Device Representation for Modeling Improvisation in Mechanical Use Situations*

Jack Hodges

Computer Science Department
University of California, Los Angeles

Abstract

Improvisation requires an understanding and application of mechanical objects in broad contexts. The capacity to interpret a situation in terms of an object's capabilities requires the integration of functional and behavioral object representations. A model is presented which describes the integration of causal interactions between these levels of abstraction. The model maintains both intentional and behavioral representations to allow inferencing at each level, but integrates them by applying an inferencing mapping between the two. This model is used to reason about simple mechanical objects in the domain of improvisational mechanics.

INTRODUCTION

When people have to resolve problems involving mechanical objects in real-life situations, they must make decisions based on conflicting goals and constraints at both the functional and behavioral level[1]. Even though a problem-solver may recognize a behavioral advantage of one object over another, their higher-level personal goals may cause them to try objects based on functional capabilities. Consider the following example of improvisation where these differences lead to a goal failure:

Broken Knife

A man wants to polish one of his silver candlesticks. He must therefore pry open a can of silver polish in the kitchen, but doesn't want to brave winter weather to get a screwdriver from the garage. He reasons that he can use a screwdriver-like object and decides to try a carving-knife. What he doesn't realize is the knife is not strong enough in the dimension relevant for prying. The knife blade breaks.

[1.] Functional descriptions refer to the intended use of objects, as opposed to behavioral descriptions, which describe physical interactions between objects.

There are many representational issues in **Broken Knife**, spanning the situational, intentional, functional and behavioral reasoning levels. At the situational level, planning choices are dictated by the relationships between the man and such contextual elements as the winter weather and objects available in the kitchen. On the intentional planning level, the man has chosen the Polish-Metallic plan to preserve his candlesticks. This plan requires that he have silver polish on his rag, a state which is blocked by the fact that the silver polish can is closed. Recognizing that the only resolution is to pry the can open, he realizes that the tool he usually uses for this function, a screwdriver, is in the garage. There is now a goal conflict: between his goal of preserving the candlesticks and his goal to preserve his own comfort.

Here the functional level becomes significant. A screwdriver works as a prying tool for the silver polish can because it fits into the slot between the can and lid and is strong in relation to the force necessary to pry open the lid. A carving-knife will fit into the slot and was strong enough for the functions that it was used for in the past. The knife therefore apparently matches the constraints for Pry-Object, so the man uses it.

Finally there is the behavioral level. The knife is indeed strong, but only in the context of carving and along the width of the knife's blade. Along the thickness of the knife's blade, where the force of prying will be borne, the knife is not strong - not in relation to the friction force holding the lid onto the can. The knife blade therefore breaks.

We have been interested in modeling improvisation situations like **Broken Knife** in hopes of better understanding the creative process during problem-solving. Improvisation is a kind of invention where the problem-solver is constrained by circumstance. Improvisation thus encompasses the scope of Edison, an on-going project to model the knowledge and reasoning of naive inventors, people whose knowledge and planning is based on experience rather than technical expertise [Dyer, Hodges & Flowers, 1986]. Our claim is that any approach to real-life problem-solving and decision-making must integrate each of the above levels of abstraction into a complete system. Previous object models have emphasized object representation at the functional or behavioral level, but none have integrated the two into a single representation and processing mechanism. Edison has been designed to achieve this integration, and to support the associated multi-level reasoning.

REPRESENTING OBJECT FUNCTION AND BEHAVIOR

Intentional representation models have traditionally described objects with an emphasis toward their intended uses, while behavioral models have empha-

sized their behavioral capabilities. Each model type has been successful in its respective domain, but either could benefit from the capabilities of the other.

Intentional and Functional Object Models

Intentional object models, such as conceptual-dependency (CD) [Schank & Abelson, 1977], describe objects by an agent's intentions and how an object's function affects the outcome of those intentions (i.e. objects are black boxes). Using CD notation, the act of throwing a ball in a game of catch is represented by the thrower Propeling the ball toward the catcher while unGrasping it. The resulting state enables the ball to Ptrans from the thrower's location to the catcher's location. With this kind of model inferences can be made about the relationship between the people playing (e.g. "John threw the ball to Bill." vs. "John threw the ball at Bill."), but not about the ball involved (e.g. what if the ball never reaches Bill). This limitation presents a problem for predicting and explaining how plans are affected by object function and behavior.

Lehnert's object primitives [Lehnert, 1978] and Rieger's common sense algorithm (CSA) [Rieger, 1985] integrated object functionality into CD to describe how and when objects are used. These models introduced the idea of a functional representation level, between intentional and mechanical levels, which had properties found in both. Unfortunately, both models had weak behavioral representations and blurred the distinction between object function and behavior. They were therefore unable to take full advantage of their functional representations.

