

January 1983
Also numbered AIM-349

Report No. STAN-G-82-950

Learning Physical Description from Functional Definitions, Examples and Precedents

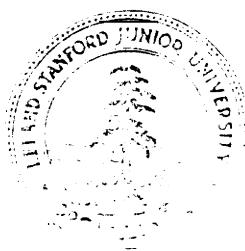
by

Patrick H. Winston, Thomas O. **Binford**, Boris Katz

and Michael Lowry

Department of Computer Science

Stanford University
Stanford, CA 94305



```

Rule
  RULE-Z
  if
    [OBJECT-10 HQ LIGHT]
    [CONCAVITY-8 PHYSICAL-PART-OF OBJECT-10]
    [BODY-9 PHYSICAL-PART-OF OBJECT-10-J]
    [BOTTOM-8 PHYSICAL-PART-OF OBJECT-10]
    [CONCAVITY-8 AKO CONCAVITY]
    [CONCAVITY-8 HQ UPWARD-POINTING]
    [BODY -9 AKO BODY]
    [BODY-9 HQ CYLINDRICAL]
    [BODY -9 HQ SMALL]
    [BOTTOM-8 AKO BOTTOM]
    [BOTTOM-8 HQ FLAT]
  then
    [OBJECT-10 AKO CUP]
  unless
    [ [OBJECT-10 AKO OPEN-VESSEL] HQ FALSE]
    [ [OBJECT-10 HQ LIFTABLE] HQ FALSE]
    [ [OBJECT-10 HQ STABLE] HQ FALSE]
    [ [BODY-9 HQ GRASPABLE] HQ FALSE]
  case
    DEFINITION-1 DESCRIPTION-2 DESCRIPTION-4 DESCRIPTION-1

Should I index it as a rule?
> y

```

Artificial Intelligence Laboratory

Stanford University

in collaboration with

Artificial Intelligence Laboratory
Massachusetts Institute of Technology

AIM Memo 349
STAN-CS-82-950

November 1982
Revised January 1983

Learning Physical Descriptions From Functional Definitions, Examples, and Precedents

by

Patrick H. Winston*

Thomas O. Binford

Boris Katz*

and Michael Lowry

*Patrick H. Winston and Boris Katz are at the Massachusetts Institute of Technology Artificial Intelligence lab.
This paper is also available from Massachusetts Institute of Technology as M.I.T. AI Memo 679

To use DEFINITION-1 I need to know if [OBJECT-3 **AKO** OPEN-VESSEL]
 I am trying to show [OBJECT-3 **AKO** OPEN-VESSEL]
 Supply, n, ?, r = rules, p = precedents, or a suggestion:
 > P
 I find:
 DESCRIPTION-Z < 3. Tinks >
 I note [CONCAVITY-3 PHYSICAL-PART-OF OBJECT-31 for use with **DESCRIPTION-2**
 I note [CONCAVITY-3 **AKO** CONCAVITY] for use with DESCRIPTION-2
 I note [CONCAVITY-3 HQ UPWARD-POINTING] for use with DESCRIPTION-2
 The evidence from DESCRIPTION-2 indicates [OBJECT-3 **AKO** OPEN-VESSEL]
 The evidence from DEFINITION-1 indicates [OBJECT-3 **AKO** CUP]
 Rule RULE-1 is derived from DEFINITION-1 DESCRIPTION-2 DESCRIPTION-3
 DESCRIPTION-1 and looks like this:

Rule
 RULE-1
 if
 [OBJECT-g HQ LIGHT]
 [CONCAVITY-7 PHYSICAL-PART-OF OBJECT-91]
 [HANDLE-4 PHYSICAL-PART-OF OBJECT-91]
 [BOTTOM-7 PHYSICAL-PART-OF OBJECT-91]
 [CONCAVITY-7 **AKO** CONCAVITY]
 [CONCAVITY-7 HQ UPWARD-POINTING]
 [HANDLE-4 **AKO** HANDLE]
 [BOTTOM-7 **AKO** BOTTOM]
 [BOTTOM-7 HQ FLAT]
 then
 [OBJECT-9 **AKO** CUP]
 unless
 [[OBJECT-g **AKO** OPEN-VESSEL] HQ FALSE]
 [[OBJECT-9 HQ LIFTABLE] HQ FALSE]
 [[OBJECT-9 HQ GRASPABLE] HQ FALSE]
 [[OBJECT-9 HQ STABLE] HQ FALSE]
 case
 DEFINITION-1 DESCRIPTION-2 DESCRIPTION-3 DESCRIPTION-1
 - Should I index it as a rule?
 > Y

Abstract

It is too hard to tell vision systems what things look like. It is easier to talk about purpose and what things are for. Consequently, we want vision systems to use functional descriptions to identify things, when necessary, and we want them to learn physical descriptions for themselves, when possible.

This paper describes a theory that explains how to make such systems work. The theory is a synthesis of two sets of ideas: ideas about learning from precedents and exercises developed at MIT and ideas about physical description developed at Stanford. The strength of the synthesis is illustrated by way of representative experiments. All of these experiments have been performed with an implementation system.

This research was done in part at the Artificial Intelligence Laboratory of Stanford University and in part by the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for MIT'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 N00014-80-C-0505.

Dietterich, Thomas G. and Ryszard S. Michalski., "Inductive Learning of Structural Descriptions," *Artificial Intelligence*, vol. 16, no. 3, 1981.

Freeman, P, and Allen Newell, "A Model for Functional Reasoning in Design," *Proceedings of the Second International Joint Conference on Artificial Intelligence*, London, England, 1971.

Katz, Boris and Patrick H. Winston, "Parsing and Generating English using Commutative Transformations," Artificial Intelligence Laboratory Memo No. 677, May, 1982. A version is also available as "A Two-way Natural Language Interface," in *Integrated Interactive Computing Systems*, edited by P. Degano and E. Sandewall, North-Holland, Amsterdam. 1982. Also see "A Three-Step Procedure for Language Generation," by Boris Katz, 'Artificial Intelligence Laboratory Memo No. 599, December, 1980.

Mitchell, Tom M., "Generalization as Search," *Artificial Intelligence*, vol. 18, no. 2, 1982.

Winston, Patrick Henry, "Learning Structural Descriptions from Examples," Ph.D. thesis, MIT, 1970. A shortened version is in *The Psychology of Computer Vision*, edited by Patrick Henry Winston, McGraw-Hill Book Company, New York. 1975.

Winston, Patrick Henry, "Learning and Reasoning by Analogy." *CACM*, vol. 23, no. 12. December, 1980. A version with details is available as "Learning and Reasoning by Analog! : the Details," M.I.T. Artificial Intelligence Laboratory Memo No. 520, April 1979.

Winston, Patrick Henry. "Learning New Principles from Precedents and Exercises," *Artificial Intelligence*, vol. 19, no. 3. A version with details is available as "Learning New Principles from Precedents and Exercises: the Details," M.I.T. Artificial Intelligence Laboratory Memo No. 632. May 1981.

Winston, Patrick Henry, "Learning by Augmenting Rules and Accumulating Censors." M.I.T. Artificial Intelligence Laboratory Memo No. 678. May 1982.

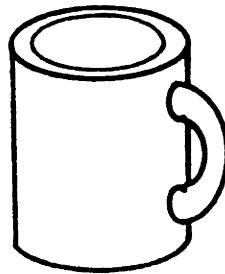


Figure 1. A cup with a handle.

Key Ideas

It is too hard to tell vision systems what things look like. It is easier to talk about purpose and what things are for. Consequently, we want vision systems to use functional descriptions to identify things, when necessary, and we want them to learn physical descriptions for themselves, when possible.

For example, there are many kinds of cups: some have handles, some do not; some have smooth cylindrical bodies, some are fluted: some are made of porcelain, others are styrofoam, and still others are metal. You could turn blue in the face describing all the physical possibilities. Functionally, however, all cups are things that are easy to drink from. Consequently, it is much easier to convey what cups are by saying what they are functionally.

To be more precise about what we are after, imagine that you are told cups are open vessels, standing stably, that you can lift. You see that the object in figure 1 has a handle, an upward pointing concavity, and a flat bottom. You happen to know it is light. Because you already know something about bowls, bricks, and suitcases, you conclude that you are looking at a cup. You also create a physical model covering this particular cup type.

Our first purpose, then, is to explain how physical identification can be done using functional definitions. Our second purpose is to show how to learn physical models using functional definitions and specific acts of identification.

It is important to note that our theory of model learning involves a physical example and some precedents in addition to the functional definition:

- The physical example is essential, for otherwise there would be no way to know which precedents are relevant.

Let E be an exercise. E is an exercise about a light object. The object's body is small. The object has a handle. The object's bottom is flat. Its concavity is upward-pointing. Its contents are hot. In E show that the object may be a cup.

Let E be an exercise. E is an exercise about a light object. The object's bottom is flat. Its body is small and cylindrical. Its concavity is upward-pointing. Its contents are hot. Its body's material is an insulator. In E show that the object may be a cup.

For the first of these two exercises, the rule requiring a handle works immediately. It is immaterial that the contents of the cup are hot.

