MASSACHUSETTS INSTITUTE OF TECHNOLOGY

ARTIFICIAL INTELLIGENCE LABORATORY

AIM-414 April 1977

LEARNING BY HYPOTHESIZING AND JUSTIFYING TRANSFER FRAMES

Patrick H. Winston

ABSTRACT

Learning is defined to be the computation done by a student when there is a transfer of information to him from a teacher. In the particular kind of learning discussed, the teacher names a *source* and a *destination*. In the sentence, "Robbie is like a fox," *fox* is the source and *Robbie* is the destination. The student, on analyzing the teacher's instruction, computes a kind of filter called a *transfer frame*. It stands between the source and the destination and determines what information is allowed to pass from one to the other.

Computing the transfer frame requires two steps: *hypothesis* and *evaluation*. In the hypothesis step, potentially useful transfer frames are produced through an analysis of the information in the source and its immediate relatives. For Robbie, a robot, the way it compares with other robots would be noted. In the evaluation step, the better of the hypothesized frames are selected through a study of the destination frame, its relatives, and the general context.

Some source-destination pairs may be generated by the student acting alone. There is also the possibility of making notes that are useful in deciding if conclusion makes sense.

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's artificial intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-75-C-0643.

THE PROBLEM

Learning remains an enigma. In spite of strong efforts by good people, we have barely scratched the surface. It is hard to write computer programs that can learn even simple things. It is even hard to be precise about what learning is.

For this paper, learning is defined to be the computation done by a student when there is a transfer of information to him from a teacher. Normally, both must do some work. The amount of work done by the two participants in the transfer can vary between two extreems, however. As illustrated in figure 1, there is a spectrum starting with learning by being told, extending through learning by studying samples, and ending with learning by self-sufficient discovery. To some people, only learning by discovery counts as legitimate learning, but such a posture seems extreme.

This paper concentrates on middle ground in the vicinity of learning by being told. It offers a theory of learning by hypothesizing and evaluating certain structures that will be called transfer frames.

The methodology. Since learning is such a broad, complex phenomenon, it is sensible to be very precise about the nature of the attack. This is an adptation of the approach used by Marr in his fundatmental work on vision [Marr]:

- First, it is necessary to observe or define some learning competance to be understood.

- Second, a representation should be selected or invented that is capable of capturing the knowledge to be learned.

- Third, the first and second items should be translated into a precisely defined computation problem to be solved.

- Fourth, algorithms should be devised that preform the desired computation.

- And fifth, the results so far should be validated either by successful computer implementation and experimentation or by appropriate psychological inquiries.

All this seems obvious, but there are strong temptations that often throw research out of proper perspective. One is to be caught up with an attraction to a particular representation. Worse yet, there may be an attachment to some particular algorithm, with a corrolary failure to understand that many algorithms usually can be devised once a computational problem is properly laid out.

Therefore, let us begin with a synopsis concentrating on the definition of a kind of learning competance and on the selection of a representation that seems appropriate to it. Then we will turn to the details of the algoritms so far devised, implemented on a computer, and experimented with.

Defining the competance to be understood. Consider the following statement:

Robbie is a robot.

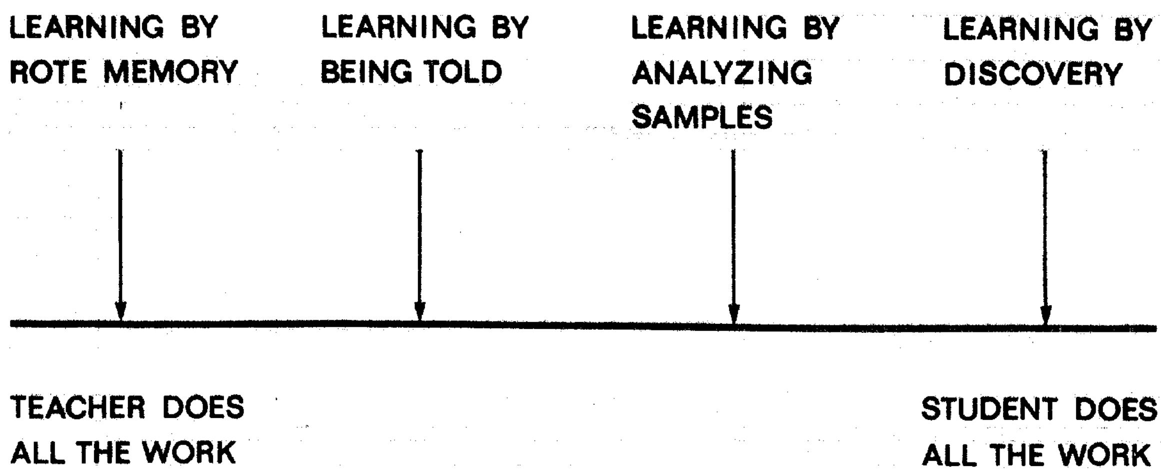

Figure 1. Learning may involve little work on the part of the learner or a lot. For there to be genuine learning, most people demand that the learner activly participate in the learning process, The simplest learning, really not learning at all, is learning by being programmed, with the learner doing nothing save submitting to the program surgery performed by the teacher. Learner participation begins when the learning is by being told or by understanding a series of samples. In the extreme, the participation of the learner is total, to the exclusion of the teacher, and there is learning by self-discovery.

Hearing this, a human or computer student should assume some facts about Robbie and becomes curious about others. Robbie is probably made of metal, and it would be interesting to know if he is intelligent.

Now consider these:

Robbie has a very high degree of cleverness.

Robbie is clever.

Robbie is clever like a fox.

Robbie is like a fox.

All convey approximately, but not exactly the same idea. The first is the most precise. The second differs very little from it. The third adds some nauance to the meaning conveyed. And the forth requires extra work for understanding since cleverness is not mentioned explicitely.

Why do we use "Robbie is like a fox," instead of "Robbie has a very high degree of cleverness?" It is not just a form of shorthand:

- The teacher conveys more by using examples.
- The teacher cannot know how explicit he must be since he cannot have a perfect model of the student.
- The teacher may be unable to exactly articulate what he knows. He may have a piece of knowlege without being able to present it according to the particular form required for conversational instruction. He may have to allude, rather than state.

In addition, other things happen. The student may become curious about whether Robbie is like a fox in other ways. Given that Robbie is clever, the student may wonder if Robbie is also like Suzie, another robot already known to be clever. Still another possibility is that the student may wonder if Robbie is like something or someone known to be the quintescence of cleverness.

In summary, then, the competance to be understood is the competance to absorb both explicit and simile-like instruction and the competance to be curious on learning new information.

The representation. What is to be the representation selected to be the target of learning as so circumscribed? That is to say, what conventions about symbols and their arrangement are suited to capturing the knowledge to be learned. Of the many representations available now, the frames representation seems best suited in terms of the point of view it encourages.

Roughly, a frame is a collection of properties. Here, for example, is a frame describing a fox:

FRAME NAME	SLOT	VALUE
FOX	A-KIND-OF	SMALL-MAMMAL
	COLOR	RED
	CLEVERNESS	VERY-HIGH

The frame name identifies what is to be described. Each of the properties that constitute the description is conveyed by a so-called slot-value combination. Each slot name specifies a property and the value associated with a slot dictates what is known about the corresponding property for the thing described by the frame.

From a programming point of view, this use of the frame representation scheme is exactly like an ordinary property-list representation. Indeed, the frame idea can be defined as a generalization of the property list, and the points of generalization, as programming mechanisms, are not exploited here. The frames vocabulary is used, nevertheless, because Minsky's original paper brought about a certain point of view [Minsky]. This point of view contributed strongly to the way of looking at learning offered here.

One ready objection is that programs using a frame representation can learn nothing that is not expressible in terms of frames. This seems true, but not confining. The world of objects, classes, and even processes that can be described in terms of frames seems amply large for useful learning research.

The computational problem. The key computational problem, therefore, is to fill frame slots using information given by a teacher either explicitely or in the form of simile-like instructions.

An algorithm. Here is the essense of an algorithm, to be described in detail later, that accomplishes the computation required to deal with simile-like instruction:

- The teacher, names a *source frame* and a *destination frame*. In the sentence, "Robbie is like a fox," *fox* is the source and *Robbie* is the destination. The teacher may or may not specify the exact slots in which the source and destination have the same values. He may tell the student that Robbie and a fox are alike with respect to cleverness or he may just say that they are alike.

- The student, on analyzing the teacher's instruction, computes a *transfer frame*. The transfer frame is a filter. It stands between the source and the destination as in figure 2, determining exactly what slot-value combinations are allowed to pass from one to the other.