Behavioral Object Models

Behavioral object models describe objects' physical composition and interactions in lieu of their intended purpose or context. Instead of action primitives based on some form of agency, the primitives in behavioral models are simple qualitative physical process descriptions [Forbus, 1985] which describe objects and their interactions. The actor's Propel and Grasp actions (in a game of Catch) result in Force and Constraint states which enable the process, Transmit, of force to the ball. The ball unGrasping is paralleled by a Constrain process, and the resulting Force and Constraint states enable a Move process. Behavioral models are useful for predicting, explaining and simulating the ball's behavior (e.g. when the ball's weight, force and direction are known), but not for describing how or why it was thrown in the first place. Another problem with behavioral models is that they don't utilize contextual and intentional information for disambiguating, or predicting, object function during problem-solving.

Representing Objects In Edison

Edison is an object-based representational model for

reasoning about situations like **Broken Knife** by integrating object knowledge derived from intentional and behavioral points of view. The intentional part of Edison's bi-level model considers the object as an instrument to achieving specific goals in specific contexts. The behavioral part of Edison considers the object and its behavioral dependencies with other objects. This integration is achieved by considering the structural continuity which must be maintained to support inferences between these abstraction levels, and by considering a third, *functional*, part which overlaps the intentional and behavioral abstraction levels and provides for a continuous inference path between them.

The objects described in Edison are simple mechanical devices, such as screwdrivers, hammers, knives, can openers, and nail clippers. In Edison, the representational emphasis is on the physical qualities and relations which support an object's functional description. Most of the reasoning in Edison is done at higher levels, so the simulator is only used for diagnosis and explanation. This contrasts to detailed qualitative simulators, such as [Doyle, 1988], designed for this purpose. The Edison model represents all objects as combinations of primitive devices (such as levers, springs, and wheels) which effect the leverage mechanism through the Transform process [Hodges, Dyer & Flowers, 1987. All object behavior can thus be described in terms of the transmission, translation, or magnification of force and motion.

Object *functions* refer to the tasks an object has been or could be applied to in a particular context, and have both intentional and behavioral qualities. Using a knife to carve turkey, to threaten someone, to tighten screws, or to pry can lids all describe knife functions. At the intentional level object functions describe this context sensitivity through *attributes*, which are qualities associated with an object's functional capability relative to other objects. For example, if we want to carve a turkey, then we need an object which has a *sharp* and *long* blade relative to the turkey.
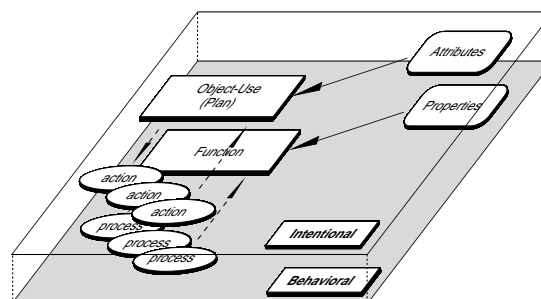

Figure 1: Knowledge structures and their causal relationships are isomorphic for intentional and mechanical representations.

2

An object's attributes direct planning choices in context by constraining applicable functions. At the behavioral level functions organize the processes (as process-state sequences) which effect the object's behavior. Processes are constrained by an object's physical properties and its relationships with other objects.

Intentional-Behavioral Representational Continuity

The relationship that object function plays in integrating intentional and behavioral models is depicted in figure 1. Whether viewed intentionally or behaviorally, the same object function is represented in a given situation. Each point of view provides different inferences about the object, so in Edison the causal relationships are kept distinct. For example, in the game of catch we may want to make inferences about the ball Ptransing (such as why it was thrown), or its Moving (such as how and where it will go), depending on our goals. If we simply merge the representation levels one set of inferences is lost.

It is also important to remember that plans and functions in a given situation both describe the same behavior, but simply at different levels of abstraction. In Edison these relationships are maintained by a structural isomorphism between intentional and behavioral knowledge structures. For example, consider the *plan-action-state* relationship which describes causality at the intentional level. This has a one-to-one correspondence with the *function-process-state* relationship at the behavioral level.

Figure 2: The bi-level representation for a game of catch shows the continuous, albeit separate, inference path between intentional and mechanical abstraction levels.

The bi-level model is designed to describe situations like the game of catch introduced above and depicted in figure 2. The intentional representation is shown on the upper level and the behavioral representation is shown on the lower level. The intentional description has a causal "gap" after the thrower's unGrasp action, whereupon the ball Ptranses to the catcher. The be-

havioral representation overlaps at this point, with the enabling and constraining conditions for the Transport function, and continues to describe the ball's behavioral path (paralleling the Ptrans action) until the Transport function terminates (i.e. the ball's motion ceases). The Transport terminating state is identical to the Catch plan's resulting state (arrival at the intended location). By integrating intentional and behavioral representations this way inferences can be made about object function and behavior not possible with either level alone.