For the second, the rule requiring a small, cylindrical body works immediately. Again it is immaterial that the contents of the cup are hot since nothing is known about the links among content temperature, graspability, and insulating materials. Proving some knowledge about these things by way of some censors makes identification more interesting.

Suppose, for example, that we teach or tell the machine that an object with hot contents will not have a graspable body, given no reason to doubt that the object's body is hot. Further suppose that we teach or tell the machine that an object's body is not hot, even if its contents are, if the body is made from an insulator. All this is captured by the following censor rules, each of which can make a simple physical deduction:

Let C1 be a Censor. C1 is a censor about an object. The object's body is not graspable because its contents are hot unless its body is not hot. Make C1 a censor using the object's body is not graspable.

Let C2 be a censor. C2 is a censor about an object. The object's contents are hot. Its body is not hot because its body's material is an insulator. Make C2 a censor using the object's body is not hot.

Repeating the second exercise now evokes the following scenario:

Asking whether the object is a cup activates the rule about cups without handles. The ifconditions of the rule are satisfied.

The *unless* conditions of the rule are checked. One of these conditions states that the object's body must not be plainly ungraspable.

Asking about graspability activates the censor relating graspability to hot contents. The censor's *if* condition is satisfied, and the censor is about to block the cup-identifying rule. The censor's *unless* condition must be checked first, however.

The censor's *unless* condition pertains to hot bodies. This condition activates a second censor, the one denying that a body is hot if it is made of an insulator. This second censor's *if* condition is satisfied, and there are no *unless* conditions.

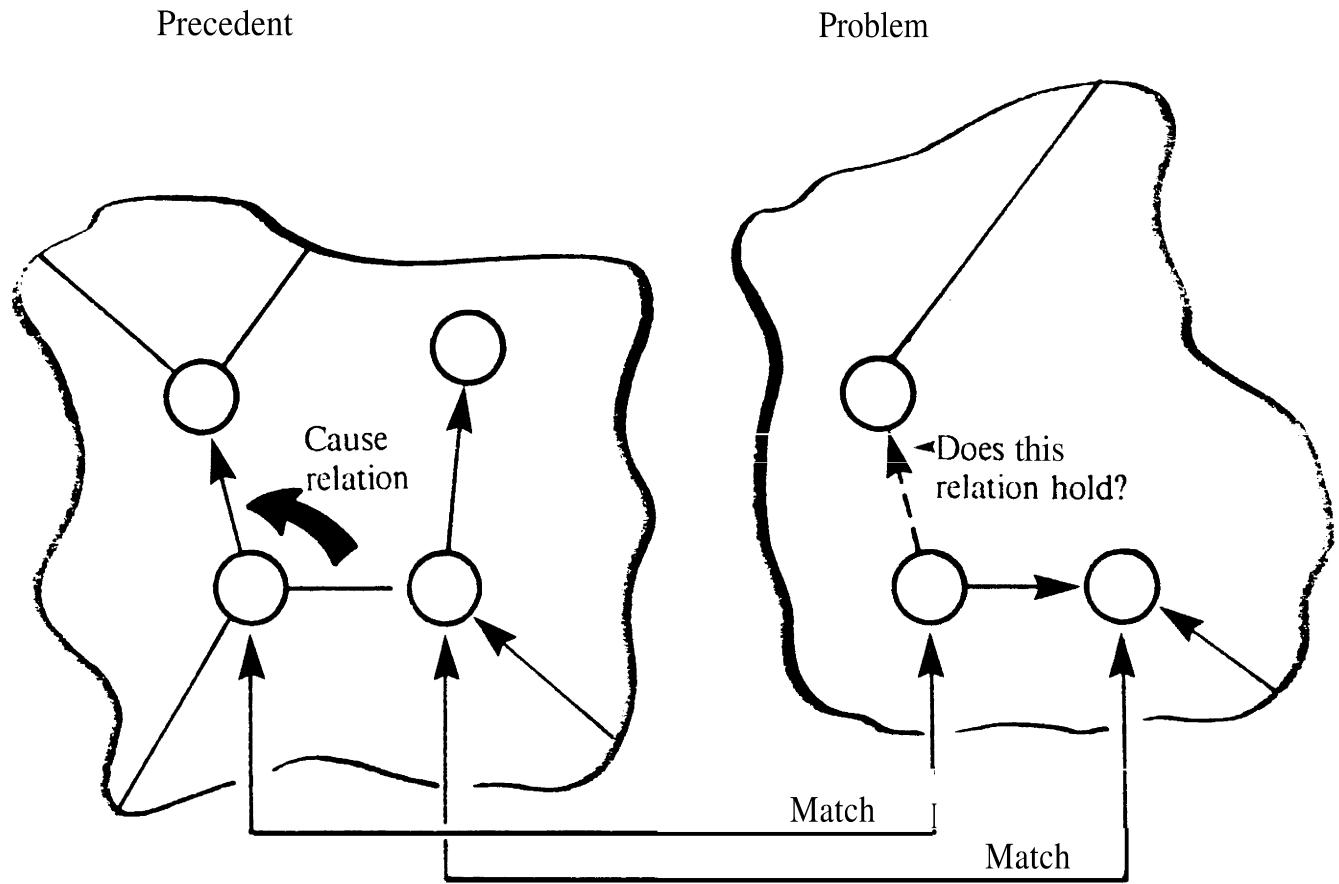


Figure 2. The matcher determines part correspondence using the links that populate the precedent and the problem. The matcher pays particular attention to links that are contained in the CAUSE structure of the precedent.

- ACRONYM uses generalized cylinders to describe how objects fill space.

A *generalized cylinder* is formed when a planar cross section moves along a curve in space, sweeping out a volume. The size of the planar cross section may change as it moves. The angle between the planar cross section and the curve is held constant, typically at 90°. Figure 4a shows some examples.

- ACRONYM uses ribbons and ellipses to represent what a viewer sees.

Ribbons are two-dimensional analogs to generalized cylinders. A *ribbon* is formed when a line is moved along a two-dimensional curve, perhaps changing size as it moves. The angle between the line and the curve is held constant, typically at 90°. Figure 4b shows some examples.

