtraction are listed below:

1. Goal: If EMPTY (BOT) then return TOP as the answer (i.e., $n-0=n$).

2. The decrease in TOP must EQUAL the decrease in BOT (i.e., a recursive definition of subtraction).

3. a is EQUAL to b (i.e., all blocks are equal).

4. Over the action (Y ← PICK/FROM(X)), the decrease in X is EQUAL to the increase in Y (i.e., blocks are conserved over the picking up action).

5. Over the action (PUT Y ONTO X), the increase in X is EQUAL to the decrease in Y (i.e., blocks are conserved over the putting down action).

6. The action (Y ← PICK/FROM(X)) requires EMPTY (Y) beforehand (i.e., the hand must be empty before picking up a block).

7. The action (PUT Y ONTO X) entails EMPTY (Y) afterwards (i.e., the putting down action always empties the hand completely).

8. ~ EMPTY (X) before the action (Y ← PICK/FROM(X)) entails that afterwards there exists a, such that a is the contents of Y. (i.e., the hand picks up exactly one block).

The meaning of the primitives is as follows. TOP and BOT are placemats on the TABLE. The subtraction problem *n-m* would begin with *n* base-1 blocks on TOP, and *m* on BOT (n.b., this is not the way base-1 block subtraction is ordinarily posed in the classroom. See note 5). There are two hands, LH and RH, which can perform two kinds of actions, namely picking up one block (PICK/FROM) or putting a block being held down (PUT/ONTO). The primitive predicate EQUAL takes two piles of blocks and says whether they designate the same number. EQUAL is not executable, and can not appear in the final plan.

The constraints above describe the mathematical goals of the procedure, the objects it works with, and the physical manifold that it operates within. The mathematical content of subtraction is expressed in constraints 1 and 2: TOP minus BOT is TOP whenever BOT is empty of blocks, but any changes in the number of blocks on BOT must be echoed by an equal change in the contents of TOP. The objects the procedure manipulates are base-1 blocks. Since these are very simple, constraint 3 suffices to describe them. (By convention, a lower case letter stands for an arbitrary block, while an upper case letter stands for an arbitrary placemat or hand.) The remaining constraints define the physical manifold that the procedure will operate within. Constraints 4 and 5 ensure that blocks are *conserved* by the actions PICK/FROM and PUT/ONTO. Constraints 6, 7 and 8 describe how the hands that manipulate the blocks work. A complete description of the workspace would require several more constraints, but these will do for purposes of illustration (for some comments on how this particular set of constraints was chosen, see note 6).

The constraints describe *domain-dependent* knowledge. If the procedure's goals or implementation environment change, then the constraints must be changed to reflect this. For example, if one used Dienes Blocks instead of base-1 blocks, then constraint 3 would be replaced by a new constraint, namely

3'. a is EQUAL to b if and only if SHAPE(a)=SHAPE(b).

If one wished to plan an addition procedure instead of a subtraction procedure, then constraint 2 would become

2'. The increase in TOP must EQUAL the decrease in BOT.

Heuristics are presupposed to be *doman-independent* knowledge. They represent common sense planning knowledge, such as "when you need to accomplish two things, and it doesn't matter which comes first, then pick one arbitrarily, do it first, then the other." We include this distinction between constraints and heuristics only because it is traditional; nothing in our theory turns on this distinction.

3.6 Planning a base-1 subtraction procedure

The planning of the base-1 subtraction procedure involved 12 steps. Each step is an application of a constraint or a planning heuristic. The planning begins with a flow chart initialized to the constraint that is marked as the "goal" of subtraction.

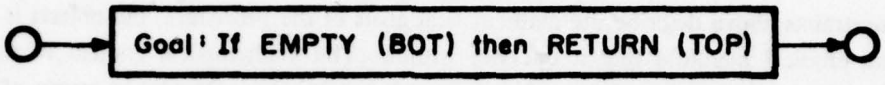

Planning proceeds by progressive refinement of goals to subgoals, or by checking the current plan against the constraints. (n.b., Because we are only interested in *having* a correct planning net for a procedure, not in *finding* one, we are going to ignore a few of the subtle issues -- see note 7).

Planning Nets

Step 1: At the outset the Implication Reduction planning heuristic which reduces an implication, (A⊃B) to a sequence of subgoals, (A,B) can be applied. The second subgoal in this case is a primitive of the workspace. So the output of Step 1 is a plan with just one subgoal:

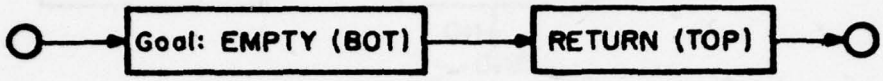

Step 2: A venerable planning heuristic, traditionally called Hill Climbing (Newell and Simon 1972), reduces the goal to a loop. The loop test sees if the goal has been achieved, and if not, it takes a step "up the hill," so to speak.

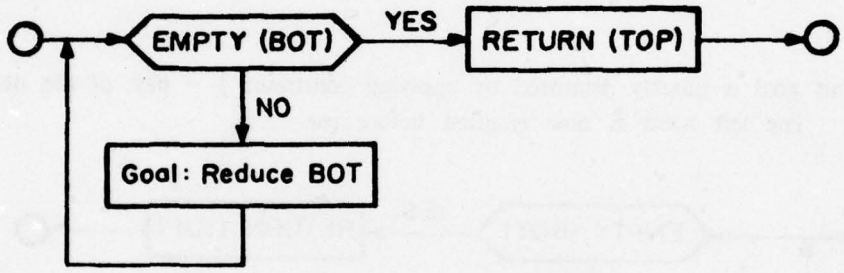

Step 3: The goal matches part of constraint 4 -- the definition of PICK/FROM. So the constraint is applied, and the plan is now fully reduced to primitives actions:

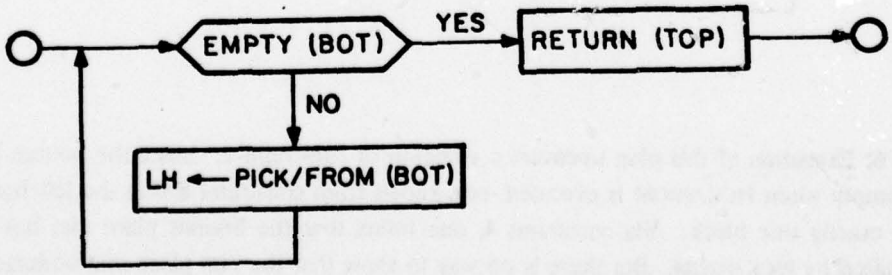

Step 4: Execution of this plan reveals a violation of constraint 6: the left hand must be empty before one can pick something up. So a new goal is created:

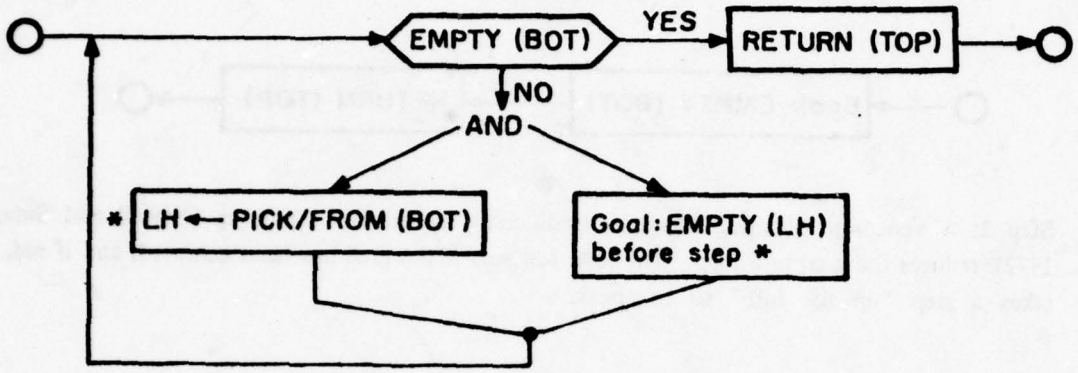

Step 5: This goal is quickly dismissed by applying constraint 7 -- part of the definition of PUT/ONTO. The left hand is now emptied before use.

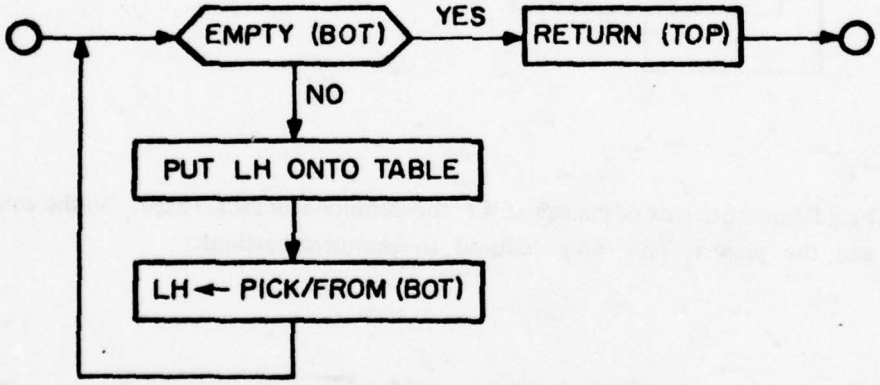

Step 6: Execution of this plan uncovers a violation of constraint 2. Since the bottom place mat is not empty when PICK/FROM is executed, one knows from constraint 8 that the left hand comes to hold exactly one block. Via constraint 4, one infers that the bottom place mat has its contents decreased by PICK/FROM. But there is no way to show that the TOP place mat undergoes an equal change. So, constraint 2 is violated, and a new goal must be created. The goal says that there must be a change in TOP to equal the change in BOT.

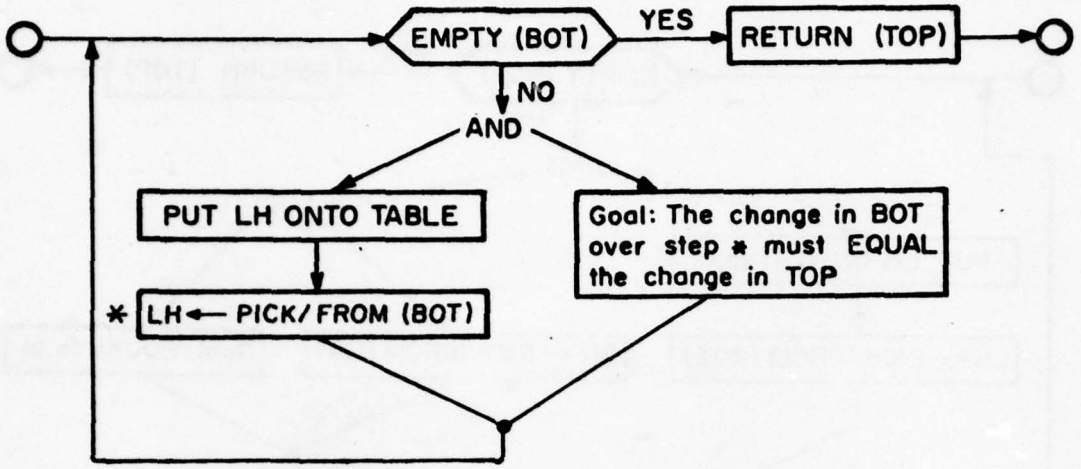

Step 7: Part of this goal matches constraint 4, the definition of PICK/FROM. A new picking up action is instantiated for the top place mat. This reduces the goal of equal changes to the goal of equal contents of the left and right hands.