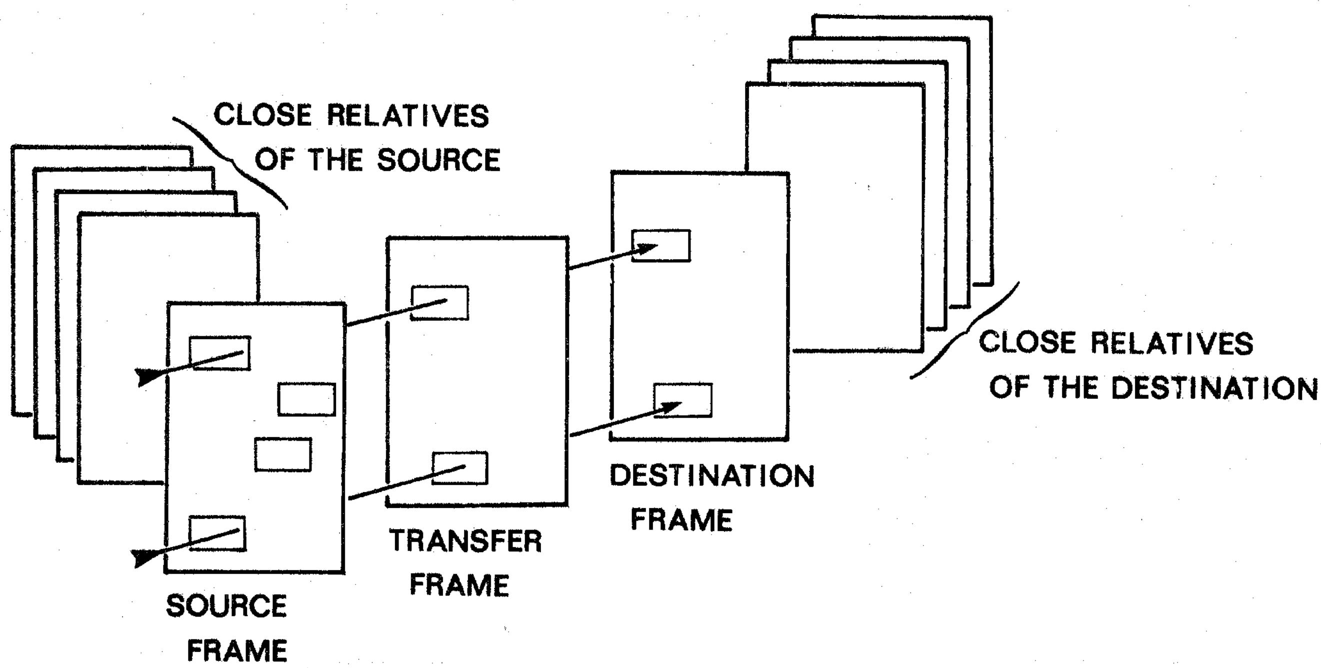

Figure 2. The basic idea behind learning using transfer frames. The teacher specifies a source and a destination and possibly the slots that are relevant. The student analyzes the source, the destination, and other aspects of the situation to discover and use a transfer frame.

- Computing the transfer frame requires two steps: *hypothesis* and *evaluation*. In the hypothesis step, potentially useful transfer frames are produced through an analysis of the information in the source frame and its immediate relatives. For Robbie, a robot, the way it compares with other robots would be noted. For a fox, other small common forest mammals would be used. In the evaluation step, the better of the hypothesized frames are selected through a study of the destination frame, its relatives, and the general context that exists by way of previous instruction.

This preview is given only to provide a flavor. Much more will be said about these procedures as well as others that deal with justification of transfers and internal generation of transfer possibilities.

Validation. The procedures described in this paper have been implemented and tested on the examples to be given. Exceptions are clearly noted. No claim is made about psychological validation, however. When the words *teacher* and *student* are used, the following is to be understood:

- The *teacher* is a human instructor.

- The *student* is an experimental collection of algorithms implemented as computer programs.

The programs are in LISP. Listings are available.

In a moment, we will look at the details of a running program that performs some simple learning that is in accord with the points of competence proposed here. To keep our own knowledge from getting too much in the way of thinking about the ideas, a semantically deprived world is used for the explanation. One consequence is that we too will have to work at understanding what is to be learned.

HYPHOTHESIS AND EVALUATION

If the source, destination, and transfer frame are given, there is nothing left to do but rush the slot values through the transfer frame. But since the transfer frame is usually not given, the learner must do some work to dig out the meaning and acquire new knowledge. It is this active participation of the listener that makes the learning interesting. To illustrate how transfer frames can be hypothesized and evaluated, we now look at some very simple examples from the blocks world shown in figure 3 and figure 4. This world is used specifically to make it easy to construct examples that illustrate all of the methods. Note that Figure 4 shows how the concepts are linked by AKO relationships, short for A-KIND-OF. INSTANCE is the opposite of AKO.

There Are Several Ways To Hypothesize Transfer Frames

Transfer frame hypothesizing begins by collecting together all of the slots in the source frame which are filled with the values VERY-LOW or VERY-HIGH. The theory is

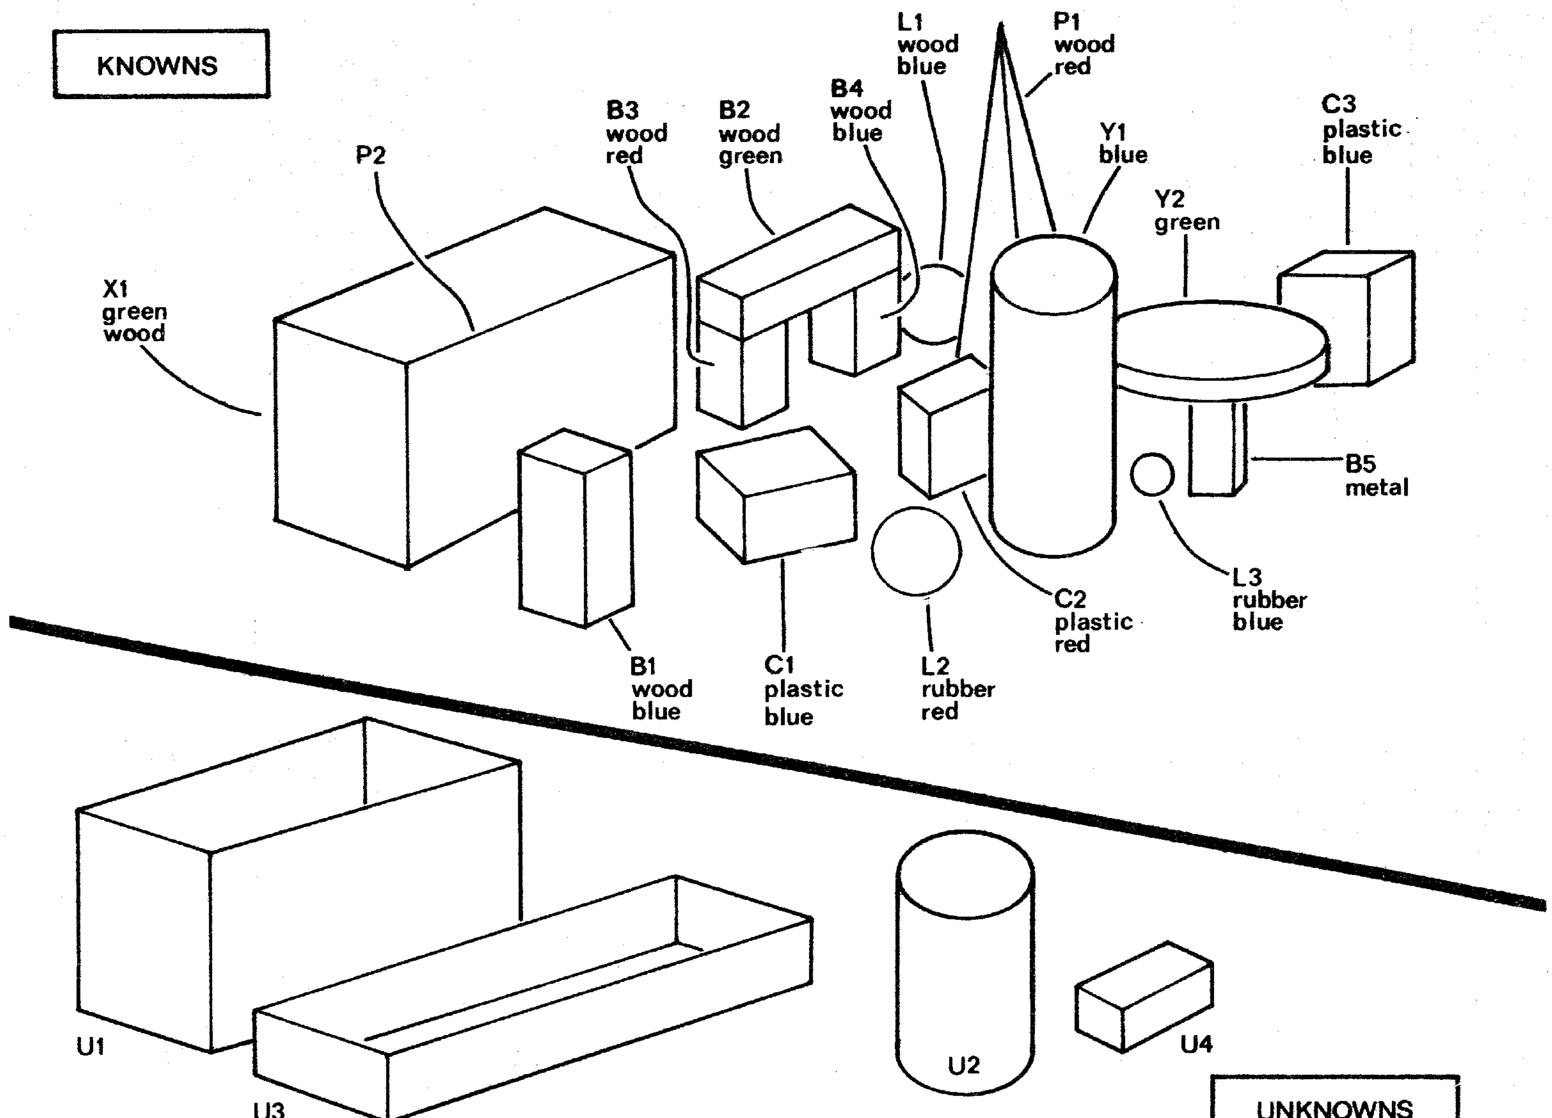

Figure 3. The objects in the blocks world.