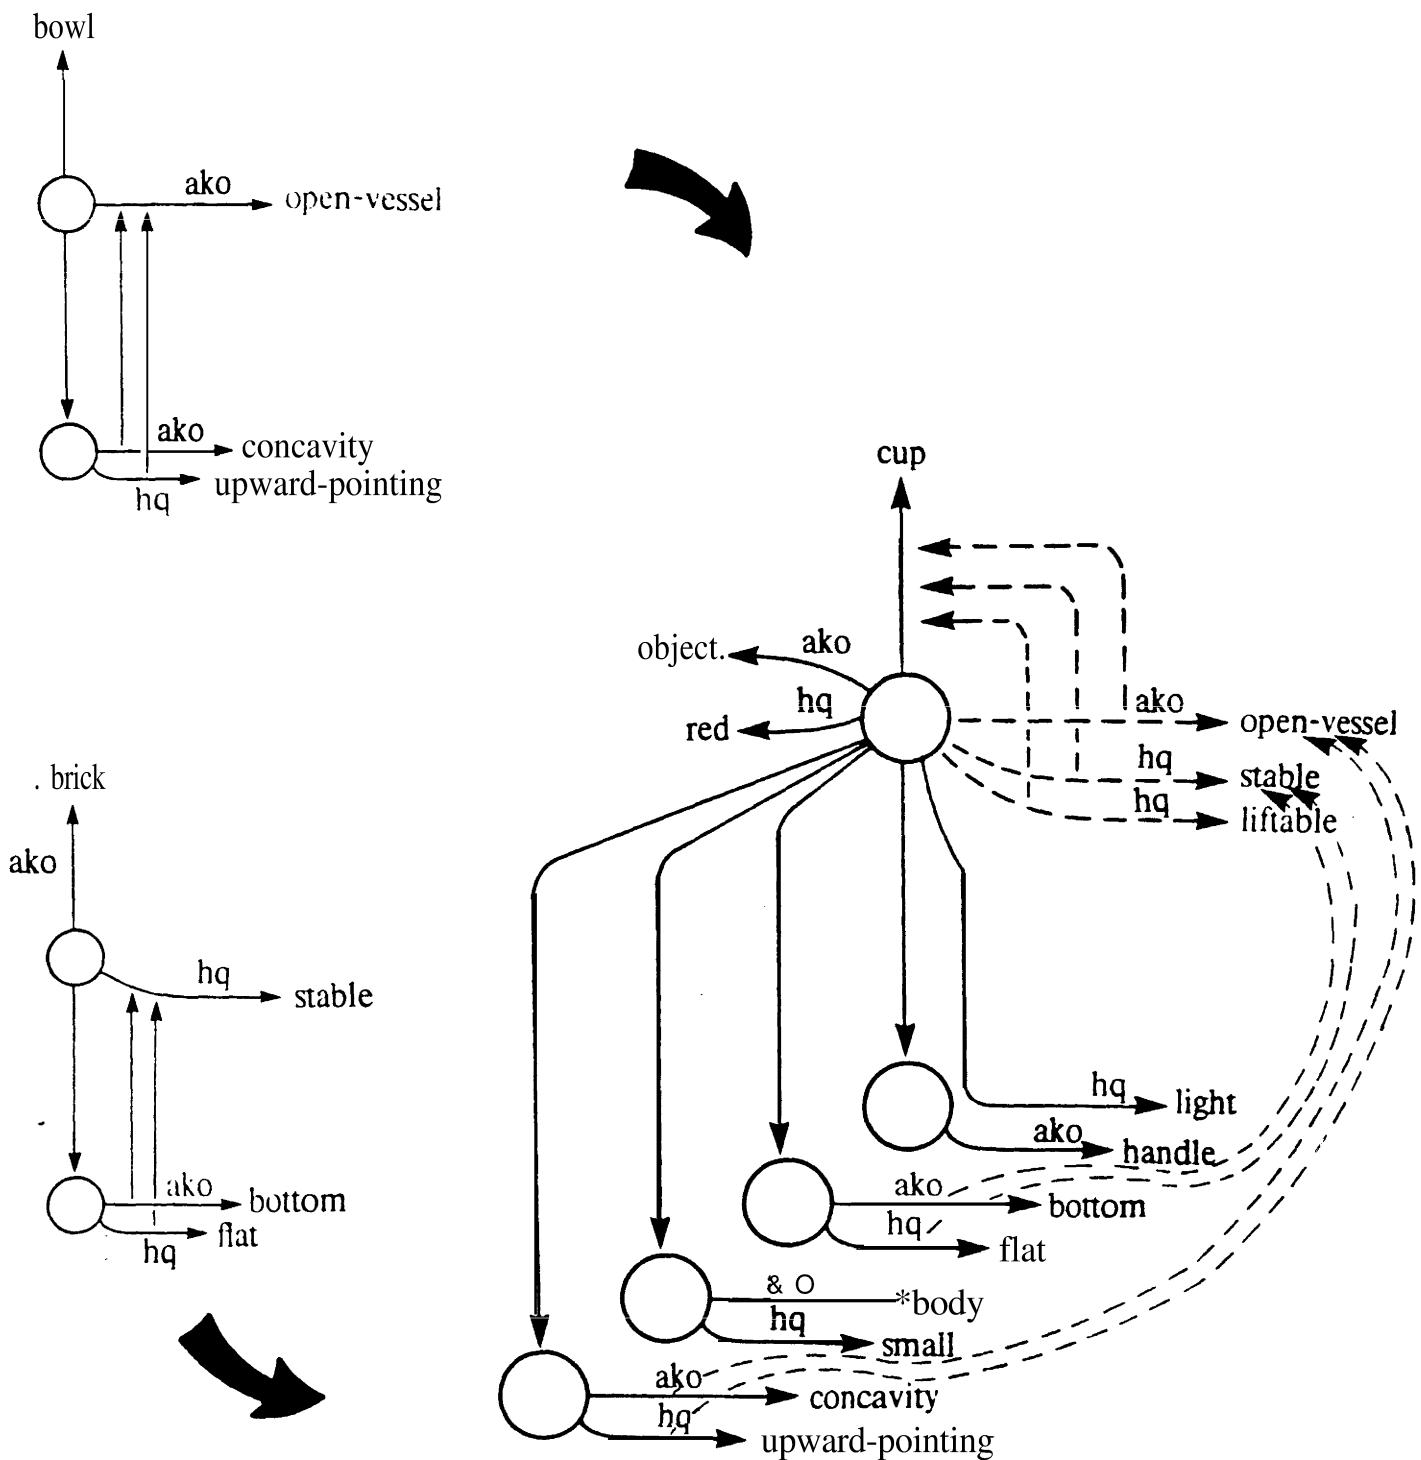


Figure 9. The brick precedent and the bowl precedent establish that the object is stable and that it is an open vessel. The cause links of the precedents are overlayed on the exercise, leading to questioned links that are immediately resolved by the facts. Overlayed structure is dashed. Many links of the precedents are not shown to avoid clutter on the diagram.

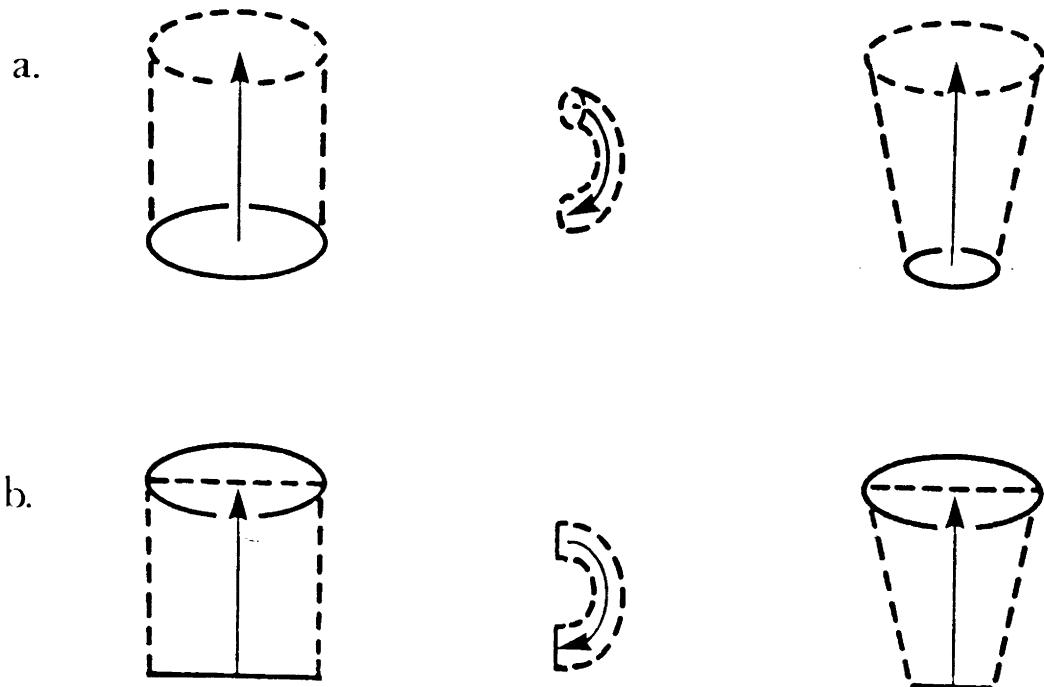


Figure 4. Part *a* shows some examples of generalized cylinders. Part *b* shows ribbons and ellipses corresponding to the objects in part *a*.

are found, the determination can be quite specific, not only about shape and size, but also about position. Having found the ribbon corresponding to the body of an airplane, for example, it is possible to predict the location, orientation, and size of the ribbons corresponding to the wings.

Brooks's landmark thesis concentrated on exactly this sort of prediction [Brooks 1981].

In principle, prediction knowledge can be used to condition the earliest vision procedures to the situation at hand. In current practice, early vision procedures operate autonomously up to the level where ribbons and ellipses are formed. Efforts are underway to push predictions further toward the pixels.