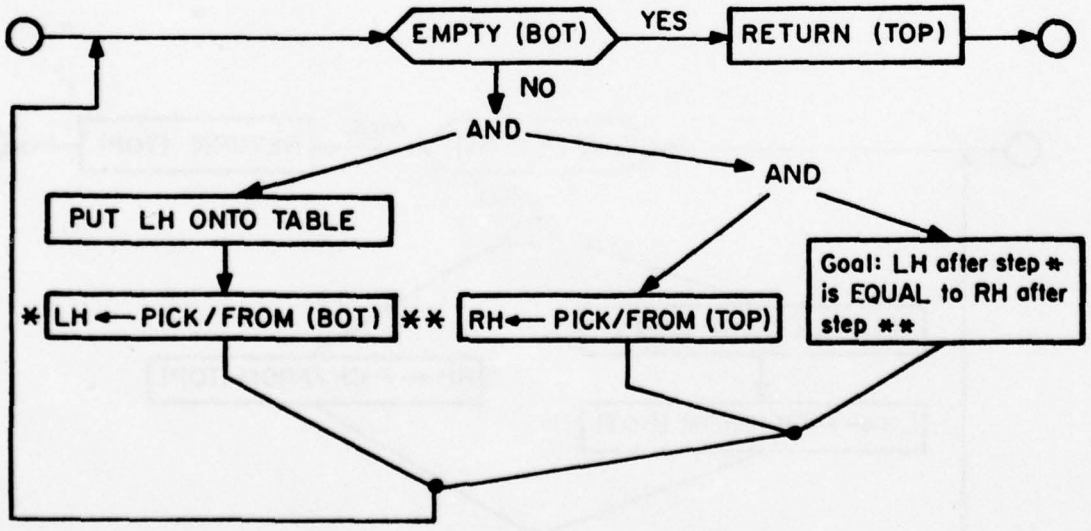

Step 8: Constraint 8 can apply twice now, once per hand. It says that only one block is picked up by PICK/FROM. Thus, the goal of EQUAL contents is replaced by equality of two arbitrarily chosen blocks.

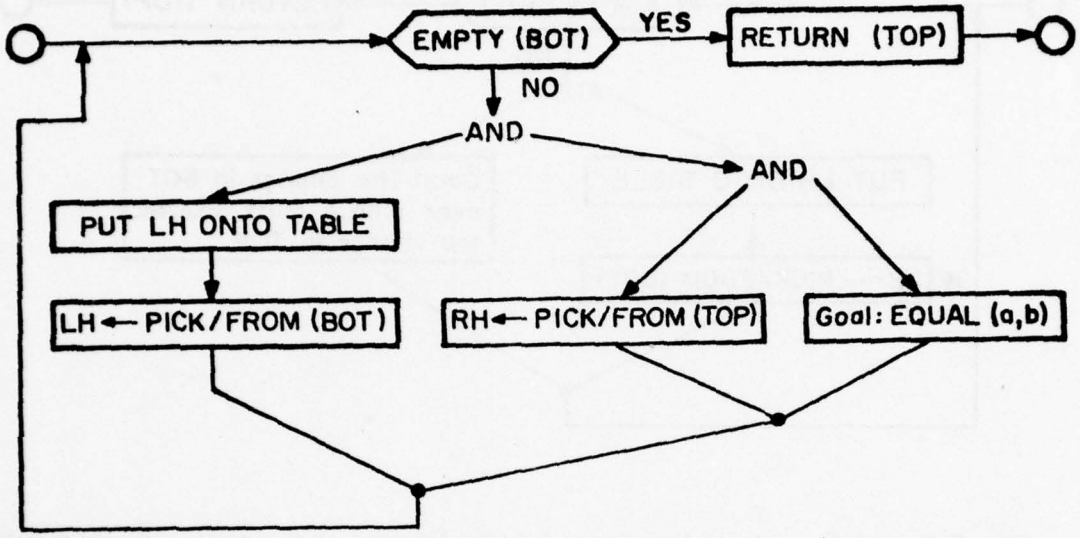

Step 9: Of course, this new goal is trivially satisfied by constraint 3 -- all blocks count the same in the base-1 number system. So the goal is simply removed from the plan.

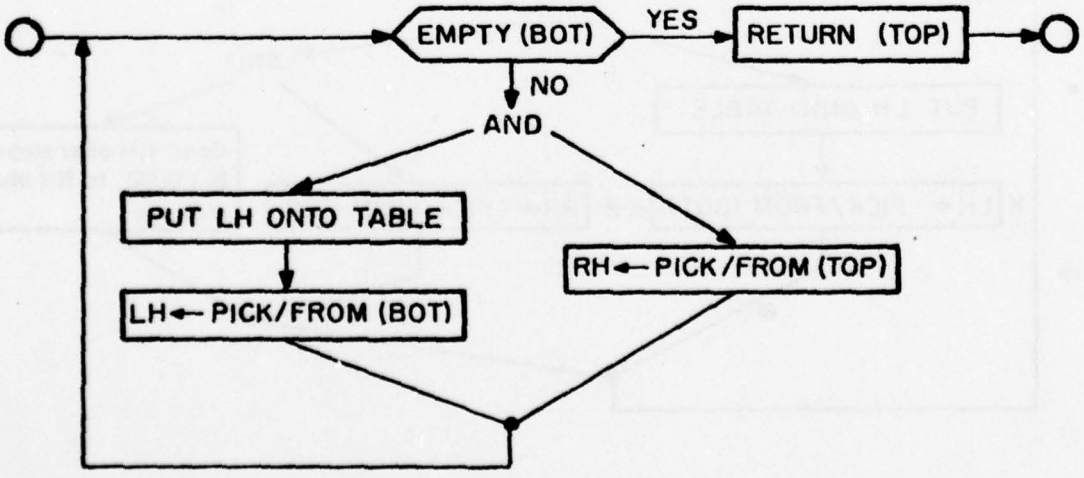

Steps 10 and 11: Execution reveals that constraint 6 is violated again, this time by the right hand. So it must be emptied before use as well, in the same two-step fashion as Steps 4 and 5.

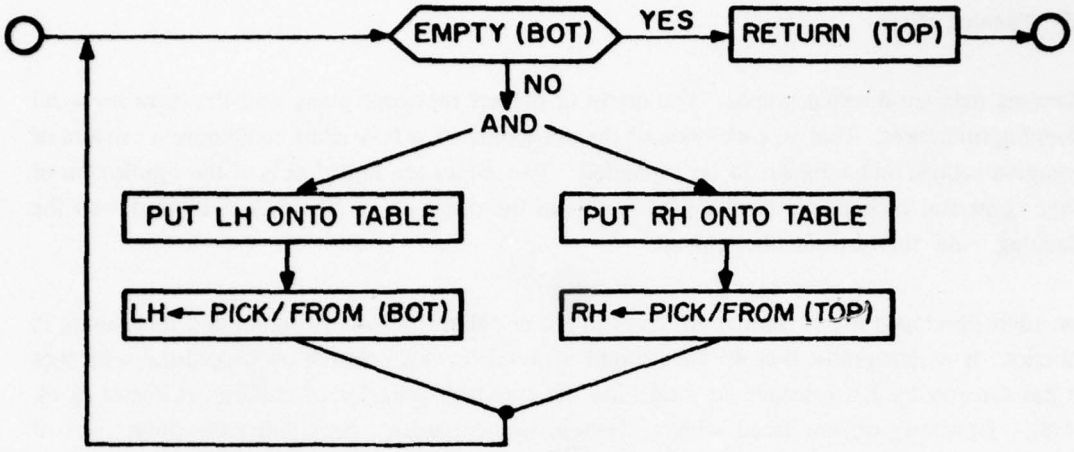

Step 12: A planning heuristic, call it Conjunction Reduction, removes the conjunction AND. The AND node is for conjoining subgoals. It makes no statements about which subgoal to achieve first. In this case, it doesn't matter how the subgoals are ordered since they turn out to be independent. So the rule arbitrarily chooses the following ordering:

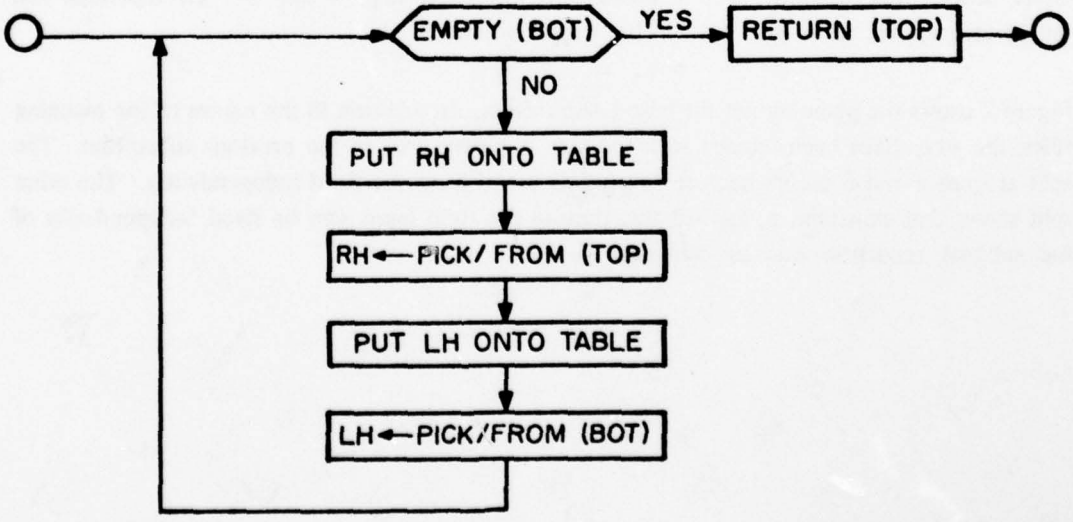

This is the final plan. Every step is a primitive, and all the constraints check out. The planning for base-1 subtraction is complete. The final plan is exactly the flow chart representation of the surface structure of the procedure.

3.7 Planning Nets

Planning nets are directed graphs. The nodes of the net represent plans, and the links represent planning inferences. That is, each node of the net stands for a flow chart containing a mixture of primitive actions and subgoals to be expanded. Two nodes are linked only if the application of some constraint or heuristic to one plan results in the other plan. The link is labeled with the planning rule that causes the change.

Sacerdoti developed a very similar structure to aid in automated task planning and monitoring in robotics. It is remarkable that we have found it useful for our research on procedural semantics as has Greeno for his research on modelling the counting behavior of children (Greeno et. al. 1978). However, we are faced with a clash in nomenclature. Sacerdoti calls these sorts of structures "procedural nets". We prefer to call them "planning nets," since their content has more to do with the planning of a procedure than the procedure itself.

Planning nets are partial orders

In fact, planning nets are generally not sequences, as the chronological presentation of the previous subsection might lead one to believe. Often, two planning inferences can be applied in either order. For example, step 6 could have preceded steps 4 and 5. To represent this independence, we allow the net to be a *partial order.*

Figure 2 shows the planning net for base-1 subtraction. In addition to the names of the planning rules, the steps have been labeled with the step numbers used in the previous subsection. The split at steps 4 and 6 occurs because constraints 2 and 6 can be fixed independently. The other split shows that constraint 6, applied this time to the right hand, can be fixed independently of the subgoal reduction due to constraint 8.