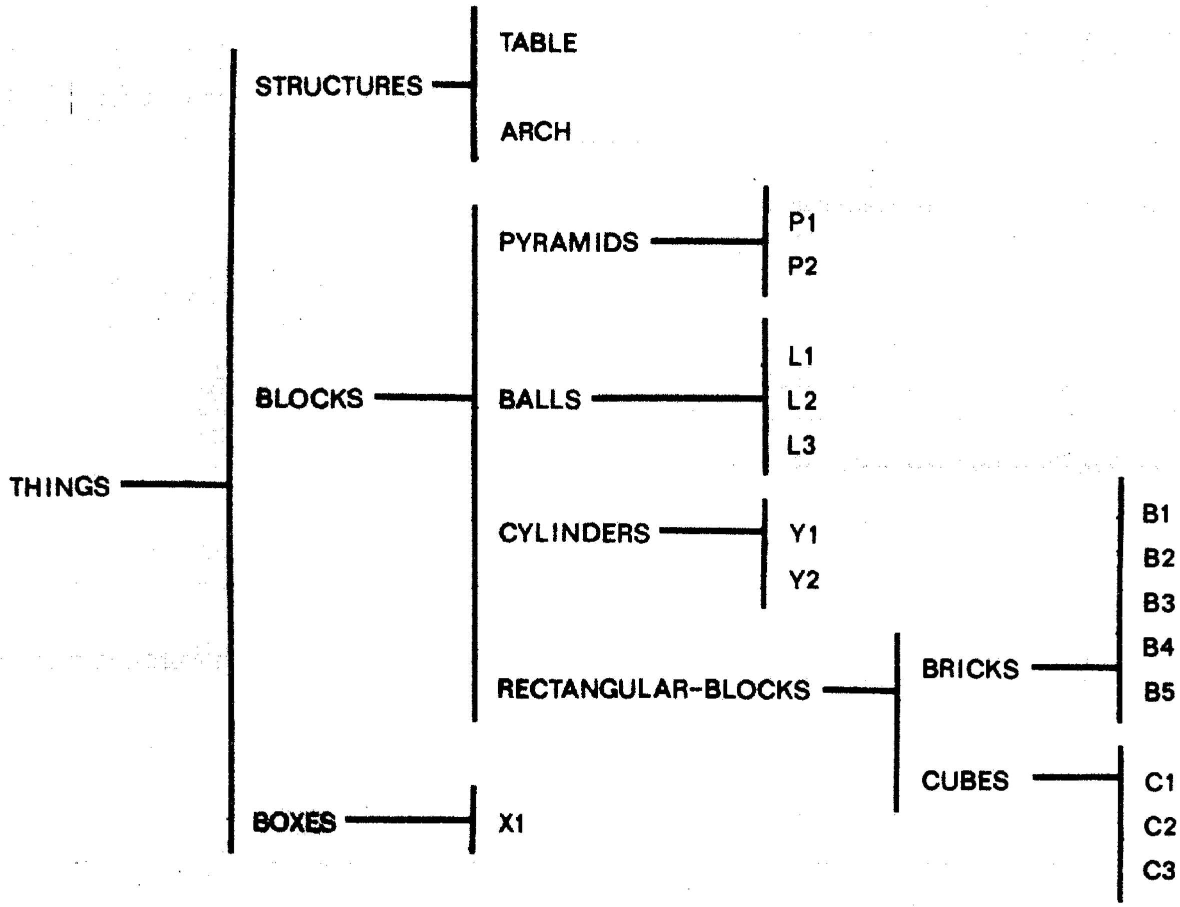

Figure 4. The hierarchical organization of the blocks world. The structure reflects how concepts are linked by the A-KIND-OF relation.

that concepts which exhibit properties to an unusual degree are potentially good sources for those properties. Suppose, for example, we have the following instruction:

Ul is like Pl.

To understand this, the student looks at the frame for Pl:

P1	AKO	PYRAMID
	HEIGHT	VERY-HIGH
	COLOR	RED
	MATERIAL	WOOD

Clearly the only slot with a VERY-HIGH value is HEIGHT. This is therefore transferred to Ul using the following transfer frame:

TF-HEIGHT	TRANSFER-SLOTS	HEIGHT

If Ul had a known height already, then the propose transfer frame would have been rejected immediately. If the first method fails to find a viable transfer frame, others are tried until one works.

The next method again searches for important slots, but this time on the basis of global knowledge. Slots whose own descriptive frames contain VERY-HIGH in their IMPORTANCE slots are deemed globally important, and they are all collected. The slot PURPOSE, for example, is globally important. Consequently the following results in learning that Ul is for storage.

Ul is like Xl.

Inspection of the Xl and PURPOSE frames shows why:

X1	AKO	BOX
	COLOR	GREEN
	MATERIAL	WOOD
	PURPOSE	STORAGE
PURPOSE	AKO	FUNCTIONAL-PROPERTY
	IMPORTANCE	VERY-HIGH

Having dispensed with slots filled with exceptional values and slots known to be globally important, the next method concentrates on slots which are unusual for concepts in the same class as the source. Thus the slot MATERIAL, found in Ll, would be judged important there because there are three balls, Ll, L2, and L3, and of these, only Ll has WOOD in the MATERIAL slot, which for balls is unusual:

L1	AKO	BALL
	SIZE	MEDIUM
	COLOR	BLUE
	MATERIAL	WOOD
L2	AKO	BALL
	SIZE	MEDIUM
	COLOR	RED
	MATERIAL	RUBBER
L3	AKO	BALL
	SIZE	MEDIUM
	COLOR	BLUE
	MATERIAL	RUBBER

Consequently the following results in knowing that U1 has WOOD in the MATERIAL slot:

> U1 is like L1.

Now suppose that we move to U2 and offer the following information:

> U2 is a CYLINDER.
> U2 is like B1.

B1, unfortunately, is rather undistinguished:

B1	AKO	BRICK
	SIZE	MEDIUM
	COLOR	BLUE
	MATERIAL	WOOD

Consequently, none of the hypothesizing methods given so far find anything, and the learner must simply gather up all the slots.

Note that after all of the slots are collected, they could be assembled together into a single transfer frame, but it seems better to group them together according to the property categories involved. This is true no matter what hypothesizing method is used to collect them. Thus B1's SIZE, COLOR, and MATERIAL slots, none of which are closely related as figure 5 shows, form three corresponding transfer frames.

There Are Several Ways To Evaluate The Transfer Frames

It is now up to the evaluation methods to determine which transfer frame to use. Several of these methods examine relatives of the destination, looking carefully for evidence that can pull the better transfer frames out of the pack. Consequently, it is important to know that U2 is a kind of cylinder and that Y1 and Y2 are too:

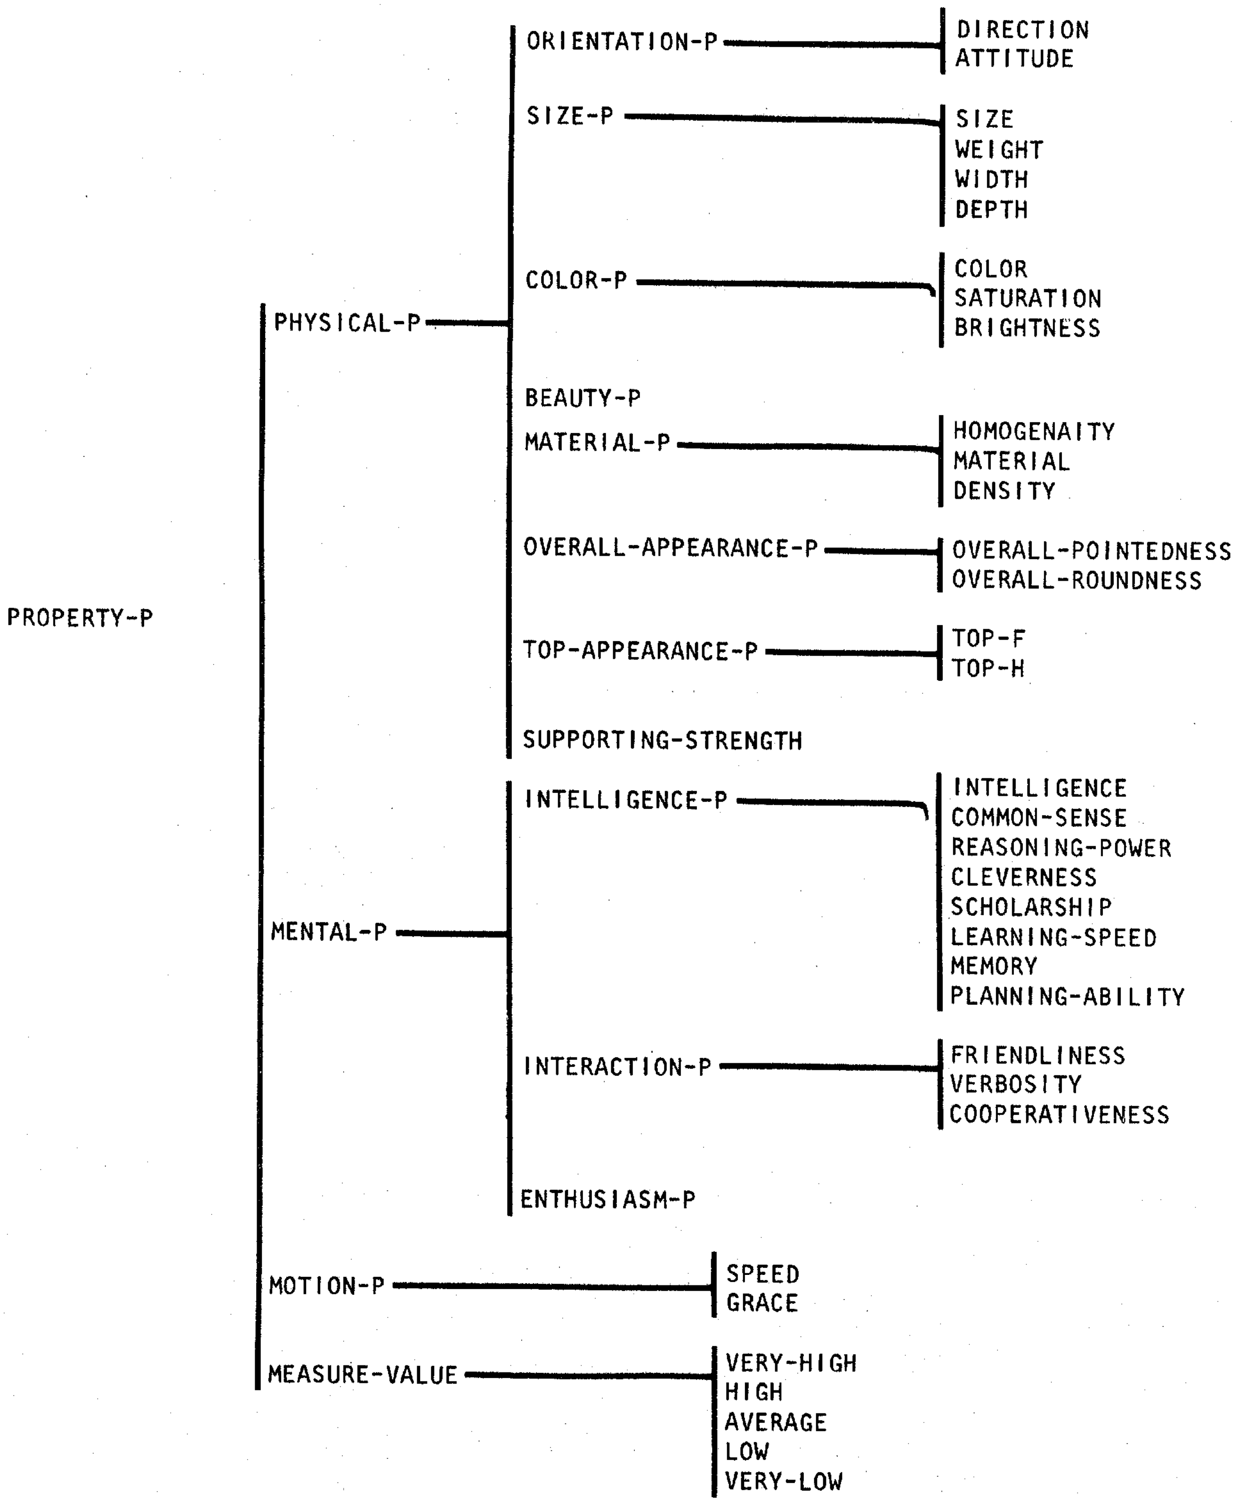

Figure 5. The hierarchical organization of the properties used in the examples.

```
U2             AKO          CYLINDER

Y1             AKO          CYLINDER
               COLOR        BLUE
               SIZE         HIGH

Y2             AKO          CYLINDER
               COLOR        GREEN

CYLINDER       AKO          THING
               INSTANCE     Y1
                            Y2
                            U2
               TYPICAL-INSTANCE
                            TI-CYLINDER
```

The typical instance is a frame created to describe how the instances are alike. Following earlier work, typical instances are computed as follows:

- If a slot-value combination appears in more than some fraction of the instances, T_{sv}, put that combination in the typical instance.
- If a slot appears in more than some fraction of the instances, T_s, but is not filled uniformly enough to pass the first test, put it in the typical instance without a value.

At the moment, both thresholds are set at 65%. Hence for the given cylinders, the typical instance is very simple:

```
TI-CYLINDER    COLOR
```

Thus the typical thing in the cylinder class has some color. The first transfer frame evaluation method exploits this information to pick out the transfer frame with the COLOR slot since the typical instance indicates that color is a commonly filled slot, one that is therefore wanted in some sense by the destination.

As of now, we therefore have the following frames:

```
B1             AKO          BRICK
               SIZE         MEDIUM
               COLOR        BLUE
               MATERIAL     WOOD

U2             AKO          CYLINDER
               COLOR        BLUE          TRANSFERRED-FROM        B1
```

Note that the COLOR slot of U2 has the BLUE value augmented by a comment specifying where the value came from. This exercises more of the Goldstein-Roberts frame system.

Now suppose the following is given again:

U2 is like B1.

Only the slots SIZE and MATERIAL emerge because COLOR is already filled. These form two frames, neither of which is better than the other with respect to the typical instance. Consequently another, weaker, method is used. This other method notices that some sibling of U2 has a SIZE slot, namely Y1. On the other hand no sibling has a MATERIAL slot. Hence the evidence favors using SIZE since it is more likely to apply than MATERIAL. Evidently U2 is medium in size.

Next, to expose still another evaluation method, let us consider the following:

U3 is like C1.

Assume that nothing more is known about U3 and that C1 is described as follows:

C1	AKO	CUBE
	SIZE	MEDIUM
	COLOR	BLUE
	MATERIAL	PLASTIC

As in the example using B1, three frames are created, one for SIZE, COLOR, and MATERIAL. Now, however, there are no known relatives of U3 yet, so none of the previous evaluation methods work. The decision, given that the sequence is connected, goes to the frame that is most in keeping with the context determined by the last transfer. The last transfer involved size, so this one will too. Actually the context is always reset to be the node in the property tree just above the last slot used. Consequently the context established is SIZE-P, as shown in figure 5, and anything from the group SIZE, HEIGHT, WIDTH, or WEIGHT passes.

This concludes the discussion of evaluators. Certainly the implementation is preliminary and a lot of changes may be found appropriate. For example, it is not clear that the dividing line between the transfer frame hypothesizers and the transfer frame evaluators is correctly placed. The same basic strategy may be useful either way.

Near Misses May Generate Transfers As Well As Examples

Previous work stressed the idea of near misses, samples which are not like the thing being described in some important way [Winston]. Using a near miss in a teaching sequence usually resulted in the placement of some so-called emphatic relationship such as MUST-BE-AKO BRICK or MUST-NOT-BE RED. The programs being described do not deal with this important teaching idea only because the thrust is in the direction of dealing with new ideas, not because the old ones have been superseded. Indeed it is fairly clear how near-miss action could be incorperated into the current system:

- Use the same hypothesis methods without change.

- Use the same evaluation methods, except that slots are not to be rejected merely because they happen to be filled in the destination frame.

- Revise the way the transfer frame is used to carry slot information from the source to the destination. Instead of adding to the VALUE facet of the slot, add to the MUST-BE facet or the MUST-NOT-BE facet instead. Or, since the Goldstein-Roberts frame system does not have these facets, it might be easier to get their effect by placing little programs into the existing REQUIRE facet.

With this it would be possible to give the following:

ARCH is *not* like TABLE.

The expected result would be the placement of EAT and WRITE in the MUST-NOT-BE facet of the PURPOSE slot of ARCH. This would happen even if ARCH already had something in the VALUE facet of the PURPOSE slot.

Summary

The hypothesis and evaluation methods are as follows:

- Use a remembered transfer frame. Examples of this will be given later.

- Make a transfer frame using slots with extreme values.

- Make one using slots that are known to be important in general.

- Look for slots that are unique with respect to the source's siblings.

- Look for slots that have unique values in them with respect to the siblings.

- Use all of the source's slots.

Thus hypothesis methods concentrate on looking at the source and its context.

- Weed out the slots that are already filled in the destination.

- Group the slots using the property hierarchy tree.

- Prefer transfer frames that have slots that are present in the typical instance associated with the destination.

- Prefer those that have slots that some sibling of the destination exhibits.

- Prefer those that are relevant to the context established by the last transfer.