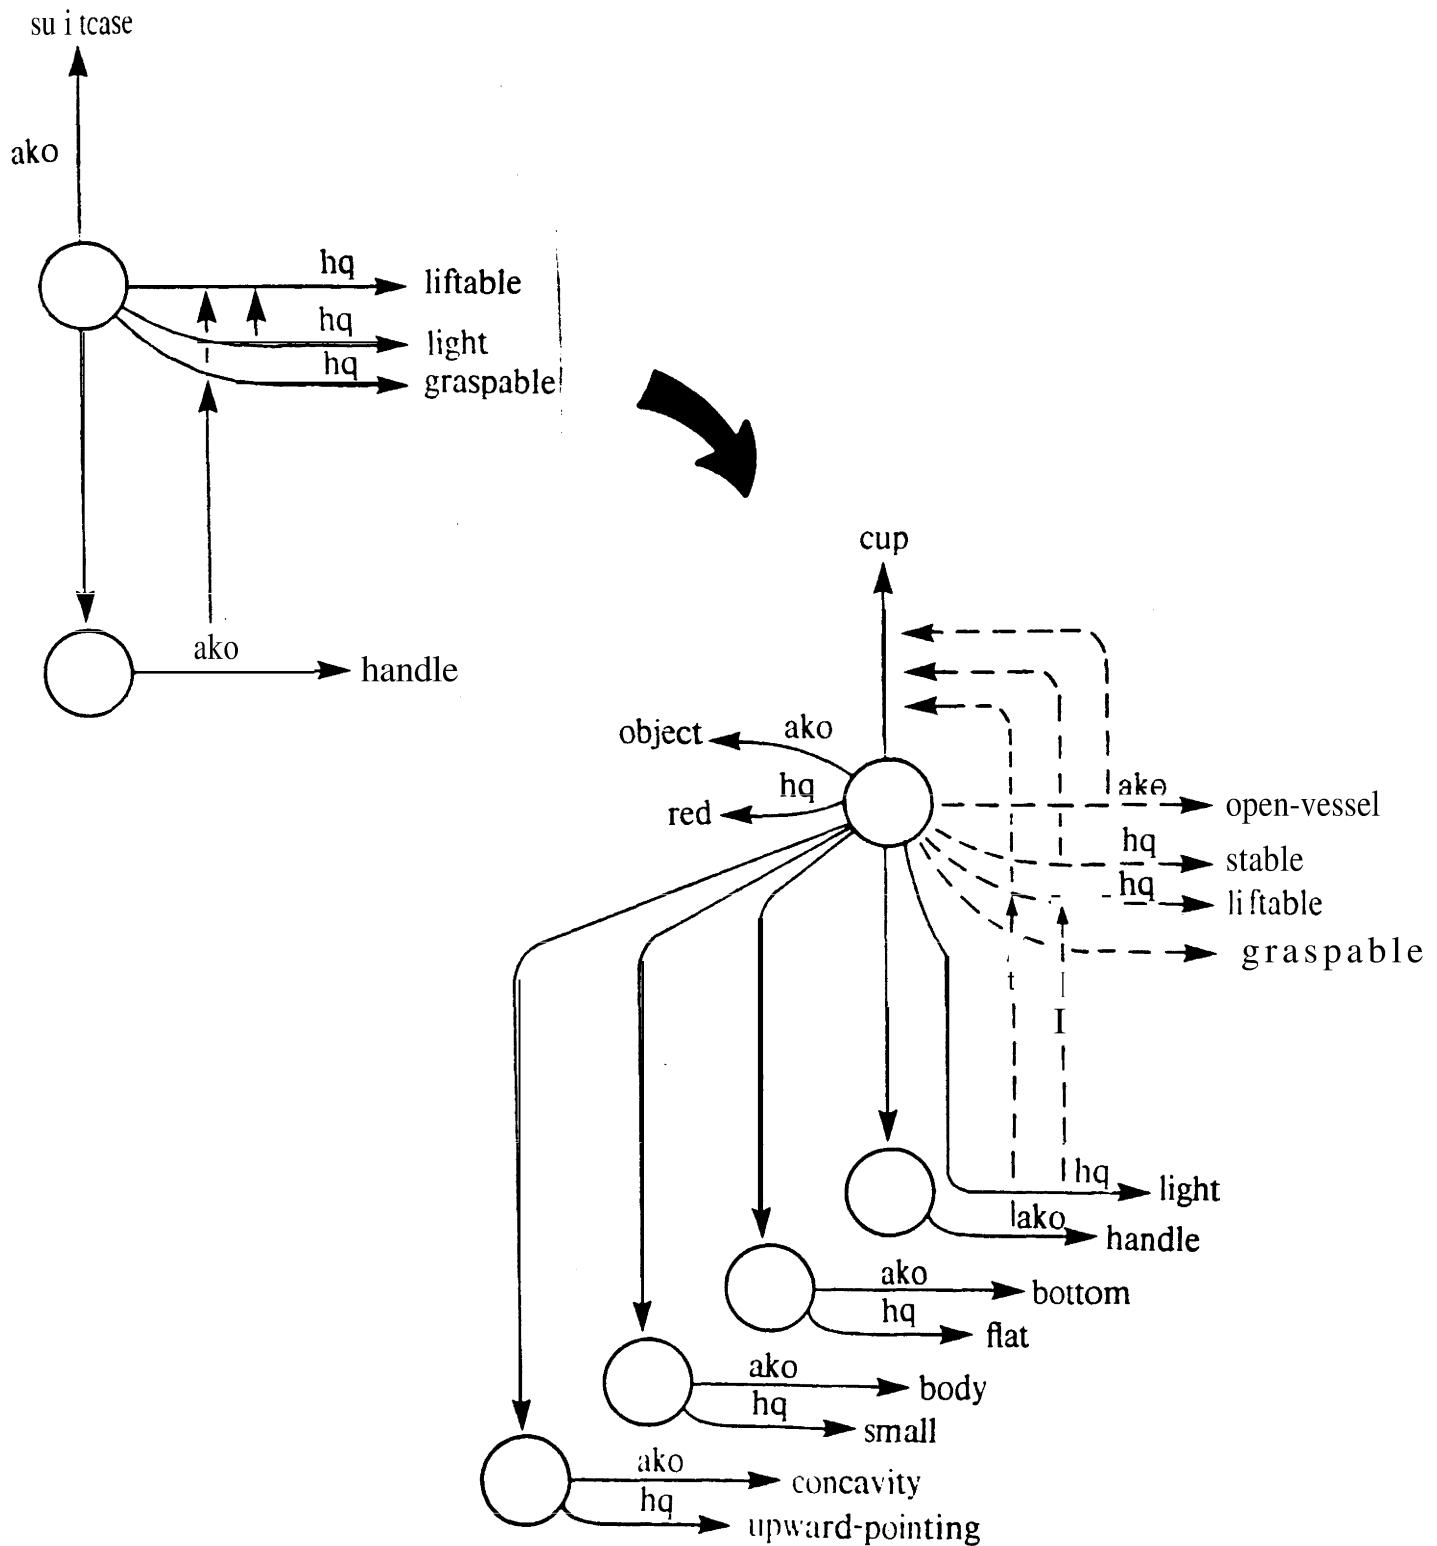


Figure 8. Cause links from the suitcase precedent are overlaid on the exercise, leading to questions about whether the object is light and whether the object has a handle. Overlaid structure is dashed. Many links of the suitcase precedent are not shown to avoid clutter on the diagram.

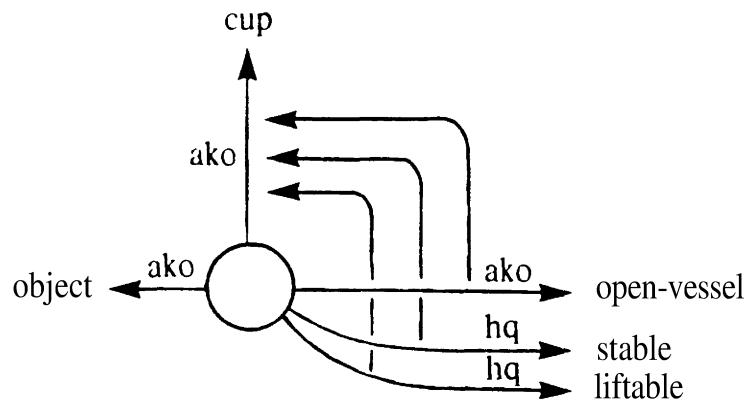


Figure 5. The functional definition of a cup. This semantic net is produced using an English description. $\Delta KO = \Delta$ Kind Of. $HQ =$ Has Quality.

Let X be a definition. X is a definition of an object. The object is a cup because it is stable, because it is liftable, and because it is an open-vessel. Remember X .

Of course, other, more elaborate definitions are possible. but this one seems to us to be good enough for the purpose of illustrating our learning theory.

The English is translated into the semantic net shown in figure 5.

The next step is to show an example of a cup, such as the one in figure 6a. ACRONYM is capable of translating such visual information into the semantic net shown in figure 6b. But inasmuch as our connection to ACRONYM is not complete, we currently bypass ACRONYM by using the following English instead.

Let E be an exercise. E is an exercise about a red object. The object's body is small. The object's bottom is flat. The object has an upward-pointing concavity.

The object has a handle.

In contrast to the definition, the qualities involved in the description of the particular cup are all physical qualities. not functional ones. (Assume that all qualities involving scales, like small size and light weight, are relative to the human body, by default, unless otherwise indicated.)

In the next step, we enhance the physical example's physical description. This enables us to specify physical properties and links that are not obtainable from vision.

The object is light.

Now it is time to show that the functional requirements are met by the enhanced physical description. To do this requires using precedents relating the

cup's functional descriptors to observed and stated physical descriptors. Three precedents are used. One indicates a way an object can be determined to be stable: another relates liftability to weight and having a handle: and still another explains what being an open-vessel means. All contain one thing that is irrelevant with respect to dealing with cups: these irrelevant things are representative of the detritus that can accompany the useful material.

Let X be a description. X is a description of a brick.

The brick is stable because the brick's bottom is flat.

The brick is hard.

Remember X.

Let X be a description. X is a description of a suitcase.

The suitcase is liftable because it is graspable and because it is light. The suitcase is graspable because it has a handle.

The suitcase is useful because it is a portable container for clothes.

Remember X.

Let X be a description. X is a description of a bowl.

The bowl is an open-vessel because it has a concavity and because the concavity is upward-pointing.

The bowl contains tomato soup.

Remember X.

With the functional definition in hand, together with relevant precedents, the analogy apparatus is ready to work as soon as it is stimulated by the following challenge:

In E, show that the object may be a cup.

This initiates a search for precedents relevant to showing something is a cup. The functional definition is retrieved. Next, a matcher determines the correspondence between parts of the exercise and the parts of the functional definition, a trivial task in this instance. Now the verifier overlays the cause links of the functional definition onto the exercise. Tracing through these overlayed cause links raises three questions: is the observed object stable, is it an open vessel, and is it liftable. All this is illustrated in figure 7.