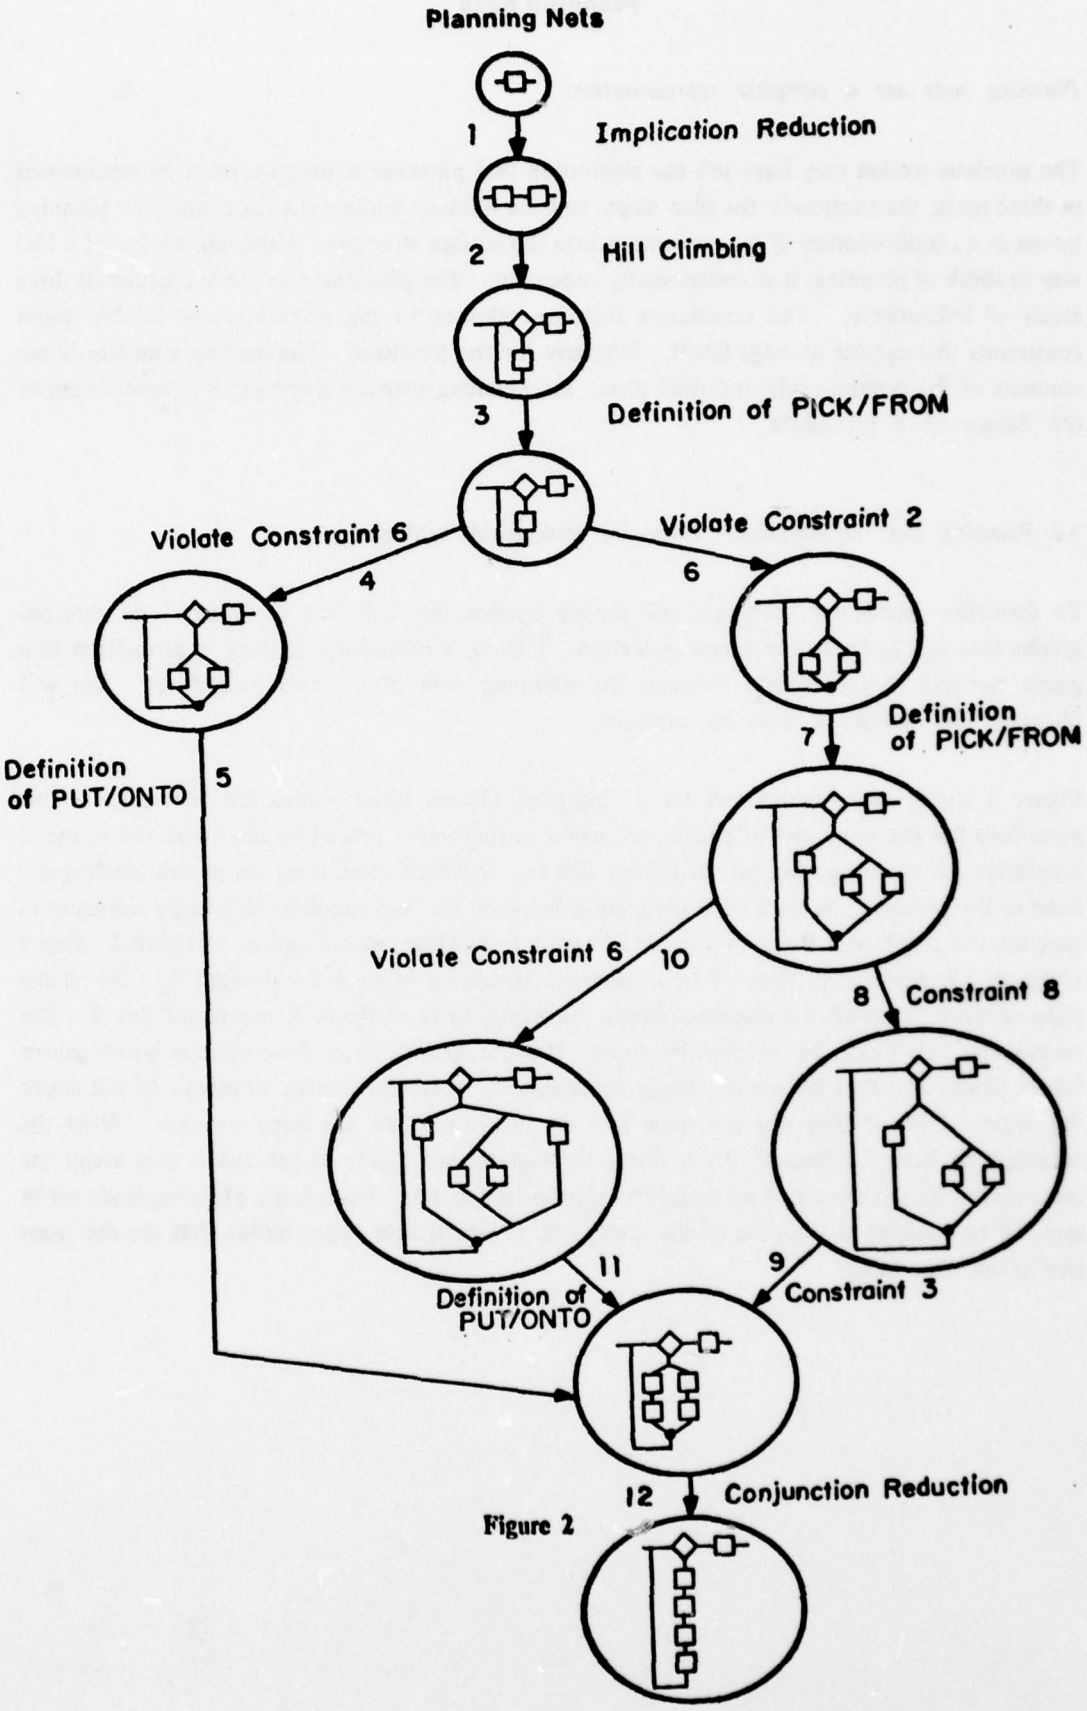
Planning Nets
1
Implication Reduction
2
Hill Climbing
3
Definition of PICK/FROM
Violate Constraint 6
4
Violate Constraint 2
6
Definition of PUT/ONTO
5
7
Definition of PICK/FROM
Violate Constraint 6
10
8
Constraint 8
11
Definition of PUT/ONTO
9
Constraint 3
12
Conjunction Reduction

Figure 2

Planning nets are a complete representation

The previous section may have left the impression that planning knowledge must be represented in three parts: the contraints, the plan steps, and the ultimate surface structure, and that planning serves as a transformation of the constraints into the surface structure. Although this is not a bad way to think of planning, it is unnecessarily redundant. The planning nets alone capture all three kinds of information. The constraints that are relevant to the procedure are exactly those constraints that appear as edge labels. Similarly for the heuristics. The surface structure is the contents of the bottom node, the final plan. So, planning nets are a complete representation of the design of a procedure.

3.8 Planning net mp-morphisms formalize procedural analogies

To formalize procedural analogies, one merely applies the definition of "match" for directed graphs that was given in the previous section. That is, a procedural analogy is formalized as a graph theoretic mp-morphism between the planning nets of the two procedures. We will illustrate this definition with an example.

Figure 3 shows the planning net for a "big-pile" Dienes Block subtraction procedure. This procedure has the same sort of pairing-off action as the base-1 procedure discussed above, but it represents a number as a big pile of Dienes Blocks. Although space does not permit labeling the links in the planning net with their planning inferences, the step numbers should be sufficient to describe the match with the planning net of base-1 subtraction, which appears in figure 2. Step 9 of figure 2 is replaced in figure 3 by a subgraph consisting of steps 9.0 through 9.7. So all the links of figure 2 match the correspondingly numbered links in figure 3, except for link 9. The reason why link 9 can't be matched is simple: it is the application of the constraint which makes base-1 blocks all count the same, namely constraint 3. In Dienes Blocks, all blocks do not count the same. Only if they are the same size do they designate the same number. What the subgraph of steps 9.0 through 9.7 is doing is planning out a way to get blocks that aren't the same size to be the same size by doing the appropriate trading. In fact, the planning leads off in step 9.0 by noticing a violation of the constraint 3', which says "only blocks that are the same size count the same."

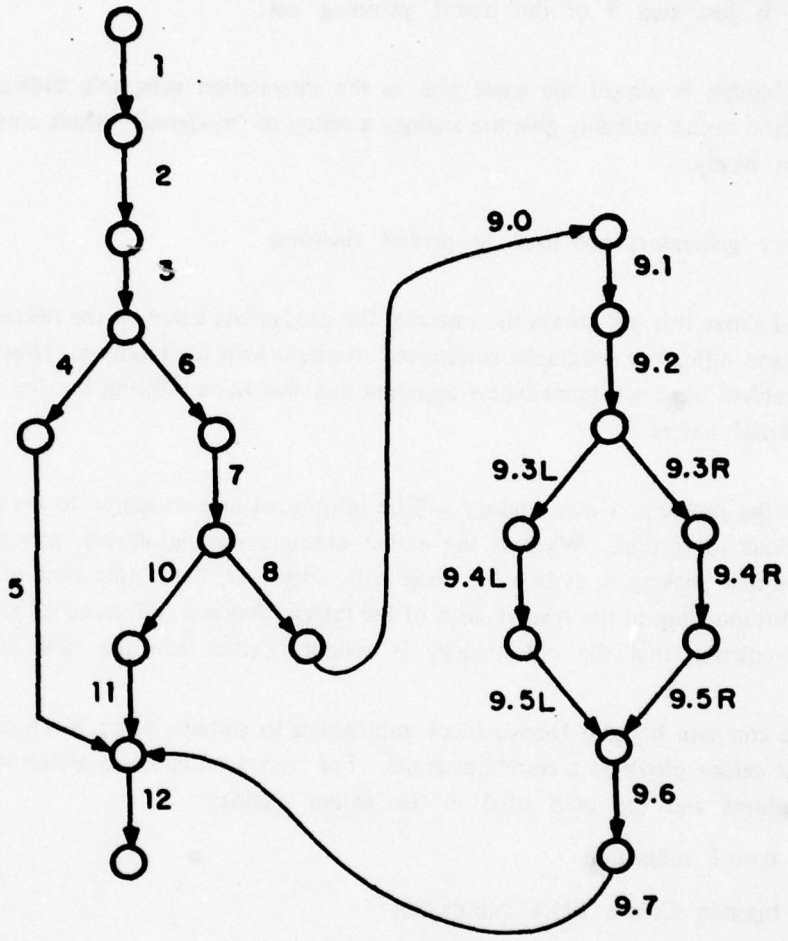

Figure 3

The mp-morphism of the two planning nets results in the following intersection and difference subgraphs (calling the Dienes Block procedure "A", and the base-1 procedure "B"):

A∩B is almost the whole planning net for base-1 subtraction, except the link for step 9.

A-B is the subgraph that replaces step 9, whose steps are labeled 9.0, 9.1, etc.

B-A is just step 9 of the base-1 planning net.

The A-B subgraph is almost the same size as the intersection subgraph, indicating that the closeness metric would probably give the analogy a rating of "moderate", which corresponds with the intuition nicely.

3.9 Difference generators are used to predict closeness

As we hinted above it is not always the case that the predictions based on the relative sizes of the intersection and difference subgraphs correspond so nicely with the intuition. However, in those cases the problem has been immediately apparent and was fixed utilizing the fact that planning nets are *partial orders.*

To illustrate the problem, a new analogy will be introduced and compared to the one described in the previous subsection. Whereas the earlier example was, intuitively, a moderately close analogy, this new analogy is quite a bit closer still. However, the simple view of the closeness metric as corresponding to the relative sizes of the intersection and difference subgraphs, leads to the false prediction that the old analogy is actually closer than the new one.