Thus evaluation methods concentrate on looking at the destination and its context.

JUSTIFICATION AND CURIOSITY

Once a transfer frame is in hand, then it is necessary to decide if using it really makes sense given all that is known about the slots it will effect. If it is used, then there is the further question of whether the new knowledge gained about the destination should trigger the student into further, internally generated speculation.

There Are Several Ways To Justify The Transfer Frames

Once transfer frames have been found and ordered by hypothesizers and evaluators, the next job is to decide if they are indeed legitimate. There are at least four ways to do this. They will be described in the order of increasing length of explanation.

First, of course, the student can ask the teacher directly if the frame is appropriate. Second, the student can use the restriction feature of the frame system to prevent the insertion of values that conflict with restriction knowledge that exists in the AKO hierarchy. This comes automatically with the Goldstein-Roberts frame language.

Third, there is a method involving inspection of the AKO/INSTANCE path between the source and the destination. The basic idea is that if the frames on the source-destination path do not show a kind of distance with respect to the slots and values to be transferred, then the transfer is judged to be weak. For a simple example, suppose that the following is given:

U4 is a BRICK.
U4 is like B1.

The SIZE, MATERIAL, and COLOR slots seem appropriate for transfer given the hypothesizing and evaluation methods described. Looking more closely, however, SIZE seems safest because all bricks are of medium size. The MATERIAL slot comes next because all but one are made of wood, as is B1, whereas there are many values for COLOR. The BRICK node therefore has greatest *admittance* with respect to the SIZE-MEDIUM combination. Here is a crude formula for quantifying this notion:

<admittance of node N>

$$= (a - b + 1) / (a + 1) \qquad \text{for } a > 0$$
$$= 1 \qquad \text{for } a = 0$$

where a is the number of times a slot is filled in N and the children of N and b is the number of times the slot is filled with a value different from the one in the source frame.

Note that if a slot is always filled the same way as it is in the source, b will be 0 and the admittance will be 1. On the other hand, if none of the values that appear are the same, $a = b$ and the admittance is $1/b$, and this number can get small. For the bricks in this example, the known brick colors are blue, green, red, blue, and undefined, giving an admittance with respect to COLOR-BLUE of $(4 - 2 + 1) / (4 + 1) = .6$. All but one of the bricks are made of wood, however, so the admittance with respect to MATERIAL is $(5 - 1 + 1) / (5 + 1) = .83$. SIZE wins because its admittance is 1. It is safer to transfer the SIZE-MEDIUM combination from B1 than it is to transfer MATERIAL-WOOD or COLOR-RED.

This situation for which the transfer is among siblings is particularly simple because the admittance is just a function of the parent and all the siblings. For the following, more must be done:

U4 is like L3.

Again SIZE, MATERIAL, and COLOR are the candidate slots. To compare them, at the moment, it seems sensible to calculate the admittance for all nodes intervening between U4 and L3, to multiply the results together, and to let that be the so-called path admittance. For the example, there are four nodes to deal with as shown in figure 4, namely BRICK, RECTANGULAR-BLOCK, BLOCK, and BALL. But two of them have neither SIZE, MATERIAL, nor COLOR slots and contribute nothing. Now COLOR-BLUE wins, for its *path admittance* is $.75 \times .6 = .45$ while that for MATERIAL-RUBBER is $.75 \times .16 = .125$ and that for SIZE-SMALL is $.5 \times .166 = .083$.

In point of fact, it makes some sense to use not just the slots in the proposed transfer frame, but also the siblings of those slots, on the ground that similar properties tend to be coherent or dispersive together. Thus the admittance for a transfer frame with only a BRIGHTNESS slot causes a admittance measurement with respect to COLOR, SATURATION, and BRIGHTNESS.

Now we turn to a fourth method for judging the quality of a proposed transfer, one that requires the student to take notes on why transfers seem to work and to create justification frames that can be matched against a proposed destination to see if the destination exhibits apparently essential slot values.

Suppose, for example, that the student knows C1 has VERY-HIGH in the slots TOP-F and TOP-H. Further suppose that the teacher gives this:

C1 is like TABLE.

Since the teacher presses home the similarity, certainly the intent must be that it is possible to eat from or write on C1, just as it is with a table, since there are now the following frames:

TABLE	AKO	STRUCTURE
	PURPOSE	EAT
		WRITE
	SIZE	MEDIUM
	TOP-F	VERY-HIGH
	TOP-H	VERY-HIGH
	HAS-PART	B5
		Y2

```
C1            AKO                CUBE
              SIZE               MEDIUM
              COLOR              BLUE
              MATERIAL           PLASTIC
              TOP-F              VERY-HIGH
              TOP-H              VERY-HIGH
```

PURPOSE is the only possible slot for transfer, but it would have been selected by several hypothesis methods anyway. After the transfer, the student, on the request of the teacher, looks to see how the source and the destination resemble one another, remembers the transfer frame, and constructs a justification frame that reflects the similarity:

```
TABLE         AKO                STRUCTURE
              .
              .

TF-1          AKO                TRANSFER-FRAME
              TRANSFER-SLOTS     PURPOSE
              TRANSFERRED-FROM
                                 TABLE
              TRANSFERRED-TO
                                 TABLE
              JUSTIFICATION-FRAME
                                 JF-1

JF-1          AKO                JUSTIFICATION-FRAME
              SIZE               MEDIUM
              TOP-F              VERY-HIGH
              TOP-H              VERY-HIGH
```

Now TABLE has become a standard source of particular values for the PURPOSE slot, namely EAT and WRITE, through the skillful selection of circumstance by the teacher.

Henceforward, a new hypothesizer will be the first to work. It will look for values in the TRANSFER-FRAME slot. In this example, it finds one for TABLE, namely TF-1.

More importantly perhaps, the student now has a justification frame attached to this standard transfer frame. This justification frame must be a subframe of a proposed new destination if the new destination is to pass. In this example, for PURPOSE to be transferred to a destination from TABLE, the destination must have the SIZE, TOP-F, and TOP-H slot values dictated by the justification frame, JF-1.

Consider this:

C2 is like TABLE.

There will be a justified PURPOSE transfer if the student's C2 frame has the three key justification frame slots properly filled or if the student can get proper values from the teacher or from his own sensory apparatus.

In general, this is really only a mechanism for getting a first idea of why a given transfer is justified. Further refinement of the justification frame is possible by direct telling or by fresh transfers to it as a destination.

While all this student note taking is going on, information is also added to the TOP-F and TOP-H and SIZE frames:

TOP-F	AKO	TOP-APPEARANCE-P	
	TRIGGER-VALUE	VERY-HIGH	TRANSFER-SOURCE TABLE
TOP-H	AKO	TOP-APPEARANCE-P	
	TRIGGER-VALUE	VERY-HIGH	TRANSFER-SOURCE TABLE
SIZE	AKO	SIZE-P	
	TRIGGER-VALUE	MEDIUM TRANSFER-SOURCE TABLE	

Of course the TRIGGER-VALUE slot for SIZE will become gorged far sooner than for TOP-F and TOP-H since SIZE in a more common property. This means that SIZE will not be as useful as the other two with respect to the use of trigger values about to be described.

Filling A Slot May Induce Curiosity

It is reasonable for the student, having just learned something, to make conjectures based on the new knowledge. Often these conjectures will be wrong since they are generated interally using rather flimsy heuristic evidence. Hence it will be more important than usual to use the various justification methods to confirm the conjectures.

To see how other conjecture methods work, suppose that U5 has the following description:

U5	AKO	BRICK
	SIZE	SMALL

The first conjecture method uses information placed when justification frames are made. Suppose that the following is given:

U5 has VERY-HIGH in the TOP-F slot.

From this, and permission of the teacher to think a bit, it is reasonable for the student to examine the TOP-F frame for clues about other properties of U5. The TRIGGER-VALUE slot of TOP-F contains the value VERY-HIGH along with a comment to the effect that the value was placed while constructing a justification frame involving a transfer from TABLE. Since VERY-HIGH in the TOP-F slot evidently helped justify a transfer from TABLE in the past, it is reasonable for the student to try a transfer from TABLE to U5 again. Thankfully the trigger value information only exists if a justification frame also exists. The student therefore has a justification frame that he can use to decide if the transfer makes sense, possibly asking the teacher some questions along the way about the slots that the justification frame specifies.

A second conjecture method uses siblings. Suppose that the following is given:

U5 has BLUE in the COLOR slot.

Using this, the student may want to look for siblings that are also blue with the hope that U5 and such a sibling may be alike in other ways. Indeed, this happens. Siblings with BLUE in the COLOR slot are collected and the most typical one becomes a conjectured source.

The most typical blue sibling is determined using a frame similarity computation defined as follows:

<frame similarity between X and Y>
= a/b for $b > 0$
= 0 for $b = 0$

where a is the number of slot-value combinations that appear in both X and Y and b is the number of slot-value combinations that appear in either.

If all of the slots in X and Y have different values, then the frame similarity will be zero. If all of the slots in X and Y have the same values, then it will be one.

In this example, both B1 and B4 are blue. B1 is judged the more typical of the two because the frame similarity between B1 and the typical instance frame for BRICK, TI-BRICK, is .66, whereas the frame similarity between B4 and TI-BRICK is only .5. The difference is the result of a PURPOSE-SUPPORT slot-value combination present in B4 but missing in B1 and TI-BRICK.