Questioning if the object is liftable leads to a second search for a precedent, this time one that relates function to form, causing the suitcase description to be retrieved. The suitcase description, shown in figure 8, is matched to the exercise, its causal structure is overlayed on the exercise, and other questions are raised: is the observed object light and does it have a handle. Since it is light and does have a handle, the suitcase description suffices to deal with the liftable issue, leaving open the stability and open-vessel questions.

cup's functional descriptors to observed and stated physical descriptors. Three precedents are used. One indicates a way an object can be determined to be stable; another relates liftability to weight and having a handle; and still another explains what being an open-vessel means. All contain one thing that is irrelevant with respect to dealing with cups: these irrelevant things are representative of the detritus that can accompany the useful material.

Let X be a description. X is a description of a brick.
The brick is stable because the brick's bottom is flat.

The brick is hard.

Remember X.

Let X be a description. X is a description of a suitcase.
The suitcase is liftable because it is graspable and because
it is light. The suitcase is graspable because it has a
handle.

The suitcase is useful because it is a portable container
for clothes.

Remember X.

Let X be a description. X is a description of a bowl.
The bowl is an open-vessel because it has a concavity
and because the concavity is upward-pointing.

The bowl contains tomato soup.

Remember X.

With the functional definition in hand, together with relevant precedents, the analogy apparatus is ready to work as soon as it is stimulated by the following challenge:

In E, show that the object may be a cup.

This initiates a search for precedents relevant to showing something is a cup. The functional definition is retrieved. Next, a matcher determines the correspondence between parts of the exercise and the parts of the functional definition, a trivial task in this instance. Now the verifier overlays the cause links of the functional definition onto the exercise. Tracing through these overlayed cause links raises three questions: is the observed object stable, is it an open vessel, and is it liftable. All this is illustrated in figure 7.

Questioning if the object is liftable leads to a second search for a precedent, this time one that relates function to form, causing the suitcase description to be retrieved. The suitcase description, shown in figure 8, is matched to the exercise, its causal structure is overlayed on the exercise, and other questions are raised: is the observed object light and does it have a handle. Since it is light and does have a handle, the suitcase description suffices to deal with the liftable issue, leaving open the stability and open-vessel questions.

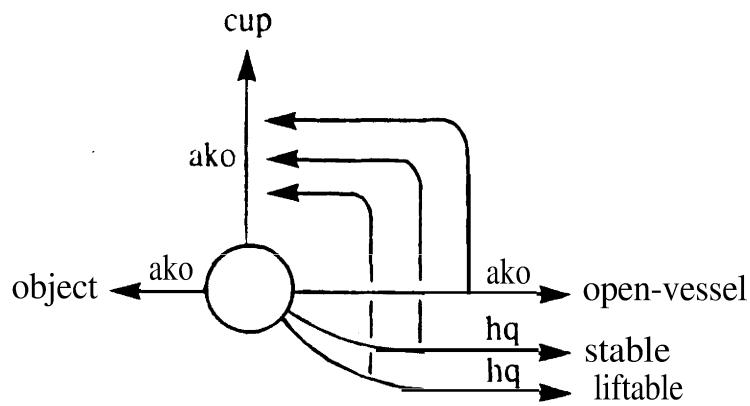


Figure 5. The functional definition of a cup. This semantic net is produced using an English description. AKO = A Kind Of, HQ = Has Quality.

Let X be a definition. X is a definition of an object. The object is a cup because it is stable, because it is liftable, and because it is an open-vessel. Remember X.

Of course, other, more elaborate definitions are possible. but this one seems to us to be good enough for the purpose of illustrating our learning theory.

The English is translated into the semantic net shown in figure 5.

The next step is to show an example of a cup, such as the one in figure 6a. ACRONYM is capable of translating such visual information into the semantic net shown in figure 6b. But inasmuch as our connection to ACRONYM is not complete, we currently bypass ACRONYM by using the following English instead.

Let E be an exercise. E is an exercise about a red object. The object's body is small. The object's bottom is flat. The object has an upward-pointing concavity. The object has a handle.

In contrast to the definition, the qualities involved in the description of the particular cup are all physical qualities, not functional ones. (Assume that all qualities involving scales, like small size and light weight, are relative to the human body, by default, unless otherwise indicated.)

In the next step, we enhance the physical example's physical description. This enables us to specify physical properties and links that are not obtainable from vision.

The object is light.

Now it is time to show that the functional requirements are met by the enhanced physical description. To do this requires using precedents relating the