Suppose we compare big-pile Dienes Block subtraction to sorted Dienes Block subtraction, an analogy that earlier provided a counterexample. For convenience, let us attach some letters to these procedures and the ones used in the earlier analogy:

A: base-1 subtaction

B: big-pile Dienes Block subtraction

C: sorted Dienes Block subtraction

The BC analogy is intuitively rather close. However, when the planning nets are compared, we find a huge subgraph of C that isn't matched, namely all the design that has to do with maintaining the sort. Indeed, this difference subgraph, C-B, is much larger than B-A and A-B together. Subgraph B-C is also quite large. Hence, even though B∩A is somewhat smaller than B∩C, any reasonable metric would predict that analogy AB should be closer than analogy BC, contrary to the intuition that big-pile Dienes Block subtraction is more similar to sorted Dienes Block subtraction than to base-1 block subtraction. There is a mismatch between predictions of the theory and judgements of closeness.

But closer examination of subgraph C-B reveals it has only one entering link, just like link 9.0 of figure 2. This link is labelled "Violates Constraint 11: keep blocks sorted by size". In other words, it appears that one plan inference is *causing* all the others. We can capture this notion of causation by utilizing the *topology* of planning nets.

As discussed above, planning nets are partial orders. Any subgraph of a partial order is also a partial order. In particular, the difference subgraphs are always partial orders. Any partial order has a unique set of *minimal elements*. This set is the smallest set of links that dominate all the other links in the subgraph. These mathematical facts insure that the following terms are well-defined:

> Where X and Y are any two planning nets, let d(X-Y) be the links that are the minimal elements of the difference subgraph X-Y, and let d(Y-X) be the links that are the minimal elements of Y-X. Call these two sets the *difference generators* of mp-morphism XY.

Difference generators are a formal representation of what is causing the difference between two procedures. Intuitively, what the difference generators of mp-morphism represent are the *crucial ideas* that separate the two procedures. All the other differences between the two procedures stem from these few crucial ones.

To illustrate this notion of "crucial ideas", take the analogy between base-1 and big-pile Dienes Blocks, which we were calling analogy AB in the previous section. d(B-A) is a graph with just one link, labelled "Step 9: Constraint 3 -- all blocks are EQUAL." d(A-B) is a link labelled "Step 9.0: Constraint 3' -- two blocks are EQUAL if and only if they have the same SHAPE." Replacing constraint 3 by constraint 3' is about as clear a statement of the difference between base-1 blocks and Dienes Blocks as one can hope to make.

Because difference generators capture the distinctions between procedures so succinctly, they seem highly appropriate as the inputs (or arguments) to the closeness metric. They are decoupled from the unimportant details that fill flow charts, procedural nets and planning nets, details which obscure the essence of analogy by inflating difference subgraphs with derived, less meaningful structure. Indeed, the comparison of analogy AB to analogy BC (i.e. the big-pile vs. sorted analogy) now agrees with intuition: all four difference generators, namely d(A-B), d(B-A), d(B-C) and d(C-B), are about one link big. On the other hand, the intersection subgraphs are as before, with $A \cap B$ being smaller than $B \cap C$. Since the difference generators are about the same size, the intersection sets are more important in the closeness metric. Hence, a reasonable metric would report that BC is closer than AB, which corresponds with the intuition that big-pile Dienes Blocks subtraction is closer to sorted Dienes Block subtraction than to base-1 blocks subtraction. At last, we seem to have found a level of abstraction for procedures where intuitions of closeness correspond to the relative sizes of the inputs to the closeness metric.

3.10 Discussion

The main point of this section has been that planning nets provide a basis for a theory of analogy that can predict the judgements of experts on the closeness of analogies between procedures. Moreover, all the aspects of the theory have very natural, almost elegant sources. The deep structure used came naturally from Sacerdoti's work in robotics, mp-morphisms are a general purpose concept, and the notion of difference generators came naturally from the topology of planning nets.

We have always been struck by how much of the design of a procedure like subtraction is governed by the design of the representation of the objects manipulated by the procedure (e.g., the place-value number system). In fact, many of the actions in any of the elementary arithmetic procedures concern not the mathematical operation per se but rather *how the object representations are manipulated.* This impression is reinforced by experience in computer programming, which is often a constant interplay between the design of the object (i.e., data) representation and the code, even at the highest levels. Anyone who has tried to understand a program that he did not write can vouch for the importance of understanding the data representation. In the process of judging the closeness of an analogy, a popular strategy is to first look at each procedure's object representation, and then build the understanding of the overall analogy on the basis of the analogy between object representations. In short, it appears to us that a large portion of the "understanding" of a procedure consists of an understanding of the implications of the procedure's object representation.

This view of procedural understanding is entirely consistent with the planning net formalism. The constraints and heuristics that appear in the net represent are, in some sense, the *essence* of the procedure. If object representations were unimportant, then none of the planning inferences would be "about" the object representation. But in fact, many planning inferences do deal with the object representation. Even in the base-1 blocks procedure above, with its extremely simple object representation, we find constraint 3 addressed solely to the object representation. In more complex procedures, using Dienes Blocks or written numerals, an even larger portion of the constraints concern the object representation. In short, although planning nets abstract out the less important aspects of a procedure, they leave behind the design of the object representation, which is quite compatible with the view that, as a representation of "understanding" of procedures, a fair portion of the constraints should model the "understanding" of the object representation.

We have not discussed the exact definition of the closeness metric, even though some definition would be necessary to methodically verify the correlations we claimed above. There are many difficulties and fine points involved in determining such a definition. In particular, it is plausible that the weight of some planning inferences is quite close to zero. We have in mind the common sense heuristics, such as Implication Reduction, that play an almost invisible role in the

planning. Also, some planning rules are applied more than once in a planning net; one may perhaps wish to avoid giving such rules an inappropriate prominence by only counting their first occurrence in the difference generators or the intersection sets. These are just two of the many points that one would have to consider in defining a closeness metric.

The reader has no doubt noticed the incredible amount of work that goes into analyzing a procedure in terms of its planning. First one constructs the flow chart, then the constraints and a sequential plan for the flow chart, and lastly calculates the planning net by noting which planning inferences are not ordered with respect to each other. This large amount of work leaves much room for error on the part of the theorists. However, each level of abstraction is well defined, and can be checked for consistency by a computer. Thus, one next step is to build a computer system of utilities to aid in the analysis of procedures. However, there is a certain amount of intuition that goes into some parts of the analysis, notably the formulation of a set of constraints, that we doubt could ever be successfully mechanized.

4. Analogies and Teleologic Semantics in Educational Research

In this section we consider some of the issues involved in explaining (or teaching) the knowledge we explained in the first previous section--teleologic semantics. Briefly, teleologic semantics is the kind of knowledge that concerns the purpose of each part of the procedure as well as the motivation behind the set of constraints that defines the particular representation for the objects. In particular, we consider how an individual piece of teleology can be explained, and how such individual explanations can be combined into an integrated explanation.

The section closes with a discussion of some issues involved in microworld-based curricula. These issues turn out to be intimately related to those involved in teaching teleologic semantics.

4.1 Local explanations: manifestation and motivation

An important property of the planning net formalization is that there is a *natural* notion of how to explain a small piece of a procedure's teleologic semantics. By "piece" we mean a constraint (or a small set of constraints) that is used in the planning net. To "explain" it, one uses a *minimal contrasting pair* of procedures -- one with the constraint, and one without it -- that compute the same "operation" as the given target procedure. In other words, we use analogies to illustrate constraints. We believe that using a concrete surface structure illustration for each deep structure concept to be explained is a very important explanatory technique that naturally falls out of this development. For example, this method of explanation frees us from having to explain the planning formalism to the student--a task potentially more difficult than teaching the procedure itself.

More formally, to illustrate some given constraint(s), one uses *two analogous procedures such that one of the difference generators of the mp-morphism between them is exactly the given constraint(s).* If the pair of procedures forms a minimal contrasting pair then the mp-morphism constituting the analogy is *elementary.*

Of course, this technique works just as well for explaining heuristics. However, heuristics are often such common sense knowledge that an explanation of them is unnecessary. So we will call the planning inferences to be explained "constraints," avoiding the cumbersome phrase "constraints or heuristics." Also, our terminology will reflect the fact that it is often possible to provide a minimal constrasting pair for each constraint individually (this observation is discussed below); we use "constraint" in place of "a small set of constraints."

An important realization is that minimal contrasting pairs can be used in two different ways in an explanation. They can be used to show how the constraint is *manifested* on the surface, and they can also be used to *motivate* the inclusion of the constraint in the ultimate design of the

procedure. Probably the best way to illustrate the differences between these two uses is with an example.

Explaining the canonicity constraint

The particular constraint that will be used in this example is one of the most subtle and influential in arithmetic, namely the *canonicity constraint.* To show how the planning net representation can aid in explaining procedures, the constraint will be presented as the "answer" to a non-trivial teleologic question.

What is the purpose of carrying? More specifically, if the problem is 52+49, why bother to carry ten? Why not leave 11 in the units place? It is not because there is no symbol for the "digit" eleven -- we could invent one if we wanted. In Dienes Block addition, the question is even clearer. Why not leave the answer in the form of 9 longs and 11 units? Why bother carrying?

The answer is that carrying maintains the *canonicity* of the representation of numbers. A canonical representation puts the representational objects in one-to-one correspondence with the real objects they represent. The Hindu-Arabic representation of numbers is canonical since there is a unique, distinct numeral for each number. Dienes Blocks are not necessarily a canonical representation, since most numbers can be represented several ways. For instance, eleven can be represented as a long and a unit, or as eleven units. The purpose of carrying is to canonicalize the sum, by making sure that there are no more than nine blocks of any given shape. In other words, carrying is the *manifestation* of the canonicity constraint.

But suppose the questioner rejoins by asking what the purpose of the canonicity constraint is. The answer involves another arithmetic subprocedure -- comparison.

It is much more efficient to find out which numeral represents a given large number if the representation is canonical. Let us use a Dienes Blocks comparison procedure to illustrate the gain in efficiency. In a non-canonical representation, the comparison procedure must compare all the piles, since a very large pile of small blocks can make up for a deficit of larger blocks. In a canonical representation, the comparison procedure needn't check all the piles. If it finds that one numeral has more flats than the other numeral, then it needn't compare the longs or units; even if the other numeral has the maximum number of longs and units allowed, namely 9 each, the first numeral will still represent the larger number. Imposing the canonicity constraint makes the comparison procedure much more efficient, because it allows the procedure to stop earlier. But the canonicity constraint is a constraint on the representation of numbers, and so all arithmetic procedures must obey it. Even though the constraint makes part of the addition procedure somewhat less efficient, it makes comparison so much more efficient that it is worth

having. This appeal to efficiency is the ultimate endpoint in the explanation of the *motivation* for carrying and the canonicity constraint.

In this mini-explanation of carrying, we have seen two important facets of teleologic knowledge. In the addition procedure, the canonicity constraint was manifested as a *carry* subprocedure. But the motivation for adopting the constraint lay in another procedure, *comparison*. Each of these two facets, which we will begin calling *local* explanation since they explain just one constraint, was illustrated with a minimal contrasting pair of procedures. One member of the pair was a fully operational version of the procedure that lacked the constraint being discussed, while the other member adopted the constraint. But the manifestation part of the explanation involved a minimal contrasting pair which was different from the pair used to motivate the constraint (i.e., addition vs. comparison). As will be discussed later, it is preferable to have a pair of analogous procedures which illustrates both the manifestation and the motivation of teleologic concepts but this is not always possible.

It is our belief that the concreteness of this minimal contrasting pair paradigm of explanation is of crucial importance in making teleologic semantics clear. The learner can *see* in very concrete terms how adopting a constraint effects the procedure. Winston show that a similar example-based paradigm was sufficient to teach the abstract concepts necessary to recognize toy block constructions, such as an arch (Winston 1975, 1978).