Thus, learning that U5 is blue may result in a transfer from B1 which would assert that U5 is made of wood. Again, whatever justification methods are available should be used. Moreover, it would be more sensible to transfer only those properties from B1 that are closely related to the COLOR slot since the conjecturing method is so tenuous. This, however, has not been implemented in the existing system.

Making these conjectures is one kind of "curiosity." Another can come from obvious reaction to learning that an unknown is a kind of something. Consider this:

U6 is a BRICK.

Without further fuss, it would make sense for the student to assume that U6 has all the typical slot-value combinations that are in the typical instance frame, assuming that U6 is of medium size and is made of wood as a result. But the student should also know that typical instances may have unfilled slots that get there when a slot is common but does not appear with the same value often enough for a value to acommpany it to the typical member. At the moment, the teacher is asked to supply values for these slots either explicitly or by reference to a source with the proper value. For the example given, then, the following is printed:

> I assume U6 has the slot-value pairs
> SIZE-MEDIUM and
> MATERIAL-WOOD.
>
> I am curious about a value for COLOR.

If the teacher supplies no help at all, it is conceivable that the student still may successfully fill the slot by guessing a suitable source. Two methods come to mind that parallel the methods just described for responding to a new slot instantiation. Unlike the other methods, neither have been implemented.

The first source to consider is found by again appealing to justification-frame information. The justification-frame construction program already places JUSTIFIED-VALUE slot information just as it now places TRIGGER-VALUE INFORMATION. For the table example, we have this:

```
PURPOSE          JUSTIFIED-VALUE EAT      TRANSFERRED-FROM         TABLE
```

Wondering about the value for the PURPOSE slot of something, the student can use TABLE as a possible transfer source to be attempted. If there are many justified values in the slot's frame, then the student might well want to screen them by looking for the known purposes of the siblings of the destination frame. There is a greater chance that the destination will have a purpose similar to one of its siblings than to something entirely removed.

The second source to consider for transfer is the most typical sibling of the destination frame that has the slot in question filled. This would not be done if the parent's admittance with respect to the slot is low.

Generalizing, the student could move up the AKO tree, looking for a suitable sibling of the more remote ancestors, not just the parent, until one is found that a lot is known about. There will be examples of this when we discuss animals. There, this will amount to trying a transfer from ROBOT, the thing most is known about, through intervening nodes to the destination. This will feel good only if the similarity of ROBOT and the destination is high with respect to the desired property. Similarity can be measured by the existing path-admittance justification method.

The Blocks World May Be Deceptively Small

The small number of properties associated with each object may be a cause for some uneasiness. Is it possible that the examples work only because of the careful arrangement of the slots and their small number? Maybe. Indeed one important question to be addressed is that of how much complexity can be coped with before the system breaks down. Meanwhile, two points are probably worth observing:

- The fact that things are immersed in an AKO tree will tend to keep the clutter down. Presumably most property values are obtained by defaulting to higher and higher level concepts.

- Good teaching normally requires using examples with relatively few prominent properties. Good examples are the ones for which the computation required for deciphering is low.

Indeed the reason the simple physical world is a good source of general metaphors, some of which reach the social world, the mental world, and various expert problem solving worlds, may be because its simplicity makes the metaphor understanding problem easier.

EXAMPLES FROM THE ANIMAL WORLD

Animal world is shown in figure 6 and figure 7. We will use it first to review some basic hypothesis and justification ideas, then we will turn fleetingly to an example involving analogy, and finally, we will look at the use of standard transfer sources such as people or robots.

Jack And Jill Can Be Described By Animal Metaphors

Let's look at a sample sequence:

> Jack is like a fox.

Since fox has a very high value for cleverness, it is concluded that Jack does too. The context becomes intelligence and the use of the fox as a metaphor for cleverness will be noted.

> Jill is like an elephant.

Since an elephant has several slots, there are several possibilities, namely memory, weight, and grace. Good memory is the winner though, since the context is intelligence.

> Jill is also like a cheetah.

Evidently Jill is fast. The context now has to do with motion properties.

> Robbie is like an elephant.

The context now singles out grace and transfers a low value because the context now has to do with motion, not weight or memory.

> Robbie is a robot.
> Robbie is like an elephant.

Robbie already has a grace property. The transfer must have something to do with either weight or memory. Knowing that Robbie is a robot helps because the other robots have values in the memory slot but not in the weight slot. Evidently Robbie has