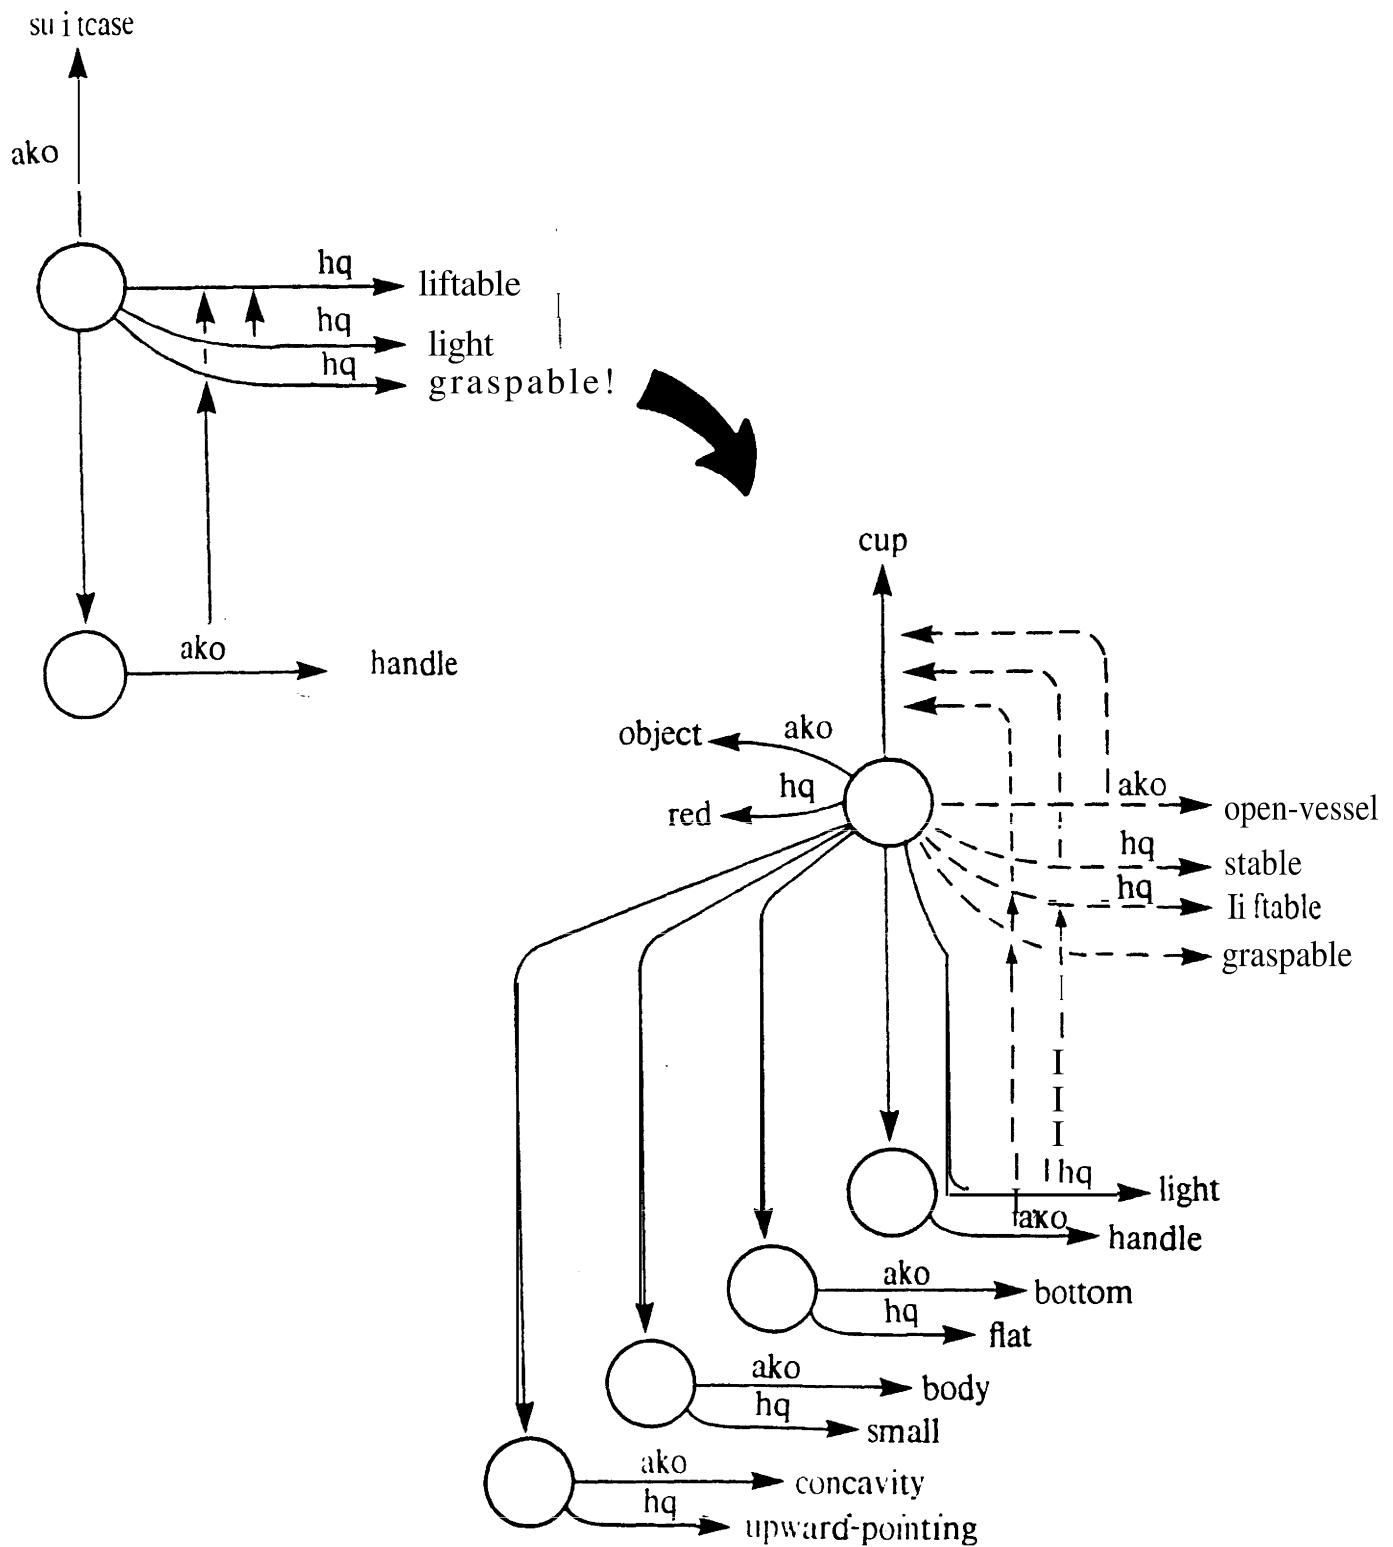


Figure 8. Cause links from the suitcase precedent are overlaid on the exercise, leading to questions about whether the object is light and whether the object has a handle. Overlaid structure is dashed. Many links of the suitcase precedent are not shown to avoid clutter on the diagram.

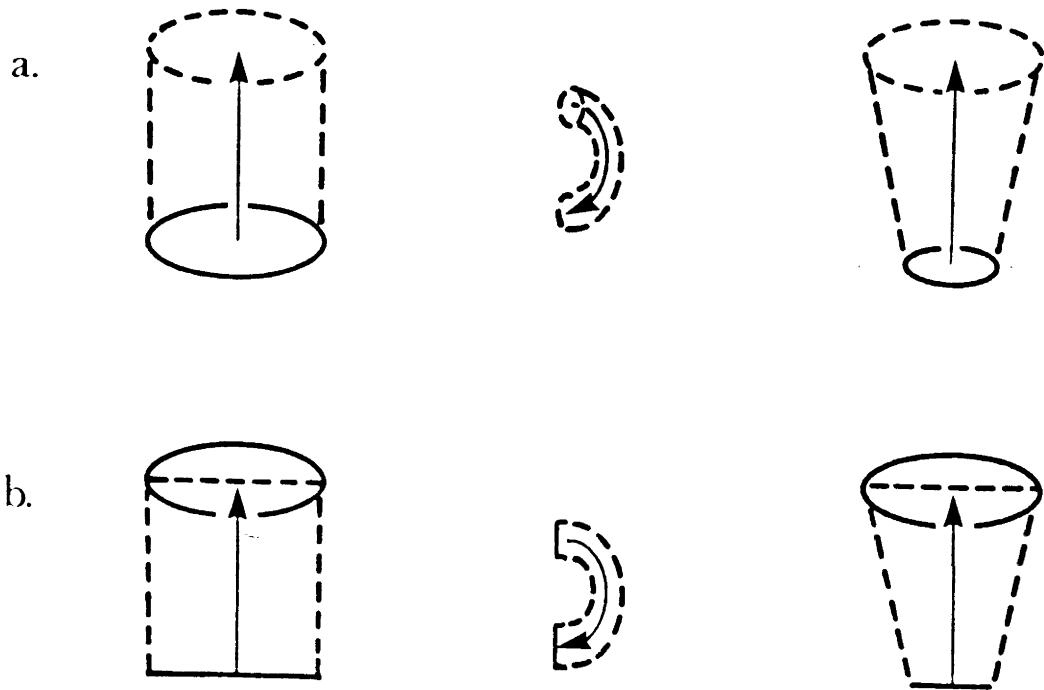


Figure 4. Part a shows some examples of generalized cylinders. Part b shows ribbons and ellipses corresponding to the objects in part a.

are found, the determination can be quite specific, not only about shape and size, but also about position. Having found the ribbon corresponding to the body of an airplane, for example, it is possible to predict the location, orientation, and size of the ribbons corresponding to the wings.

Brooks's landmark thesis concentrated on exactly this sort of prediction [Brooks 1981].

In principle, prediction knowledge can be used to condition the earliest vision procedures to the situation at hand. In current practice, early vision procedures operate autonomously up to the level where ribbons and ellipses are formed. Efforts are underway to push predictions further toward the pixels.