In fact, many minimal contrasting pairs that manifest the given constraint are available, depending on which of the remaining constraints are adopted. If all the constraints of a given target procedure are adopted, then one member of the pair is the target procedure itself. Otherwise, the contrast is exhibited across a pair of *model* procedures which still satisfy the mathematical constraints of the target procedure. Using model procedures often highlights the contrast, making it much easier to see the constraint under discussion. Such was the case with the canonicity constraint, where Dienes blocks allowed us use non-canonical numbers without inventing new digit symbols.

However, model procedures must be used with some care, as the following example illustrates.

The impact of efficiency metrics on "loop jamming"

Consider the difference between the standard carry subprocedure and the two-pass version described in the introduction, where carrying was deferred while all the columns were added, then performed on a second pass over the columns. This difference is a constraint that was called *loop jamming*, after the compiler optimization technique of the same name which weaves two loops into one (Allen and Cocke 1972).

One can not use Dienes Blocks procedures to motivate loop jamming, since exactly the same number of hand motions, fact table lookups, etc. are required by each procedure. So, Dienes Blocks are an inappropriate model domain for discussing this constraint.

However, when implemented with written numerals, loop jamming does create a difference in efficiency. (See note 8 for details) The two-pass implementation of carrying requires more writing than the standard implementation. Thus written arithmetic turns out to be an appropriate domain for discussing the loop jamming constraint.

The important point to notice about this example is that the choice of the model has some impact on the local explanation. In particular, a model which clearly displays the *manifestation* of the constraint in the procedure may not be able to demonstrate the *motivation* for the constraint. For example, since one doesn't have to worry about how to write the intermediate column sums which may be greater than nine with Dienes Blocks, we can use them to implement both the one and two pass addition procedures and thus use them to illustrate the manifestation of loop jamming. But, unfortunately, they cannot be used to motivate loop jamming since the resulting procedure is no more efficient.

Another point to notice about the preceding example is the use of *efficiency metrics* in motivating design choices. An efficiency metric is some weighted sum of hand motions, fact table lookups, table size, amount of paper used, etc. The weighting of efficiency metrics is very important. For example, if reducing memory load is more desirable than decreasing the number of write operations, then the discussion of loop jamming ends with the opposite conclusion, that two-pass carrying is better than the standard subprocedure (see note 9 for details). The two-pass version uses less short term memory but more pencil lead. So, exactly what efficiency metric is used greatly affects the local explanation. We do not look upon efficiency metrics as a regrettable new variable that must be tied down and parameterized with careful experimentation, but rather as a source of flexibility that can be used to tailor the teaching paradigm to the needs of particular students.

4.2 Principles for sequencing local explanations

For moderately complex procedures, such as subtraction, the number of constraints can be high enough to cause problems of presentation. Our current best estimate of the number of constraints of subtraction is 17. To explain this many constraints, each with its own manifestation and motivation, may seem a difficult task. However, with the planning net formalism, we can investigate how to "optimally" sequence a collection of "model" procedures; the first procedure (or "model") in the sequence would be a very, very simple version of the skill, and the last procedure in the sequence would be the target procedure. For example, in subtraction the first procedure might be base-1 block subtraction and the last, standard written subtraction. But how should the intermediate models be sequenced?

Planning Nets

Using the formalisms developed above, principles for sequencing local explanations can be stated precisely. Several such principles are stated below that we believe will lead to sequences that better enable assimilation of the overall teleology of a procedure from the explanations of its parts. Each one of them falls out quite naturally from the planning net formalism.

It will be convenient in what follows to say that such sequences run from left to right -- the target procedure is the procedure on the far right. This allows us to talk of the left and right procedures of a mp-morphism. Also, we will speak of the left and right difference generators of a mp-morphism; if A is left of B, then d(A-B) is the left difference generator.

Introduce each constraint

As we saw in the previous subsection, it is best to illustrate each constraint with a minimal constrasting pair of analogous procedures. This is probably the most important sequencing principle, that each constraint be illustrated individually. However, it is probably also true that it is better to introduce the constraint rather than take it away. This gives the sequence an air of progression toward the target procedure. Putting this principle formally, we have: each constraint is the sole contents of the right difference generator of some mp-morphism in the sequence. That is,

> Principle 1. For each constraint *C* in the target procedure's *planning net*, there exists i such that $d(P_i - P_{i-1}) = \{ C \}$,

where the procedures are numbered from left to right (first to last).

Starting with a very simple procedure would, hopefully, tap a person's intuitive understanding. Then, since each of the analogies (mp-morphisms) is very close (or at worst, moderate -- we are guaranteed only that one of difference generators is a singleton set, namely the constraint being introduced), it should be easy to transfer that understanding along, *augmenting* it only *slightly* as each new procedure is presented.

Only introduce target procedure constraints

Occasionally, it is necessary to "build" a left procedure to illustrate some constraint. This occurs when one can not adjust the sequence so that the right procedure of some other constraint is this constraint's left procedure. In this case, one ends up with an adjacent pair of procedures that do not illustrate a constraint from the target procedure. Although the person (or computer) doing the explaining can mention that this analogy isn't so important, it would be better if the sequence didn't have such pairs. So, another optimization principle to shoot for is:

> Principle 2. For each i in the sequence, there exists a constraint C in the target procedure's planning net, such that $d(P_i - P_{i-1}) = \{ C \}$.

Minimize redundancy

One should not remove a constraint that has been introduced previously, or introduce a constraint twice. Although one could argue that the redundancy of seeing the constraint illustrated in several different contexts (i.e. with different model procedures) serves to reinforce the local explanation, we are of the opinion that this would create confusion, rather than dispel it, and in addition, it would create the impression that the sequence was meandering.

More formally, we propose that the sequence obey the following conditions:

Principle 3. For any $i \neq j$ $d(P_i - P_{i-1}) \cap d(P_j - P_{j-1}) = \emptyset$

Principle 4. For any $i \neq j$ $d(P_{i-1} - P_i) \cap d(P_{j-1} - P_j) = \emptyset$

Principle 5. For any i,j $d(P_i - P_{i-1}) \cap d(P_{j-1} - P_j) = \emptyset$

The first condition advises one not to introduce a constraint twice, and the second condition advises one to avoid removing a constraint twice. The third condition says that once a constraint is introduced (the first term) it can never be taken out (the second term). Actually, it also says that once a constraint is removed, it shouldn't be reinserted, which is also a plausible condition to impose for aiding the cogency of the sequence.

Efficiency should increase monotonically

We mentioned above that a minimal contrasting pair for a constraint does not necessarily show an increase in efficiency. That is, all ways of manifesting a constraint do not necessarily motivate it as well. One condition on a sequence is that the model procedures be chosen and sequenced so that efficiency always increases as the target constraints are adopted. That is,

Principle 6. For all i, P_i is more efficient than P_{i-1}.

Since there are many minimal contrasting pairs that manifest a constraint, it is usually not difficult to find some pair that motivates it as well, but putting that pair into a sequence with the other constraint's pairs can be somewhat difficult. We know of only one constraint for addition or subtraction, namely the canonicity constraint, where the motivation pair *must* be distinct from the manifestation pair. This is inevitable since canonicity is basically designed to improve the efficiency of comparison, not the other arithmetic operations. Thus, if one were only interested in a sequence of addition procedures or subtraction procedures, then the pair for the canonicity constraint would necessary violate this sequence principle. However, with this one exception, it has been easy to find *some* minimal contrasting pair that serves to both manifest and motivate a constraint for subtraction.

However, putting such pairs into a sequence requires some care. Switching the order of two constraints in a sequence often alters the relative efficiency of the minimal contrasting pair of procedures that manifest the unit. Under one ordering, both constraints might improve efficiency. But under the reverse order, adopting one of the units may result in no increase in efficiency, or even a decrease in efficiency. This might seem strange, so let us pause a moment for an example.

Consider ordering the canonicity constraint versus the constraint that Dienes Blocks be kept sorted by size. First, suppose that the canonicity constraint precedes the sort-by-size constraint in the sequence. Under this ordering, the efficiency increases between each procedure; imposing the canonicity constraint forces the procedure to search through the big pile of Dienes blocks to check that there are no more than ten blocks of any given shape. Hence, adopting the sort-by-size constraint greatly improves efficiency by eliminating rumaging around through the big pile in favor of simply counting up the number of blocks in each of the small piles.

Now suppose the order in the sequence were reversed, and sort-by-size were imposed before canonicity. The minimal constrasting pair for sort-by-size consists of (a) adding two big piles of Dienes Blocks together by simply forming the union, versus (b) adding each of the small piles together in a series of separate union operations. Now the introduction of the constraint actually decreases the efficiency of addition. Since no carrying is required (canonicity not being imposed yet), there is no use in the separation by size. Maintaining the constraint creates extra work with no reward. So, modifying the order of two constraints in the sequence can impact the ability to motivate them.

Although it may be a difficult condition to achieve, if a manifestation-based sequence has monotonically increasing efficiency, the viewer can see with no additional examples not only *what* each constraint is, but *why* it exists (i.e. what good it is).

Telescoping sequences

Occasionally, one finds mp-morphisms which introduce a constraint but don't need to remove any constraints. The canonicity constraint can be illustrated with a mp-morphism whose left difference subgraph is null (for addition, one could use the two-pass addition procedure described in the introduction as the right hand procedure, and the first pass of it for the left procedure). That is, the mp-morphism is *total* with respect to the left planning net. It seems plausible that mp-morphisms which never removed constraints would create a very strong sense of progression toward a target procedure. Such sequences are characterized by the condition

Principle 7. For any i, $d(P_{i-1} \cdot P_i) = \emptyset$

4.3 A space of mp-morphisms

Needless to say, it will rarely be possible for a sequence to satisfy all the sequencing principles mentioned above. Indeed, we may only be able to satisfy some principles along part of its length, and different principles along another part. We need some way to study the relative contribution of the various principles to ease of explanation.

Ultimately, we would like to develop a representation of all principled sequences to a given target procedure. These sequences could be represented in an economical way by a directed graph whose nodes would represent planning nets. There would be a link from node A to node B only if they appeared as an adjacent pair in some sequence that was considered a plausible explanation sequence, perhaps because it met some minimum number of the principles listed above. (In particular, one might include all (known) minimal constrasting pairs for the target constraints -- this would correspond to using principle number one as a threshold for inclusion in the space.) This directed graph has the property that any sequence from a "most primitive version" node to the "target" node would be a possible sequence for explaining the teleology of the target procedure. We tend to think of this graph as a *space of* mp-morphisms.

One clear problem that could be attacked with such a space is improving on the *naturalness* of teleologic explanations. Presenting the 17 or so mp-morphisms (or procedural models) for place-value subtraction is bound to be very confusing unless they can somehow be aligned along the individual's own cognitive structures (see Appendix 1 for a detailed example of one such chain of models). We have already mentioned seven principles that probably contribute to better comprehension of such explanations. Each of these principles would be incorporated into the space, perhaps as annotations on the basic partial order. Hopefully, experience and experiment will lead to the discovery of other factors that improve the naturalness of teleologic explanations.

4.4 Using the mp-morphisms space in microworld-based curricula

In a microworld-based curriculum, the student explores a rich environment, hopefully inventing something analogous to the target skill. For example, a student might be given Dienes Blocks and a puzzle which requires using multi-digit arithmetic to solve it. Actually, how students are motivated to do the arithmetic is not an issue here. The point is that students are not given the sequence of actions that implement arithmetic for the given representation of numbers. Instead, they must invent it themselves. Invention is the essence of a microworld-based curriculum.

Tracking a student's discovery process

The space of mp-morphisms could be quite useful as a way to "track" a student's discovery process. The basic idea is that an observer (possibly a computer) analyzes the procedures that

the student invents in terms of planning nets. The nodes in the space that correspond to the plans of these procedures are marked. The student's progress is then expressed as the shortest sequence along the constraints that connect the marked nodes. This provides a strong hypothesis concerning what the student has learned during the discovery process.