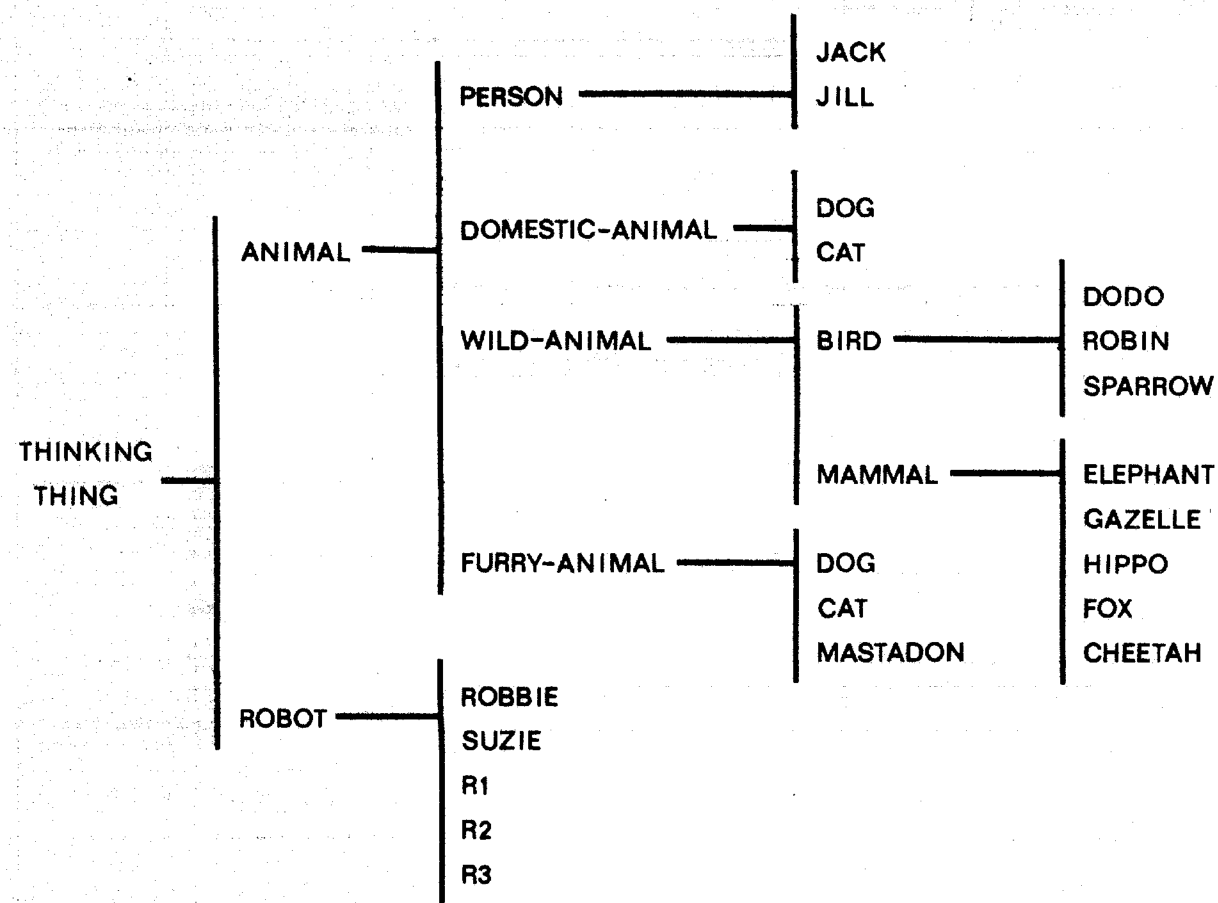

Figure 6. The hierarchical organization of a simple animal world.

```
(DEFRAME ANIMAL
         (AKO ($VAL (THINKING-THING)))
         (INSTANCE ($VAL (PERSON)
                         (DOMESTIC-ANIMAL)
                         (WILD-ANIMAL)
                         (FURRY-ANIMAL)
                         (EXTINCT-ANIMAL))))

(DEFRAME BIRD
         (AKO ($VAL (WILD-ANIMAL)))
         (INSTANCE ($VAL (DODO) (ROBIN) (SPARROW)))
         (INTELLIGENCE ($VAL (VERY-LOW)))
         (SIZE ($VAL (VERY-LOW))))

(DEFRAME CAT (AKO ($VAL (DOMESTIC-ANIMAL) (FURRY-ANIMAL))) (SIZE ($VAL (LOW))))

(DEFRAME CHEETAH (AKO ($VAL (MAMMAL))) (SPEED ($VAL (VERY-HIGH))))

(DEFRAME DODO (AKO ($VAL (BIRD))) (EXTINCT ($VAL (YES))) (INTELLIGENCE ($VAL (LOW))))

(DEFRAME DOG (AKO ($VAL (DOMESTIC-ANIMAL) (FURRY-ANIMAL))))

(DEFRAME DOMESTIC-ANIMAL
         (AKO ($VAL (ANIMAL)))
         (INSTANCE ($VAL (DOG) (CAT)))
         (INTELLIGENCE ($VAL (HIGH)))
         (SIZE ($VAL (LOW))))

(DEFRAME ELEPHANT
         (AKO ($VAL (MAMMAL)))
         (MEMORY ($VAL (VERY-HIGH)))
         (GRACE ($VAL (VERY-LOW)))
         (WEIGHT ($VAL (VERY-HIGH)))
         (SIZE ($VAL (VERY-LARGE))))

(DEFRAME FOX
         (AKO ($VAL (MAMMAL) (FURRY-ANIMAL)))
         (CLEVERNESS ($VAL (VERY-HIGH)))
         (INTELLIGENCE ($VAL (HIGH)))
         (SIZE ($VAL (LOW))))

(DEFRAME FURRY-ANIMAL
         (AKO ($VAL (ANIMAL)))
         (INSTANCE ($VAL (DOG) (CAT) (FOX) (MASTADON)))
         (INTELLIGENCE ($VAL (AVERAGE))))

(DEFRAME GAZELLE (AKO ($VAL (MAMMAL))) (GRACEFULNESS ($VAL (HIGH))) (SIZE ($VAL (AVERAGE))))

(DEFRAME HIPPO (AKO ($VAL (MAMMAL))) (WEIGHT ($VAL (HIGH))) (SIZE ($VAL (VERY-LARGE))))
```

Figure 7. The frames that define animal world.

```
(

(DEFRAME JACK (AKO ($VAL (PERSON))))

(DEFRAME JILL (AKO ($VAL (PERSON))))

(DEFRAME MAMMAL
         (AKO ($VAL (WILD-ANIMAL)))
         (INSTANCE ($VAL (ELEPHANT) (CHEETAH) (FOX) (GAZELLE) (HIPPO)))
         (INTELLIGENCE ($VAL (VERY-HIGH)))
         (SIZE ($VAL (AVERAGE))))

(DEFRAME MASTADON (AKO ($VAL (FURRY-ANIMAL))))

(DEFRAME PERSON
         (AKO ($VAL (ANIMAL)))
         (INSTANCE ($VAL (JACK) (JILL)))
         (INTELLIGENCE ($VAL (VERY-HIGH)))
         (SIZE ($VAL (AVERAGE))))

(DEFRAME R1 (AKO ($VAL (ROBOT))) (MEMORY ($VAL (AVERAGE))))

(DEFRAME R2 (AKO ($VAL (ROBOT))))

(DEFRAME R3 (AKO ($VAL (ROBOT))))

(DEFRAME ROBBIE (AKO ($VAL (ROBOT))))

(DEFRAME ROBIN (AKO ($VAL (BIRD))) (SIZE ($VAL (VERY-LOW))))

(DEFRAME ROBOT
         (AKO ($VAL (THINKING-THING)))
         (INSTANCE ($VAL (ROBBIE) (SUZZIE) (R1) (R2) (R3)))
         (INTELLIGENCE ($VAL (AVERAGE)))
         (MEMORY ($VAL (HIGH)))
         (COMMON-SENSE ($VAL (AVERAGE)))
         (REASONING-POWER ($VAL (VERY-LOW)))
         (VERBOSITY ($VAL (LOW)))
         (SIZE ($VAL (AVERAGE))))

(DEFRAME SHEEP (INTELLIGENCE ($VAL (LOW))) (SIZE ($VAL (AVERAGE))))
```

Figure 7 (continued). The frames that define animal world.

```
(DEFRAME SPARROW (AKO ($VAL (BIRD))) (SIZE ($VAL (VERY-LOW))))




(DEFRAME SUZZIE (AKO ($VAL (ROBOT))))


(DEFRAME THING
         (AKO ($IF-ADDED (ADD-INSTANCE)) ($IF-REMOVED (REMOVE-INSTANCE)))
         (INSTANCE ($TO-ADD ((INSTANTIATE-A-FRAME)))
                   ($IF-ADDED (ADD-AKO))
                   ($IF-REMOVED (REMOVE-AKO))
                   ($VAL (THINKING-THING) (OBJECT) (PROPERTY)))
         (SELF ($DISCUSS ((DESCRIBE-FRAME)))
               ($ORDER ((SLOTS-TO-BE-INSTANTIATED)))
               ($TYPE)
               ($SLOTS (@(SETMINUS1 (HERITAGE-SLOTS :FRAME) (FSLOTS 'THING))))))


(DEFRAME THINKING-THING (AKO ($VAL (THING))) (INSTANCE ($VAL (ROBOT) (ANIMAL))))





(DEFRAME WILD-ANIMAL
         (AKO ($VAL (ANIMAL)))
         (INSTANCE ($VAL (BIRD) (MAMMAL)))
         (INTELLIGENCE ($VAL (AVERAGE)))
         (SIZE ($VAL (AVERAGE))))
```

Figure 7 (continued). The frames that define animal world.

a good memory.

Robbie is like an elephant.

The third time around, only weight is left. The context becomes size.

Now for the next example, suppose the frame for robot has the following information:

ROBOT	AKO	THING
	INTELLIGENCE	MEDIUM
	MEMORY	HIGH
	COMMON-SENSE	MEDIUM
	REASONING-POWER	LOW
	VERBOSITY	LOW

These properties make two groups: one deals with intelligence, memory, common sense, and reasoning power, all aspects of the general notion of intelligence, and the other deals with verbosity, a dimension of personality. If X is unknown, then two transfer frames will be proposed in response to the following item:

X is like a robot.

Suppose the transfers are allowed and the transfer frames are recorded. Then consider the following sequence:

Y has medium common sense.
Y is like a robot.

What properties of robots are preferred for the next transfer? Intelligence, common sense and reasoning power could be relevant or verbosity might be right. But since X's memory is already known to be good, the choice is to pass values through intelligence, common sense, and reasoning power since these qualities have been transferred earlier as a group from the robot concept along with memory which already has a value in Y.

Transferring intelligence, common sense, and reasoning power properties is the preferred action because having one fact about intelligence makes acquiring more a likely possibility. So far Y has no personality properties and it would be more risky to transfer through the verbosity slot.

A Transformation May Be Specified Directly Or By Analogy

Of course a value need not slither through a transfer frame unscathed. Generally, it may be subjected to some sort of value transformation. VERY-HIGH becomes VERY-LOW if T-OPPOSITE is the transformation in effect. MEDIUM becomes HIGH if T-MORE is the transformation. An APPLE becomes FRUIT by way of T-GENERALIZE. Other, fancier things may be useful in making metaphors between worlds.

The name of the transformation may be directly specified, of course, as in the following fragment:

John is the opposite of a fox.

However, the transformation may be given by an analogy:

Jane resembles a fox in the same way John does.

After CLEVERNESS is found to be the dimension in which Jane is like a FOX, it is a simple matter to use the corresponding transfer frame to test John against FOX to find that T-OPPOSITE is the implied transformation.

Testing the transfer frame using the analogy source and the analogy destination also can help filter out wrongly conjectured transfer frames that may have survived all other filtering operations. It better be that the same transformation applies to all of the slots in the transfer frame when it is used to compare the analogy source and analogy destination frames. Otherwise, chuck it out.

Notice, incidentally, that the source, the destination, the analogy source, and the analogy destination may all be different. Notice also that these four items, together with the transfer frame and the transformation, all may or may not appear, giving a total of 63 combinatorial possibilities, the bulk of which are probably absurd.

Path Admittance Helps Decide If Slots Can Be Filled From A Standard Source

Suppose questions are asked about the size and intelligence of a cheetah in the context of the thinking-thing information shown before in the animal AKO tree and the frame listing.

As mentioned at the close of the justification and curiosity section, one way to fill the specified slots is to try some well-known thing as a source, trusting to the path-admittance method to warn against bad transfers.

Given the current information about animals, the path admisability has been calculated for certain transfers as follows:

ROBOT-TO-CHEETAH	INTELLIGENCE-MEDIUM	.07
ROBOT-TO-DODO	INTELLIGENCE-MEDIUM	.06
ROBOT-TO-CHEETAH	SIZE-MEDIUM	.28
ROBOT-TO-DODO	SIZE-MEDIUM	.14

It is a bit hard to follow this directly because the program that calculates the numbers is looking at a lot of nodes and it is looking for the siblings of INTELLIGENCE and SIZE as well as INTELLIGENCE and SIZE themselves. But evidently transfers to cheetah are more reliable than transfers to dodo and transfers of size are much more reliable than transfers of intelligence. This makes sense, both in terms of the information the student has and in terms of what we would expect a priori.

Now, for further illumination, we can follow the history of the path admittance of the robot-to-cheetah transfers as other information is added to the students knowledge.