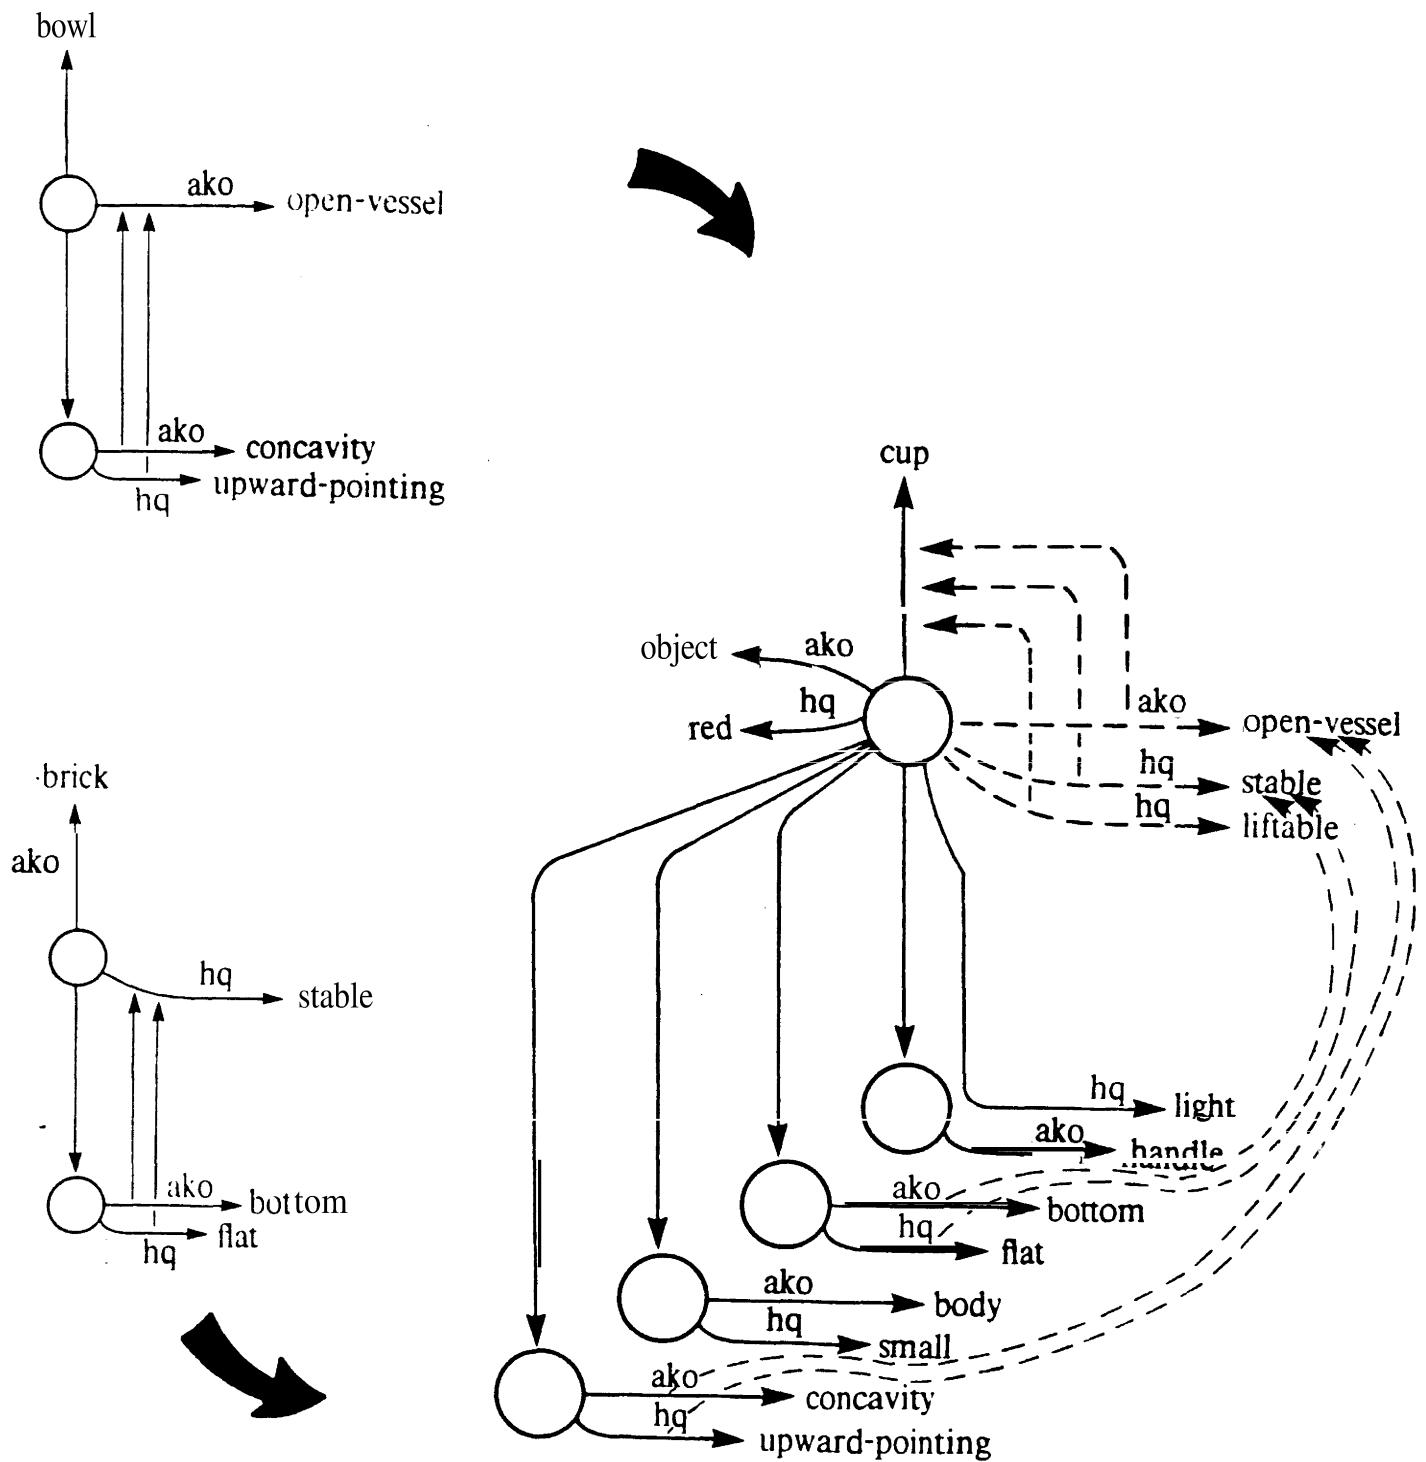


Figure 9. The brick precedent and the bowl precedent establish that the object is stable and that it is an open vessel. The cause links of the precedents are overlaid on the exercise, leading to questioned links that are immediately resolved by the facts. Overlayed structure is dashed. Many links of the precedents are not shown to avoid clutter on the diagram.

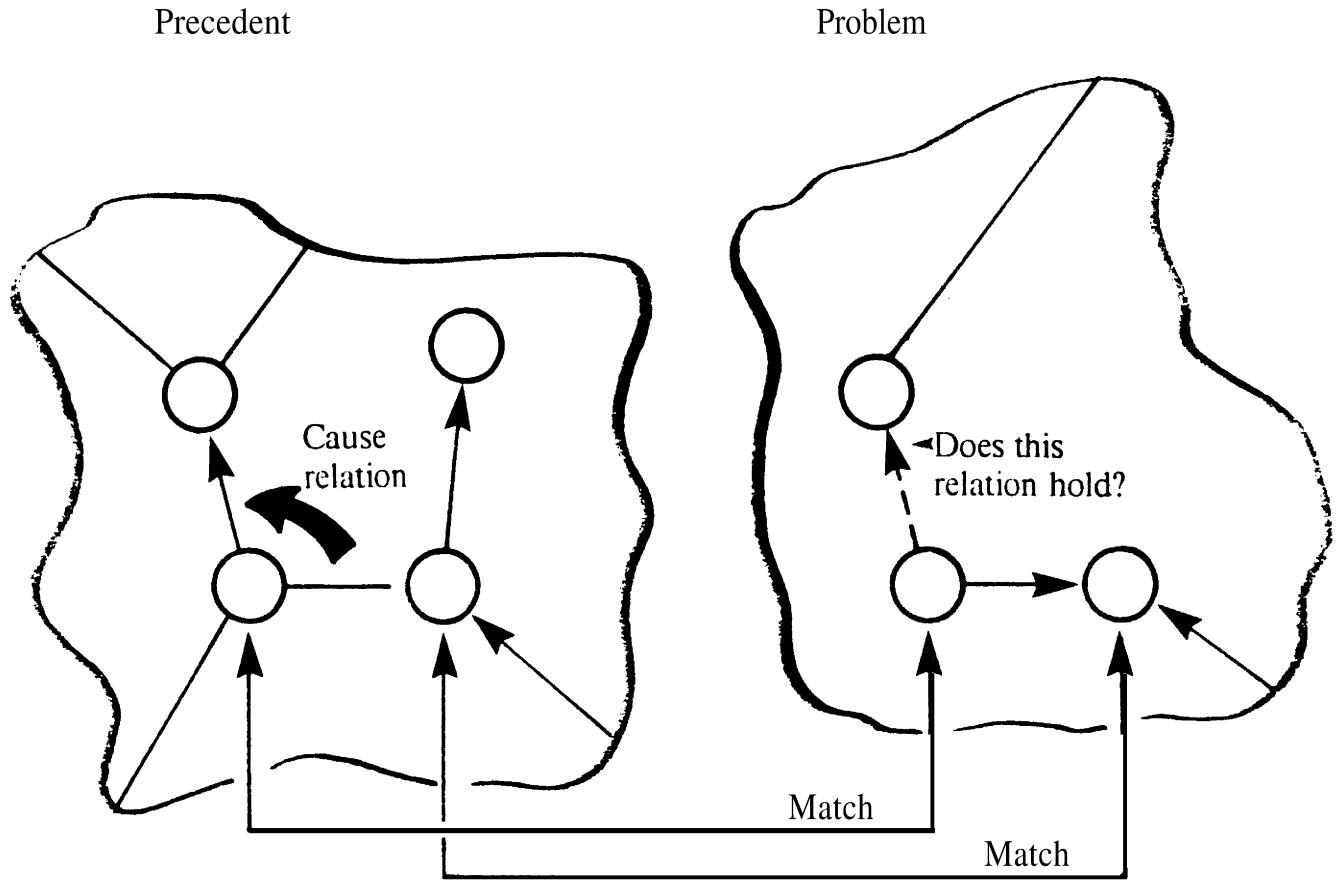


Figure 2. The matcher determines part correspondence using the links that populate the precedent and the problem. The matcher pays particular attention to links that are contained in the CAUSE structure of the precedent.

- ACRONYM uses generalized cylinders to describe how objects fill space.

A *generalized cylinder* is formed when a planar cross section moves along a curve in space, sweeping out a volume. The size of the planar cross section may change as it moves. The angle between the planar cross section and the curve is held constant, typically at 90°. Figure 4a shows some examples.

- ACRONYM uses ribbons and ellipses to represent what a viewer sees.

Ribbons are two-dimensional analogs to generalized cylinders. A *ribbon* is formed when a line is moved along a two-dimensional curve, perhaps changing size as it moves. The angle between the line and the curve is held constant, typically at 90°. Figure 4b shows some examples.

Let E be an exercise. E is an exercise about a light object. The object's body is small. The object has a handle. The object's bottom is flat. Its concavity is upward-pointing. Its contents are hot. In E show that the object may be a cup.

Let E be an exercise. E is an exercise about a light object. The object's bottom is flat. Its body is small and cylindrical. Its concavity is upward-pointing. Its contents are hot. Its body's material is an insulator. In E show that the object may be a cup.

For the first of these two exercises, the rule requiring a handle works immediately. It is immaterial that the contents of the cup are hot.

For the second, the rule requiring a small, cylindrical body works immediately. Again it is immaterial that the contents of the cup are hot since nothing is known about the links among content temperature, graspability, and insulating materials. Proving some knowledge about these things by way of some censors makes identification more interesting.

Suppose, for example, that we teach or tell the machine that an object with hot contents will not have a graspable body, given no reason to doubt that the object's body is hot. Further suppose that we teach or tell the machine that an object's body is not hot, even if its contents are, if the body is made from an insulator. All this is captured by the following censor rules, each of which can make a simple physical deduction:

Let C1 be a Censor. C1 is a censor about an object. The object's body is not graspable because its contents are hot unless its body is not hot. Make C1 a censor using the object's body is not graspable.