Such a tracking study would provide an empirical way to verify conjectures about "natural" sequences for teleologic explanations. That is, observing that students generally followed sequences that increase the efficiency of the procedure would support the conjecture that monotonically increasing efficiency is important for cogent, natural explanations.

Sequencing microworlds

A persistent problem in microworld-based curricula is how to sequence the microworlds so as to maximize the cumulation of intuitions built up while exploring the microworld and enable them to be transferred to the target procedure. One ready answer is provided by the space of mp-morphism sequences, assuming it has been annotated to show which sequences are most natural.

Sequencing microworlds obviously imposes an order on the traversal of the nodes in the mp-morphism space. Although there are many Dienes Block procedures and abacus procedures, one can't move from a Dienes Block procedure to an abacus procedure's node until one leaves the Dienes Block microworld and enters the abacus microworld. So, the most natural sequence of microworlds is the one that enables traversal of the most natural sequences through the constraint space. Let us illustrate this conjecture with an example.

Suppose one tried to teach addition with the following sequence of microworlds:

base-1 blocks, the abacus, Dienes Blocks, written numbers

One would expect the students to become frustrated when they find that the teleology associated with place-value encoding of numbers, which they laboriously invented for the abacus, is obviated by the shape-value encoding of Dienes Blocks. And when they find they must resurrect this place-value notion to move from Dienes Blocks to written numbers, one would expect them to become disgruntled, or worse yet, apply "teacher psychology" and guess that place-value couldn't possibly be part of the design because "we already had that." In comparison, reordering the sequence to be

base-1 blocks, Dienes Blocks, the abacus, written numbers

allows invention of the notion of place-value just once, in transistion from Dienes Blocks to the abacus, and then maintenance of the notion throughout the abacus microworld and on into the written numbers.

These ordering results could be predicted on the basis of one of the naturalness principles mentioned above, namely, that constraints ought to *accumulate* along the sequence. They should be added once and never removed. In the first sequence of microworlds, there is no sequence of procedures that can avoid adding the constraints that express place-value encoding during the first transition, and dropping some of them during the second transition.

What is the closest possible procedure in a given microworld to the target procedure?

Just exactly how close to standard arithmetic procedures can procedures built around a particular representation of numbers, say Dienes Blocks, be made to be? Can a Dienes Block procedure be devised that is totally isomorphic to a standard written procedure? This is a question of interest to educators. For example, it bears on the question of just how much a child can learn about standard arithmetic by inventing a good arithmetic procedure in a given microworld, such as Dienes Blocks. This in turn bears on the question of how many microworlds, and which ones, are necessary to allow the student to easily converge upon the target skill. With a formal theory of analogy between procedures, we can now precisely determine how close the best possible procedure defined over a given microworld can be to the target procedure.

Take any procedure that uses the given representation of numbers. Examine the difference generator of the analogy between it and the target procedure (e.g. written addition). If this set contains constraints can not be met because of the basic physics of the representation, then one can not construct a model procedure that is isomorphic to the target procedure. An example should make this a little clearer.

A careful examination of the planning net has shown that it is impossible to construct a Dienes Block addition procedure whose analogy with written addition is perfect (i.e. an isomorphism). One constraint that is always present in Dienes Blocks involves the shape-value encoding that is the hallmark of Dienes Blocks. There is an encoding of the relationship between position and place-value that is present in both written addition and sorted Dienes Block addition, but it is redundantly coded by the visual appearance of Dienes Blocks. If one got rid of this redundancy by evening out the sizes of the blocks, then they wouldn't be Dienes Blocks anymore. So the redundancy is inherent in the representation, and will be part of the difference generator of the analogy to written addition no matter how clever one is about inventing Dienes Block addition procedures.

As a consequence, certain subtle *shifts in representation* which occur in the standard procedure for adding written numbers can not be duplicated in any Dienes Block addition procedure (see note 10 for details). This deficit gives some bite to the inherent incompleteness; the subtlety of these shifts makes them likely candidates for misunderstandings which Dienes Blocks are apparently helpless to prevent. This essential inadequacy can be directly diagnosed if not predicted using the theory of analogy between procedures.

In similar fashion, other microworlds can be evaluated. This evaluation is, however, quite constructive. Once the inherent mismatch with the target procedure has been indentified, the gap can be filled by modifying the microworld, or adding another microworld to the curriculum, if so desired.

In short, many of the same issues appear to be involved in the teaching teleology and discovery-based teaching. Planning nets seem to provide a formal tool for investigating this relationship further.

5. Conclusions

The major claim of this paper is that planning nets provide useful formalisms for capturing the teleologic semantics of procedures. However, probably the most important thought to take away from this exposition is the importance and utility of using *planning knowledge* in the deep structure analysis of procedures.

In contrast to other work on analogy, we have ignored the process of solving an analogy problem. Instead, we have concentrated on an intuitive determination of what representation most closely models the way experts conceive of procedures in order to understand analogies. This methodology has arrived at the same conclusion that was reached by completely different method. In particular, our planning nets are very similar to Sacerdoti's "procedural nets" (Sacerdoti 1977). Sacerdoti has shown his procedural nets to be a *sufficient* representation for designing procedures, and indeed much better than other known representations. We have tried to show a similar representation to be a *sufficient* representation for judging the closeness of analogy, and indeed much better than other known representations. In short, evidence is accumulating that planning net-like representations are good for many purposes. However, we should point out once again that neither Sacerdoti nor ourselves make any claims that the process of building a planning net, either for analogy or design, exactly models the human *process* of building a planning net.

Since teleologic knowledge is a part of a certain kind of expertise, one naturally wonders how it can be taught. Planning nets provide a precise framework for constructing explanations and curricula to explicate teleology. In particular, the formalism helps answer the question of how to sequence a set of "model" procedures, with certain formal properties. Moreover, many of these same formal properties seem useful in discovery learning curricula.

Our last comment should undoubtedly be that this research is just beginning. There are many deficiencies and questions that must be addressed. Reliable empirical measurements of closeness and transferability must be made. The general precision of the theory must be improved, and its inordinate amount of detail must be tamed, hopefully with the aid of a computer. In particular, we would like a complete, precise space of mp-morphisms for all five arithmetic operations. The limitations of the theory should be tested by exercising it on examples from other domains. In other words, this paper is more a proposal to investigate a promising line of thought than a report on completed research.

References

Allen, F. and J. Cocke "A Catalog of Optimizating Transformations", in *Design and Optimization of Compilers*, Randall Rustin ed., Prentice Hall, 1972

Brown, J.S. and R.R. Burton "Diagnostic Models for Procedural Bugs in Basic Mathematical Skills," *Cognitive Science*, Ablex Publishing, Volume 2, Pages 155-192, 1978.

Card, S.K. "Studies in the Psychology of Computer Text Editing Systems", Xerox PARC report No. SSL-78-1, Xerox Palo Alto Research Center, Palo Alto, California, 1978.

Evans, T.G. "A Program for the Solution of Geometric-Analogy Intelligence Test Questions", in *Semantic Information Processing*, M. Minsky ed., MIT Press, 1968.

Fischer, G., J.S. Brown, R.R. Burton "Aspects of a Theory of Simplification, Debugging, and Coaching", BBN Report No. 3912, ICAI Report No. 10, Bolt Beranek and Newman Inc., Cambridge, Massachusetts, 1978.

Greeno, J., G.R. Gelman, M.S. Riley. "Young Children's Counting and Understanding of Principles", in preparation.

Newell, A. & H.A. Simon. *Human Problem Solving*, Englewood Cliffs, New Jersey: Prentice-Hall, 1972.

Papert, S.A. "Computer-Based Microworlds as Incubators for Powerful Ideas," MIT AI Working Paper, Massachusetts Institute of Technology, Cambridge, Massachusetts, 1978.

Sacerdoti, E. *A Structure For Plans and Behavior*, The Artificial Intelligence Series. New York: Elsevier North-Holland, 1977.

Sternberg, R.J. "Componential Processes in Analogical Reasoning", *Psychological Review*, Vol. 84, No. 4, July 1977

Tversky, A. "Features of Similarity", *Psychological Review*, Vol. 84, No. 4, July 1977

Winston, P.H. "Learning by Creating and Justifying Transfer Frames" *Aritficial Intelligence*, Vo. 10 No. 2, April 1978

Winston, P.H. "Learning Structural Descriptions from Examples" in *The Psychology of Computer Vision*, P.H. Winston (ed.), New York, New York: McGraw-Hill, 1975

Notes

Note 1: It is safe to assume that individuals will differ in their judgements of the closeness of analogies. We take the position that this is due to the different deep structures that they assign to procedures. For example, someone who is just learning addition may not find the analogy between one-pass and two-pass addition particularly close. This might be due to a lack of distinct concepts for "carrying" and "column addition". So, how one understands a procedure affects the data that the theory of analogy will be verified against. Since we are interested in telcologic semantics and since telcologic understanding is a mark of expertise, it was important to use experts as subjects.

Note 2: Tversky (1977) weighted the features in the set A-B more heavily than the features in the set B-A in order to account for certain experimental data, e.g. that "Red China is similar to North Korea" has a lower degree of intuitive similarity than "North Korea is similar to Red China."

Note 3: The judgements of closeness are those of experts on arithmetic, and so can be taken to reflect the teleologic semantics of arithmetic.

Note 4: The folklore about protocol taking, supported by a few experiments (Card 1978), is: when in doubt, use a finer grain size. If the grain size is too large, one might miss distinctions. If one errs the other way, and makes the grain size too fine, then one creates more work for oneself, yet if one is tenacious, the relevant distinctions will ultimately appear, probably as relations between *groups* of actions instead of single, individual actions. So, it appears that the grain size issue (and a very similar issue, namely the choice of primitive actions) appears to be more a practical tradeoff than an insurmountable source of uncertainty in the theory.

Note 5: Dienes Block subtraction and other block subtraction procedures are usually taught using oral or written presentations of the problems. Thus, to solve *n-m*, the first step is to translate *n* into blocks, using some "bank" as a source of blocks. Next, one translates *m* into blocks, but *uses the first pile as the source.* When one is finished translating, the first pile contains *n-m* blocks. This procedure for doing block subtraction is so dissimilar to written subtraction that we have avoided using it in this paper.

Note 6: In formulating constraints, it is very important to put as little into each constraint as possible. For example, we could have replaced constraint 2 by "decrementing BOT by 1 must echoed by decrementing TOP by 1." Although adequate for base-1 subtraction, this is not the most general statement of the constraint, and indeed, this constraint would have to be replaced to handle Dienes Block subtraction. The basic idea is to split the declarative description of the world and the goal as finely as possible, so that small variations on the procedure can be modeled by

replacement of one constraint among many small ones, rather than modification of one clause of a large, special purpose constraint.

Note 7: We have glossed over a number of very difficult issues in the presentation of the planning steps. For instance, why was the TABLE chosen in Step 5 as the location for emptying the LH? How did we know not to empty it on TOP or BOT? Only the successful reasoning has been presented -- the alternatives that didn't work weren't mentioned. Most of the research in planning for robotics has gone into improving the search for correct plans by recognizing unworkable alternatives and recovering from them gracefully. All these difficult questions involving search can be ignored because we are only interested in *having* a correct planning net for a procedure, not in *finding* one.

Note 8: In the standard version of subtraction, where the carry loop is jammed together with the add-column loop, one must write $n + m$ digits, where n is the length of the longest addend, and m is the number of carries required (it is assumed that one writes a "1" above the columns one carries into). In the two-pass version, one must write $n + 2m$ digits: one must remember from the first pass which columns are overflowing, and this requires m notes to oneself, say in the form of writing a "1" above the overflowing column. The second m operations come from rewriting the answer digit of the columns that are carried into. There may be even more rewriting if the answer carried into is a 9.