For ROBOT-TO-CHEETAH, INTELLIGENCE-MEDIUM

A cheetah is a furry animal.	.08	.40
A sheep is a furry animal.	.40	.30
A dog is very intelligent.	.30	.24

For ROBOT-TO-CHEETAH, SIZE-MEDIUM

A cheetah is a furry animal.	.28	.25
A sheep is a furry animal.	.25	.38
A dog is small.	.38	.30

Keeping the known facts in mind, all of this makes some sense. Discovering that a cheetah has a new connection into the AKO tree means the shortest path to robot is different. In one case, this drives the path admittance down slightly, in the other, up considerably. The large increase in the INTELLIGENCE-MEDIUM number reflects both the shortening of path length and the reduction in the number of siblings known to be of the wrong intelligence.

Learning that a sheep is also a furry animal makes the size transfer look better because it gives the cheetah a new sibling with the same size as the robot, but it makes the intelligence transfer look worse because a sheep is also known to be of low intelligence.

Finally, getting extra information about a dog, already known to be a furry animal, makes both transfers path admittance go down because both the dog's size and intelligence differ from the robot's.

KEY ISSUES

There has been fully too little experiment with the programs and the ideas in them to know how much can be accomplished. Ideas have been illustrated, but certainly none have been solidly demonstrated. Many more experiments and much larger, more completely specified domains are necessary to do that. Still, there is some preliminary hope that the following principles may hold:

The principle of representational parsimony. If all sorts of knowledge is represented uniformly, then it is all subject to the same learning processes. Since objects, properties, and even justifications have the same representation, all can be learned about through transfer frames. With respect to domain, any in which the objects and properties can be described in terms of frames is potentially a domain that learning using transfer frames can address.

The principle of expanding competence. The more that is known, the better learning should be, both in terms of speed and accuracy. Certainly speed and accuracy should increase with increasing knowledge when learning is by transfer frames since the more the student knows, the easier it is for the teacher to find lucid examples less subject to misinterpretation.

SPECULATIONS

Metaphor Traces Could Be Used To Find Substitutes, Notice Attributes, And Pass The Time

Look again at the example of the table transfer to the cube. Having noticed that the transfer took place while both frames were observed to be of medium size and to have flat tops, a lot of information was recorded that might be used as follows:

- Having made the transfer to C1, idle time could be spent seeing if relatives of C1 are also like TABLE when compared through the JUSTIFICATION-FRAME, JF-1. If so, proceed to learn more by making the transfer through the transfer frame, TF-1.

- To find something which would serve the same purpose as a table, note that the table's purpose was transferred earlier by using the table's recorded transfer frame and justification frame. See if anything in the physical vicinity satisfies the justification frame. Index into the frames in the vicinity, perhaps using the justification frame's slots.

- To find something whose purpose is to serve in writing, look into the frame for PURPOSE and note that TABLE has been a source of metaphors for writing. Get the TRANSFER-FRAME and JUSTIFICATION-FRAME information from TABLE.

Past Transfers Could Be Used In Generating Descriptions

The simplest ways to generate transfers is to bounce back information previously digested. The system already leaves certain information behind to enable this.

First, when transfers are used to transfer information into a concept later to be described, the source is recorded. If Sam was said to be like a fox, it would be easy to say this:

> Sam is very intelligent, like a fox.

Just having this would make conversation dull, full of triteness, but other devices could be used:

- When a frame has a slot filled with a VERY-HIGH or VERY-LOW value, the fact could be recorded in the slots frame by instantiation of the VERY-HIGH-VALUES or VERY-LOW-VALUES slots using the name of the frame where the exceptional value was recorded. This frame is then a possible transfer source.

- If a frame with an extreme value was used as a source before, it should be particularly good. It is better if it has not served as a transfer source contributing to the description of the thing to be described.

- If, in looking over what is to be said, there are many possible sources, the transfer generator can run its various source possibilities through the filters using its best guess about what the listener already knows about the concept being described. Clearly the best descriptions are the ones that allow rapid filtering down to the correct transfer frames. This means the sources specified will automatically tend to tell the listener facts it is interested in knowing about and stick to a context, among other things.

- The transfer generator may decide it is folly to do the whole description as a chain of metaphors. Instead it may be better to explicitly specify a slot or a context from time to time.

- The transfer generator can bias itself by choosing sources from either pleasant categories (fields and flowers) or unpleasant ones (fire and brimstone).

- As an additional literary device, a pointer into the AKO tree should be maintained and transfer sources should be selected from the descendants of it. This would tend to help avoid inelegant mixing of metaphors.

SUMMARY

The path has been involved. Therefore it makes sense to put the key ingredients on display now, by way of summary:

Frame-like representation. A representation is a vocabulary of symbols and a set of conventions for arranging them to describe things. Obviously representation is a central issue in attempts to understand learning, for nothing can be learned unless there is a representation that can capture the new knowledge to be learned. Consequently, when a new and powerful representation is found, it is useful to examine it with a view toward addressing the competences exhibited by learners.

Frame terminology is used here because it is rich and because the program that has been described makes use of a small number of features of FRL, the frame representation language created by Goldstein and Roberts [Roberts and Goldstein].

A simple property list is a representation in which things are described by properties that can assume values. The *frame representation*, invented by Minsky and developed by Goldstein and Roberts, is a newer representation in which the notion of property list is generalized [Minsky]. Instead of properties, there are slots. One facet of a slot is its value, but unlike properties, slots can have many facets, not just a value facet. Among these facets are places where demon-like procedures reside, waiting for insertions, deletions, or accesses that trigger them into taking a piece of the action.

For our purpose, the value facet of a slot was the most important one, and we refered to the value without being more precise. Occasionally it is useful to know that other facets exist, the restriction facet in particular. The restriction facet dictates constraints on what is allowed in the value facet. Finally, the values of some slots have comments associated with them that, for example, can give information about where the value came from.

It is also useful to know that frames are inherently arranged in hierarchies so that access to one frame can cause access to a second frame from which the first inherits information.

The *destination* frame is the thing to be learned about. It has *slots* that may assume *values*. Often values are supplied by a *source* frame that happens to have a slot value suited to the destination. Thus if it is said that Robbie is sly as a fox, then *fox* is the source, *Robbie* is the destination, *cleverness* is the slot, and *very-high* is the value.

Hypothesizing and evaluating transfer frames. Typically there are many possible ways the source may be like the destination. A combination of the known properties of the source and destination must be used, perhaps together with context, to make the correct judgement about what is to be learned. This is done using a *transfer frame*, a frame that stands between the source and the destination like a template and determines the information that is transferred. A key idea is that these transfer frames can be generated dynamically by the student using a variety of common-sense methods that access what is already known. The good teacher, knowing how these methods work and having a rough model of what is already known by the student, can teach in a way that improves both the transfer rate and accuracy.

Grouping and the typical instance. Groups are important. Groups of things tend to have the same properties, not just a single, group-defining property. Consequently, an abstract *typical instance* can be constructed for a group. The typical instance consists of a number of slots and slot values that capture the essence of the group it describes. The notion of typical instance derives from earlier work [Winston] [Davis]. The importance of typical instances in learning seems clearer now because they are the key to several ideas for hypothesizing and evaluating transfer frames. The typical instance descriptions of groups that the source belong to help hypothesize transfer frames and those for the destination help evaluate those transfer frames.

Similarly, groups of properties are important. The group dealing with size, for example, encompasses weight, height, width, depth, and general size. If one is mentioned, it helps to establish a context in which the others are expected.

Dissimilarity measurement using network path admittance. All things reside in a network of a-kind-of, or *AKO* relations. For every AKO there is a complementary *INSTANCE*. Paths between the source and destination in the AKO-INSTANCE network offer some help in deciding if a given transfer makes sense: if nodes along a path show a strong tendency for slots like those in a proposed transfer to be filled with the the right values, then the network has high *path admittance* and there is some support for the transfer.

Interesting experiments with children suggested this idea [Carey-Block]. If a small child is asked if an animal has a certain organ, it often responds with a certainty proportional to the apparent overall similarity between the given animal and the most common thing that the child knows has the organ, a human typically. A child will be less sure that a bird has a liver than that a monkey has one.

Creating and using justification frames. Keeping track of specific properties that legitimize filling slots in various ways is another way to judge a proposed transfer. These specific slot-value combinations are stored in *justification frames*. These justification frames can be accumulated by experience. They can also be acquired and honed by dialogue with the teacher, just as other frames can be.

REFERENCES

- David Marr, "Cooperative Computation of Stereo Disparity," *Science*,, Vol. 194, October, 1976. Cited here for its description of Marr's approach to the computational problems of psychology.

- Marvin Minsky, "A Framework for Representing Knowledge," in *The Psychology of Computer Vision*, edited by Patrick Henry Winston, McGraw-Hill Book Company, New York, 1975. The paper in which the vocabulary of frames was largely established.

- Bruce Roberts and Ira Goldstein, paper in preparation. Describes the frame system used by the programs described here.

- Patrick Henry Winston, "Learning Structural Descriptions from Examples," PhD thesis, in *The Psychology of Computer Vision*, edited by Patrick Henry Winston, McGraw-Hill Book Company, New York, 1975.

- Randall Davis, "Applications of Meta Level Knowledge to the Construction, Maintenance and Use of Large Knowledge Bases," PhD thesis, Stanford Artificial Intelligence Laboratory, AIM-283, 1976. Cited here for its interesting use of groups and typical members to help humans in the design of situation-action rules.

- Susan Carey-Block, paper in preparation. Will describe tendancy of children to compare animals with humans to decide if the animals have various characteristics.