Let C2 be a censor. C2 is a censor about an object. The object's contents are hot. Its body is not hot because its body's material is an insulator. Make C2 a censor using the object's body is not hot.

Repeating the second exercise now evokes the following scenario:

Asking whether the object is a cup activates the rule about cups without handles. The *if* conditions of the rule are satisfied.

The *unless* conditions of the rule are checked. One of these conditions states that the object's body must not be plainly ungraspable.

Asking about graspability activates the censor relating graspability to hot contents. The censor's *if* condition is satisfied, and the censor is about to block the cup-identifying rule. The censor's *unless* condition must be checked first, however.

The censor's *unless* condition pertains to hot bodies. This condition activates a second censor, the one denying that a body is hot if it is made of an insulator. This second censor's *if* condition is satisfied, and there are no *unless* conditions.

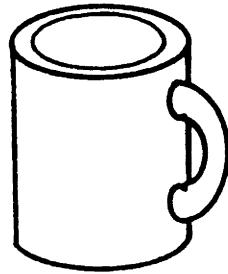


Figure 1. A cup with a handle.

Key Ideas

It is too hard to tell vision systems what things look like. It is easier to talk about **purpose** and what things are for. Consequently, we want vision systems to use functional descriptions to identify things, when necessary, and we want them to learn physical descriptions for themselves, when possible.

For example, there are many kinds of cups: some have handles, some do not; some have smooth cylindrical bodies, some are fluted: some are made of porcelain, others are styrofoam, and still others are metal. You could turn blue in the face &scribing all the physical possibilities. Functionally, however, all cups are things that are easy to drink from. Consequently, it is much easier to convey what cups are by saying what they are functionally.

To be more precise about what we are after, imagine that you are told cups are open vessels, standing stably, that you can lift. You see that the object in figure 1 has a handle, an upward pointing concavity, and a flat bottom. You happen to know it is light. Because you already know something about bowls, bricks, and suitcases, you conclude that you are looking at a cup. You also create a physical model covering this particular cup type.

Our first purpose, then, is to explain how physical identification can be done using functional definitions. Our second purpose is to show how to learn physical models using functional definitions and specific acts of identification.

It is important to note that our theory of model learning involves a **physical** example and some precedents in addition to the functional definition:

- The physical example is essential, for otherwise there would be no way to know which precedents are relevant.

Dictterich, Thomas G. and Ryszard S. Michalski., "Inductive Learning of Structural Descriptions," *Artificial Intelligence*, vol. 16, no. 3, 1981.

Freeman, P, and Allen Newell, "A Model for Functional Reasoning in Design," *Proceedings of the Second International Joint Conference on Artificial Intelligence*, London, England, 1971.

Katz, Boris and Patrick H. Winston, "Parsing and Generating English using Commutative Transformations," Artificial Intelligence Laboratory Memo No. 677, May, 1982. A version is also available as "A Two-way Natural Language Interface," in *Integrated Interactive Computing Systems*, edited by P. Degano and E. Sandewall, North-Holland, Amsterdam. 1982. Also see "A Three-Step Procedure for Language Generation," by Boris Katz, Artificial Intelligence Laboratory Memo No. 599, December, 1980.

Mitchell, Tom M., "Generalization as Search," *Artificial Intelligence*, vol. 18, no. 2, 1982.

Winston, Patrick Henry, "Learning Structural Descriptions from Examples," Ph.D. thesis, MIT, 1970. A shortened version is in *The Psychology of Computer Vision*, edited by Patrick Henry Winston, McGraw-Hill Book Company, New York. 1975.

Winston, Patrick Henry, "Learning and Reasoning by Analogy." *CACM*, vol. 23, no. 12. December. 1980. A version with details is available as "Learning and Reasoning by Analog! : the Details," M.I.T. Artificial Intelligence Laboratory Memo No. 520, April 1979.

Winston, Patrick Henry, "Learning New Principles from Precedents and Exercises," *Artificial Intelligence*, vol. 19, no. 3. A version with details is available as "Learning New Principles from Precedents and Exercises: the Details," M.I.T. Artificial Intelligence Laboratory Memo No. 632, May 1981.

Winston, Patrick Henry, "Learning by Augmenting Rules and Accumulating Censors." M.I.T. Artificial Intelligence Laboratory Memo No. 678, May 1982.

Abstract

It is too hard to tell vision systems what things look like. It is easier to talk about purpose and what things are for. Consequently, we want vision systems to use functional descriptions to identify things, when necessary, and we want them to learn physical descriptions for themselves, when possible.

This paper describes a theory that explains how to make such systems work. The theory is a synthesis of two sets of ideas: ideas about learning from precedents and exercises developed at MIT and ideas about physical description developed at Stanford. The strength of the synthesis is illustrated by way of representative experiments. All of these experiments have been performed with an implementation system.

This research was done in part at the Artificial Intelligence Laboratory of Stanford University and in part by the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for MIT'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 N00014-80-C-0505.

.
:
:

To use DEFINITION-1 I need to know if [OBJECT-J **AKO** OPEN-VESSEL]
I am trying to show [OBJECT-3 **AKO** OPEN-VESSEL]
Supply y,n,?,r = rules, p = precedents, or a suggestion:
> p
I find:
 DESCRIPTION-2 < 3. Tinks >
I note [CONCAVITY-3 PHYSICAL-PART-OF OBJECT-31 for use with **DESCRIPTION-2**
I note [CONCAVITY-3 **AKO** CONCAVITY] for use with DESCRIPTION-2
I note [CONCAVITY-3 HQ UPWARD-POINTING] for use with DESCRIPTION-2
The evidence from DESCRIPTION-2 indicates [OBJECT-3 **AKO** OPEN-VESSEL)
The evidence from DEFINITION-1 indicates [OBJECT-3 **AKO** CUP]
Rule RULE-1 is derived from DEFINITION-1 DESCRIPTION-2 DESCRIPTION-3
DESCRIPTION-1 and looks like this:

Rule
 RULE-1
if
 [OBJECT-9 HQ LIGHT]
 [CONCAVITY-7 PHYSICAL-PART-OF OBJECT-91]
 [HANDLE-4 PHYSICAL-PART-OF OBJECT-91]
 [BOTTOM-7 PHYSICAL-PART-OF OBJECT-91]
 [CONCAVITY-7 **AKO** CONCAVITY]
 [CONCAVITY-7 HQ UPWARD-POINTING]
 [HANDLE-4 **AKO** HANDLE]
 [BOTTOM-7 **AKO** BOTTOM]
 [BOTTOM-7 HQ FLAT]
then
 [OBJECT-g **AKO** CUP]
unless
 [[OBJECT-g **AKO** OPEN-VESSEL] HQ FALSE]
 [[OBJECT-9 HQ LIFTABLE] HQ FALSE]
 [[OBJECT-g HQ GRASPABLE] HQ FALSE]
 [[OBJECT-9 HQ STABLE] HQ FALSE]
case
 DEFINITION-1 DESCRIPTION-2 DESCRIPTION-3 DESCRIPTION-1

- Should I index it as a rule?
> y

Artificial Intelligence Laboratory
Stanford University

in collaboration with

Artificial Intelligence Laboratory
Massachusetts Institute of Technology

AIM Memo 349
STAN-CS-82-950

November 1982
Revised January 1983

Learning Physical Descriptions
From Functional Definitions,
Examples, and Precedents

by

Patrick H. Winston*

Thomas O. Binford

Boris Katz*

and Michael Lowry

*Patrick H. Winston and Boris Katz are at the Massachusetts Institute of Technology Artificial Intelligence lab.
This paper is also available from Massachusetts Institute of Technology as M.I.T. AI Memo 679

```
Rule
  RULE-Z
  if
    [OBJECT-10 HQ LIGHT]
    [CONCAVITY-B PHYSICAL-PART-OF OBJECT-10]
    [BODY-9 PHYSICAL-PART-OF OBJECT-10 ]
    [BOTTOM-8 PHYSICAL-PART-OF OBJECT-10)
    [CONCAVITY-8 AKO CONCAVITY]
    [CONCAVITY-8 HQ UPWARD-POINTING]
    [BODY-9 AKO BODY]
    [BODY-9 HQ CYLINDRICAL]
    [BODY -9 HQ SMALL]
    [BOTTOM-8 AKO BOTTOM]
    [BOTTOM-8 HQ FLAT]
  then
    [OBJECT-10 AKO CUP]
  unless
    [ [OBJECT-10 AKO OPEN-VESSEL] HQ FALSE]
    [ [OBJECT-10 HQ LIFTABLE] HQ FALSE]
    [ [OBJECT-10 HQ STABLE] HQ FALSE]
    [[BODY-9 HQ GRASPABLE] HQ FALSE]
  case
    DEFINITION-1 DESCRIPTION-2 DESCRIPTION-4 DESCRIPTION-1

Should I index it as a rule?
> Y
```

January 1983
Also numbered AIM-349

Report No. STAN-CS-82-950

Learning Physical Description from Functional Definitions, Examples and Precedents

by

Patrick H. Winston, Thomas O. **Binford**, Boris Katz
and Michael Lowry

Department of Computer Science

Stanford University
Stanford, CA 94305