Note 9: In the column carried into, the standard subprocedure requires adding three digits, one of which is of course the carried "1". But adding three digits requires remembering the sum of the first two digits while accessing the third digit. The two-pass subprocedure doesn't load memory this way, since the intermediate sum is written down instead.

Note 10: When one adds two large digits from a given column, one gets back a non-digit, e.g. "14". The first shift in representation is to break this number down into units and tens. Next, the units must be converted into a digit in the columns being added, while the tens must be converted into an argument to the carry subprocedure. In Dienes Block addition, the second conversion is superfluous, since the result of the column addition is already scaled up to the value of the column, so to speak. That is, an add in the tens column yields "140" in the form of 14 LONGs, not 14 UNITs.

Appendix 1

An explanation of the teleologic semantics of subtraction

To give a feel for how an explanation based on paths of minimal contrasting pairs of analogous procedures might go, an example of such a path is presented here. It begins with base-1 subtraction model, passes through some Dienes Block subtraction procedures, and ends with the standard procedure for subtraction of written numerals. Although reading these rather abbreviated descriptions can have nothing like the impact of actually handling the blocks and doing the procedures, the power of this technique to explain teleologic semantics should nonetheless be apparent.

Throughout the path, there is a certain ambivalence about the particular material that is used in the representation of number. In fact, the primitives and constraints used to describe and implement procedures really can't differentiate real, wooden Dienes Blocks from, say, drawings of Dienes Blocks, as long as they are manipulated the same way. In fact, there is no particular point where adoption of the constraints of the target procedure (written subtraction) forces us off the counting table and onto paper -- one can actually implement standard subtraction with cards bearing digits.

However, the material *does* make a difference to the efficiency metrics. In particular, some of the later constraints can only be motivated by assuming that erasing is more work than writing, which is true of paper, but hard to emulate with manipulatory materials.

We start with base-1 blocks because the mathematical semantics of this subtraction procedure are simple and concrete.

1. Polynomial Base-1 numerals are rather bulky for representing large numbers. One solution to the block management problem is to let some counters stand for a fixed number of the unit counters. This is the *polynomial* constraint (3 in the text). The next procedure of this mp-morphism is a simple version of big-pile Dienes Block subtraction.

2. Search Instead of Random Choice This mp-morphism adds the notion that searching for two blocks of the same shape is more efficient than picking two blocks at random, then trading to make them the same shape.

3. Chose Larger to Trade Down The idea here is to trade down the larger of the two blocks. If one picks an arbitrary block to trade down, but not the unit block, then eventually one will be able to match their shapes, but it will often take more trading than always picking the larger one to trade down. This range procedure requires memorization of which of two shapes stands for a

larger *multiplier*.

4. Search for Next Larger Before Trading When one can't find two blocks of equal shape, and instead has two blocks of unequal shape, then before trading down the larger one, replace it with a block that is the *next size larger* than the smaller block. If the search succeeds, one only has to trade down once. This plan step requires memorizing which shape is the next larger one than a given shape.

5. Choose TOP to Trade Down This model is motivated by observing that when the block that is traded down comes from BOT (the bottom numeral), the subtraction as a whole takes more time than it would if the block had come from TOP (the numeral that is being subtracted from). When a block from BOT is traded down, the nine smaller blocks that are left over go back into BOT. So the main loop must run nine times more. If a block comes from TOP, the nine extras go back into TOP. If BOT runs out soon, they may never be touched. So trading down a block from TOP is more efficient than trading down a block from BOT.

The goal of choosing TOP blocks creates a subgoal that the TOP block be larger than the BOT block. This subgoal is satisfied by a subgraph which is already a part of the domain planning net, namely, the union of the subgraphs generated by models 2, 3 and 4. So, the new part of the planning net underlying this procedure is just the part that satisfies the goal "chose BOT block" exclusive of the part which satisfies the subgoal.

6. Canonicity This constraint was described in the text.

7. Base Ten The canonicity constraint produces a trading pattern which is much easier to remember if all the multipliers are powers of ten (or some other base). For example, in canonical American money, which is a polynomial representation but not a base-10 representation of number, a citizen would canonicalize their pocket change by trading in five pennies for a nickel, two nickels for a dime, three dimes for a quarter and a nickel, etc.

8. Sort by Power Canonicalization (= carrying) and decanonicalization (= borrowing) are somewhat easier if numerals are sorted so that all counters of a certain power are accessible at once. Dienes Blocks, as we observed them being used in schools, lacked this constraint. In fact, Dienes Blocks also lack the canonical and base-10 constraints as well. However, teachers usually require their students to obey these two.

9. Power Represented by Location Only Numerals must take up space, either on table tops or on paper. Once powers are sorted, location in space redundantly represents the power of a counter. In this mp-morphism, that redundancy is removed by making all coefficient tokens (i.e. "digits") look the same, regardless of the power. The abacus, for example, obeys this constraint. This

allows one to represent much larger numbers, since one need not invent new token shapes when one needs to use a new, higher power. That is, one can make an abacus of arbitrary width, but Dienes Blocks, which are inherently unable to obey this constraint, are limited in practice to at most four powers.

10. Zero To use location to represent power, a prearranged pattern of locations must be used. But, such fixed patterns, like the abacus or columnar ruled paper, can't represent numbers that are larger than they have been designed to represent. Moreover, producing the patterns accurately is difficult to do free hand. A good solution to this problem is represent power with relative locations, which amounts to using *zero* as a place holder. A "relative-location abacus" could be built which lays out piles of beads in a line on the table; it would use a clear plastic bead as a place holder and piles of colored beads as non-zero "digits."

11. Alignment In setting up the subtraction problem, one insists that the numerals be aligned so that digits of the same power are in the same column. This reduces the effort necessary to locate the digits of matching power when subtracting.

12. Non-countable Coefficients It is quicker to arrange counters on a table or write coefficients symbols on paper if the number of counters or strokes is small. This motivates replacing countable coefficients with symbolic ones (e.g., digits). However, with symbolic coefficients, the PICK/FROM operation is radically altered. It is no longer possible to decrement a coefficient by picking up a piece of it (i.e., picking up a block or erasing a hash mark). Instead, a decrementation table must be memorized. That is, one must be able to count backwards from twenty.

There is no particular point where the target constraints force us off the counting table and onto paper. Manipulatory systems can be devised which use non-countable coefficients. One such manipulatory system is simply a set of cards bearing digits, which are laid out in a line on the table.

13. Memorize Pairing Off The next few minimal contrasting models are designed to minimize the manipulation of the cards in a manipulatory system, or erasing a digit and writing a new one in a written system. In the previous number systems, column subtraction was realized by pairing off decrements of the top and bottom digits. A "movie" of the card procedure doing 15-3 would be

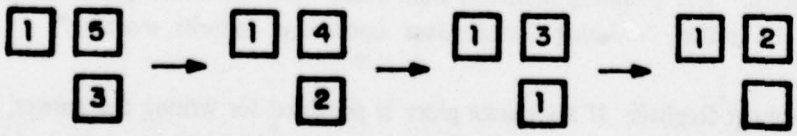

This model replaces this pairing off loop with a table lookup. A "movie" of the modified card procedure doing 25-7 is

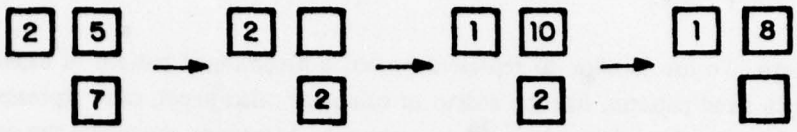

14. Memorize Comparison This model procedure replaces the two step borrowing (see movie above) with a one step borrow by looking ahead. That is, it looks ahead to see which digit will be zero -- the top or the bottom. This amounts to memorizing the greater-than table for digits. Now the movie for 25-7 is

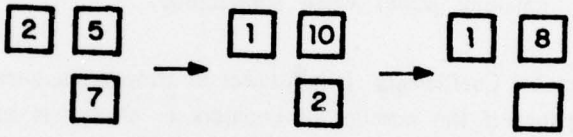

15. Memorize Teens Facts Two table lookups can be reduced to one, and two digit rewrites can be saved if a new facts table is provided for the teens facts. The new table is 10 by 9 and contains facts like 15-7=8. The movie reduces to

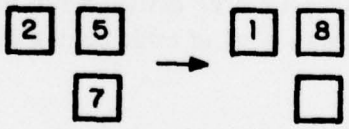

16. Sequence Columns In the previous systems, columns are processed in random order. However, this necessitates marking the columns that are done by zeroing the bottom digit. This digit rewrite can be saved if the columns are processed in some set order -- either left to right or vice versa. The planning heuristic, that is the right difference generator of this mp-morphism, could be called "ordering independent operations reduces marking".

17. Answer Register If a separate place is provided for writing the answer, then erasures of the

top digits can be reduced. This is motivated by the fact that writing a digit is easier than erasing -- a property peculiar to paper.

18. Right to Left If the columns are processed right to left, one borrows from the top digit. If the columns are processed left to right, one borrows from the answer. The numeral that gets borrowed from ends up with erasures, while the other one has no erasures. If one erases by scratching out the digit and writing the new digit above, then the numeral that's borrowed from can become a real mess. The motivation for this analogy is that there is more need for the answer numeral to be legible than the top numeral. Hence, subtraction is more efficient if one processes the columns from right to left.

At last, we have arrived at the standard subtraction algorithm via a sequence of procedures/models where each model in this sequence has a mp-morphism between it and its immediate successor thus creating a well structured sequence of analogous models converging to the desired target procedure.

Navy

1 Dr. Ed Aiken
Navy Personnel R&D Center
San Diego, CA 92152

1 Dr. Jack R. Borsting
Provost & Academic Dean
U.S. Naval Postgraduate School
Monterey, CA 93940

1 Dr. Robert Breaux
Code N-71
NAVTRAEQUIPCEN
Orlando, FL 32813

1 DR. PAT FEDERICO
NAVY PERSONNEL R&D CENTER
SAN DIEGO, CA 92152

1 Dr. John Ford
Navy Personnel R&D Center
San Diego, CA 92152

1 Dr. Steve Harris
Code L522
NAMRL
Pensacola FL 32508

1 Dr. Norman J. Kerr
Chief of Naval Technical Training
Naval Air Station Memphis (75)
Millington, TN 38054

1 CAPT Richard L. Martin
USS Francis Marion (LPA-249)
FPO New York, NY 09501

2 Dr. James McGrath
Navy Personnel R&D Center
Code 306
San Diego, CA 92152

1 DR. WILLIAM MONTAGUE
LRDC
UNIVERSITY OF PITTSBURGH
3939 O'HARA STREET
PITTSBURGH, PA 15213

Navy

1 Library
Navy Personnel R&D Center
San Diego, CA 92152

6 Commanding Officer
Naval Research Laboratory
Code 2627
Washington, DC 20390

1 JOHN OLSEN
CHIEF OF NAVAL EDUCATION &
TRAINING SUPPORT
PENSACOLA, FL 32509

1 Psychologist
ONR Branch Office
495 Summer Street
Boston, MA 02210

1 Psychologist
ONR Branch Office
536 S. Clark Street
Chicago, IL 60605

1 Code 436
Office of Naval Research
Arlington, VA 22217

1 Office of Naval Research
Code 437
800 N. Quincy SStreet
Arlington, VA 22217

5 Personnel & Training Research Programs
(Code 458)
Office of Naval Research
Arlington, VA 22217

1 Psychologist
OFFICE OF NAVAL RESEARCH BRANCH
223 OLD MARYLEBONE ROAD
LONDON, NW, 15TH ENGLAND

1 Psychologist
ONR Branch Office
1030 East Green Street
Pasadena, CA 91101

Navy

1 Scientific Director
Office of Naval Research
Scientific Liaison Group/Tokyo
American Embassy
APO San Francisco, CA 96503

1 Head, Research, Development, and Studies
(OP102X)
Office of the Chief of Naval Operations
Washington, DC 20370

1 Scientific Advisor to the Chief of
Naval Personnel (Pers-Or)
Naval Bureau of Personnel
Room 4410, Arlington Annex
Washington, DC 20370

1 DR. RICHARD A. POLLAK
ACADEMIC COMPUTING CENTER
U.S. NAVAL ACADEMY
ANNAPOLIS, MD 21402

1 A. A. SJOHOLM
TECH. SUPPORT, CODE 201
NAVY PERSONNEL R& D CENTER
SAN DIEGO, CA 92152

1 Dr. Alfred F. Smode
Training Analysis & Evaluation Group
(TAEG)
Dept. of the Navy
Orlando, FL 32813

1 CDR Charles J. Theisen, JR. MSC, USN
Head Human Factors Engineering Div.
Naval Air Development Center
Warminster, PA 18974

1 W. Gary Thomson
Naval Ocean Systems Center
Code 7132
San Diego, CA 92152

1 Dr. Ronald Weitzman
Department of Administrative Sciences
U. S. Naval Postgraduate School
Monterey, CA 93940

Army

1 HQ USAREUE & 7th Army
ODCSOPS
USAAREUE Director of GED
APO New York 09403

1 DR. RALPH DUSEK
U.S. ARMY RESEARCH INSTITUTE
5001 EISENHOWER AVENUE
ALEXANDRIA, VA 22333

1 Dr. Ed Johnson
Army Research Institute
5001 Eisenhower Blvd.
Alexandria, VA 22333

1 Dr. Michael Kaplan
U.S. ARMY RESEARCH INSTITUTE
5001 EISENHOWER AVENUE
ALEXANDRIA, VA 22333

1 Dr. Harold F. O'Neil, Jr.
ATTN: PERI-OK
5001 EISENHOWER AVENUE
ALEXANDRIA, VA 22333

Air Force

1 Dr. Marty Rockway (AFHRL/TT)
Lowry AFB
Colorado 80230

1 Jack A. Thorpe, Capt, USAF
Program Manager
Life Sciences Directorate
AFOSR
Bolling AFB, DC 20332

Marines

1 Director, Office of Manpower Utilization
HQ, Marine Corps (MPU)
BCE, Bldg. 2009
Quantico, VA 22134

1 MCDEC
Quantico Marine Corps Base
Quantico, VA 22134

1 DR. A.L. SLAFKOSKY
SCIENTIFIC ADVISOR (CODE RD-1)
HQ, U.S. MARINE CORPS
WASHINGTON, DC 20380

CoastGuard

1 MR. JOSEPH J. CCWAN, CHIEF
PSYCHOLOGICAL RESEARCH (G-P-1/62)
U.S. COAST GUARD HQ
WASHINGTON, DC 20590

Other DoD

1 Dr. Stephen Andriole
ADVANCED RESEARCH PROJECTS AGENCY
1400 WILSON BLVD.
ARLINGTON, VA 22209

12 Defense Documentation Center
Cameron Station, Bldg. 5
Alexandria, VA 22314
Attn: TC

1 Dr. Dexter Fletcher
ADVANCED RESEARCH PROJECTS AGENCY
1400 WILSON BLVD.
ARLINGTON, VA 22209

1 Military Assistant for Training and
Personnel Technology
Office of the Under Secretary of Defense
for Research & Engineering
Room 3D129, The Pentagon
Washington, DC 20301

Civil Govt

1 Dr. Susan Chipman
Basic Skills Program
National Institute of Education
1200 19th Street NW
Washington, DC 20208

1 Dr. Joseph I. Lipson
Division of Science Education
Room W-638
National Science Foundation
Washington, DC 20550

1 Dr. John Mays
National Institute of Education
1200 19th Street NW
Washington, DC 20208

1 Dr. Arthur Melmed
National Intitute of Education
1200 19th Street NW
Washington, DC 20208

1 Dr. Andrew R. Molnar
Science Education Dev.
and Research
National Science Foundation
Washington, DC 20550

1 Dr. Thomas G. Sticht
Basic Skills Program
National Institute of Education
1200 19th Street NW
Washington, DC 20208

1 Dr. Joseph L. Young, Director
Memory & Cognitive Processes
National Science Foundation
Washington, DC 20550

Non Govt

1 Dr. Earl A. Alluisi
HQ, AFHRL (AFSC)
Brooks AFB, TX 78235

1 Dr. John R. Anderson
Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15213

1 DR. MICHAEL ATWOOD
SCIENCE APPLICATIONS INSTITUTE
40 DENVER TECH. CENTER WEST
7935 E. PRENTICE AVENUE
ENGLEWOOD, CO 80110

1 1 psychological research unit
Dept. of Defense (Army Office)
Campbell Park Offices
Canberra ACT 2600, Australia

1 Dr. Alan Baddeley
Medical Research Council
Applied Psychology Unit
15 Chaucer Road
Cambridge CB2 2EF
ENGLAND

1 Dr. Nicholas A. Bond
Dept. of Psychology
Sacramento State College
600 Jay Street
Sacramento, CA 95819

1 Dr. Lyle Bourne
Department of Psychology
University of Colorado
Boulder, CO 80302

1 Dr. Kenneth Bowles
Institute for Information Sciences
University of California at San Diego
La Jolla, CA 92037

1 Dr. John S. Brown
XEROX Palo Alto Research Center
3333 Coyote Road
Palo Alto, CA 94304

Non Govt

1 DR. C. VICTOR BUNDERSON
WICAT INC.
UNIVERSITY PLAZA, SUITE 10
1160 SO. STATE ST.
OREM, UT 84057

1 Charles Myers Library
Livingstone House
Livingstone Road
Stratford
London E15 2LJ
ENGLAND

1 Dr. William Chase
Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15213

1 Dr. Micheline Chi
Learning R & D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

1 Dr. Allan M. Collins
Bolt Beranek & Newman, Inc.
50 Moulton Street
Cambridge, Ma 02138

1 Dr. Meredith Crawford
Department of Engineering Administration
George Washington University
Suite 805
2101 L Street N. W.
Washington, DC 20037

1 Dr. Hubert Dreyfus
Department of Philosophy
University of California
Berkely, CA 94720

1 MAJOR I. N. EVONIC
CANADIAN FORCES PERS. APPLIED RESEARCH
1107 AVENUE ROAD
TORONTO, ONTARIO, CANADA

Non Govt

1 Dr. Ed Feigenbaum
Department of Computer Science
Stanford University
Stanford, CA 94305

1 Mr. Wallace Feurzeig
Bolt Beranek & Newman, Inc.
50 Moulton St.
Cambridge, MA 02138

1 Dr. Victor Fields
Dept. of Psychology
Montgomery College
Rockville, MD 20850

1 Dr. John R. Frederiksen
Bolt Beranek & Newman
50 Moulton Street
Cambridge, MA 02138

1 DR. ROBERT GLASER
LRDC
UNIVERSITY OF PITTSBURGH
3939 O'HARA STREET
PITTSBURGH, PA 15213

1 Dr. Ira Goldstein
XEROX Palo Alto Research Center
3333 Coyote Road
Palo Alto, CA 94304

1 Dr. Ron Hambleton
School of Education
University of Massechusetts
Amherst, MA 01002

2 Dr. Barbara Hayes-Roth
The Rand Corporation
1700 Main Street
Santa Monica, CA 90406

1 Library
HumRRO/Western Division
27857 Berwick Drive
Carmel, CA 93921

Non Govt

1 Dr. Earl Hunt
Dept. of Psychology
University of Washington
Seattle, WA 98105

1 Mr. Gary Irving
Data Sciences Division
Technology Services Corporation
2811 Wilshire Blvd.
Santa Monica CA 90403

1 Dr. Steven W. Keele
Dept. of Psychology
University of Oregon
Eugene, OR 97403

1 Dr. Walter Kintsch
Department of Psychology
University of Colorado
Boulder, CO 80302

1 Dr. David Kieras
Department of Psychology
University of Arizona
Tuscon, AZ 85721

1 Mr. Marlin Kroger
1117 Via Goleta
Palos Verdes Estates, CA 90274

1 LCOL. C.R.J. LAFLEUR
PERSONNEL APPLIED RESEARCH
NATIONAL DEFENSE HQS
101 COLONEL BY DRIVE
OTTAWA, CANADA K1A 0K2

1 Dr. Jill Larkin
SESAME
c/o Physics Department
University of California
Berkely, CA 94720

1 Dr. Alan Lesgold
Learning R&D Center
University of Pittsburgh
Pittsburgh, PA 15260

Non Govt

1 Dr. Michael Levine
Department of Psychology
University of Illinois
Champaign, IL 61820

1 Dr. Robert A. Levit
Manager, Behavioral Sciences
The BDM Corporation
7915 Jones Branch Drive
McClean, VA 22101

1 Dr. Mark Miller
Systems and Information Sciences Laborat
Central Research Laboratories
TEXAS INSTRUMENTS, INC.
Mail Station 5
Post Office Box 5936
Dallas, TX 75222

1 Dr. Richard B. Millward
Dept. of Psychology
Hunter Lab.
Brown University
Providence, RI 02912

1 Dr. Allen Munro
Univ. of So. California
Behavioral Technology Labs
3717 South Hope Street
Los Angeles, CA 90007

1 Dr. Donald A Norman
Dept. of Psychology C-009
Univ. of California, San Diego
La Jolla, CA 92093

1 Dr. Seymour A. Papert
Massachusetts Institute of Technology
Artificial Intelligence Lab
545 Technology Square
Cambridge, MA 02139

1 Dr. James A. Paulson
Portland State University
P.O. Box 751
Portland, OR 97207

Non Govt

1 MR. LUIGI PETRULLO
2431 N. EDGEWOOD STREET
ARLINGTON, VA 22207

1 DR. PETER POLSON
DEPT. OF PSYCHOLOGY
UNIVERSITY OF COLORADO
BOULDER, CO 80302

1 Dr. Peter B. Read
Social Science Research Council
605 Third Avenue
New York, NY 10016

1 Dr. Fred Reif
SESAME
c/o Physics Department
University of California
Berkely, CA 94720

1 Dr. Ernst Z. Rothkopf
Bell Laboratories
600 Mountain Avenue
Murray Hill, NJ 07974

1 Dr. Allen Schoenfeld
SESAME
c/o Physics Department
University of California
Berkely, CA 94720

1 Dr. Robert Sternberg
Dept. of Psychology
Yale University
Box 11A, Yale Station
New Haven, CT 06520

1 DR. ALBERT STEVENS
BOLT BERANEK & NEWMAN, INC.
50 MOULTON STREET
CAMBRIDGE, MA 02138

1 DR. PATRICK SUPPES
INSTITUTE FOR MATHEMATICAL STUDIES IN
THE SOCIAL SCIENCES
STANFORD UNIVERSITY
STANFORD, CA 94305

Non Govt

1 Dr. John Thomas
IBM Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

1 DR. PERRY THORNDYKE
THE RAND CORPORATION
1700 MAIN STREET
SANTA MONICA, CA 90406

1 Dr. David J. Weiss
N660 Elliott Hall
University of Minnesota
75 E. River Road
Minneapolis, MN 55455

1 Dr. Karl Zinn
Center for research on Learning
and Teaching
University of Michigan
Ann Arbor, MI 48104