Intentional-Behavioral Inference Continuity

There is a difference in generality between intentional and behavioral reasoning levels which, despite the structural continuity, obviates direct inferences between the two levels. However, because the same object is considered at both inference levels, its functions provide the necessary inference continuity through the associated constraining attributes and properties.

At the intentional level, attributes describe functional capabilities of an object learned through experience, and are specific to particular objects in particular contexts. Knowing the attribute enables high-level inferences about its functional capabilities if the context is reinstated. For example, knowing that a carving-knife was successful in transmitting force for carving a turkey, one might have concluded that the knife is a strong object (w.r.t. the turkey). The *strong* attribute of the knife is a relative term between like property values of the knife and bird, and is only valid for this situation. Other situations requiring *strong* objects, however, might remind the problem-solver of the carving-knife. Attributes thus affect planning, providing grist between context and a problem-solver's associated interpretation. Figure 3 depicts the relationship between different attributes, such as *strong* and *thin*, and how they constrain Knife-Use via the knife functions Pry-Object and Slice.

3
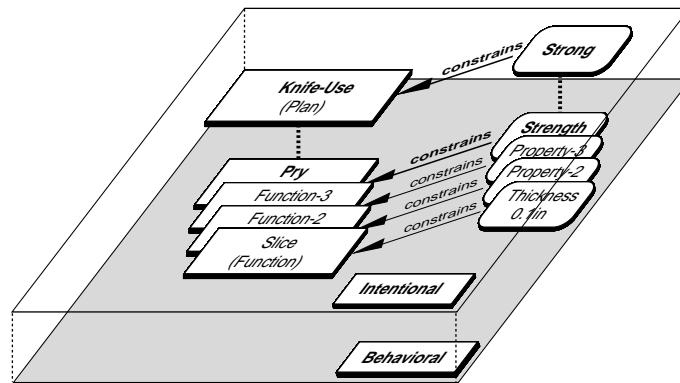
Figure 3: Functional attributes like *strong* are causally related to plan application through the constraints they place on object functionality. Likewise property values constrain the underlying processes. Different attributes will be associated with different situations, and different property values will support the associated functions.

At the behavioral level the *strong* attribute is associated with the knife's value for the breaking-strength[1] property, which directly constrains the Pry-Object function's processes. Knowing the knife's value for breaking-strength guarantees inferences about its capacity to pry. The correspondence between the *strong* attribute and the breaking-strength property value enables inferences between levels. The difference between object functionality based on the attribute, *strong,* and that based on the property, breaking-strength, is that dimensionality (i.e. detail) is lost. If the problem-solver retrieves the knife based on the higher-level functionality (for example during planning), then the dimension of strength is unlikely to be remembered. In Broken Knife this leads to failure. However, the fact that a screwdriver was *strong* for its intended function for tightening screws, leads to an inference that it will be *strong* for other functions as well, such as prying a varnish-can or punching an oil-can for which it is an effective tool. If the knife's behavior is the basis for retrieval (for example during problem-solving experimentation), then dimension is remembered and predictions, inferences, or explanations can be made with confidence.

Attribute-Property Relationships

The ability to make correspondences between attributes and property values is important because of the different inferences that can be made at the functional and behavioral levels, respectively. If the correspondence is made, then the inferences can be compared and behavior modified. Each attribute defines a *range* in a property's quantity space. The two

---

[1]. The equivalent force an object can withstand prior to failure.

attributes, *light* and *heavy*, which describe the weight property of an object, illustrate a many-to-one relationship which is characteristic between attributes and property values. Many attributes are associated with object function through a specific property, such as *strong* to strength, or *long* to length. Attributes can also be described by combinations of properties or other attributes. The attribute *metallic*, for example, is described by the attributes *shiny, smooth, cold* and *hard*.

There are no exact correspondences between an attribute and its associated property value, since attributes are context-dependent. Nevertheless, some comparisons can be made based on how properties and attributes are represented. In Edison property values are defined as (property, dimension, value) triples, and attributes as (property, reference) doubles. These relationships are illustrated in figure 4 for the carving-knife's *strong* attribute in **Broken Knife**.

Figure 4: The attributes *weak* and *strong* illustrate the relative breaking-strengths (shown in pounds) of objects and their context dependency.

The attributes *weak* and *strong* map onto the material property describing breaking-strength. The numbered line segment in figure 4 represents a portion of quantity space describing breaking-strength values, with its central value being the Friction-Force attribute reference. There are two ways that attributes are referenced to property values in quantity-space:

1. To a known reference point or value (e.g. silver-pol-ish can-to-lid friction force).

2. To a boundary value (e.g. the full-open position of a water faucet).

The reference point defines the context which an attribute is directly applicable, and is always found in the situational context. For instance, in **Broken Knife** the reference can be the can's or screwdriver's, breaking-strength, or the friction force between the Canlip and the Lidlip. Either way the knife is comparably weak. The shaded bars in figure 4 represent attributes, and show the variation of the terms *weak* and *strong* with respect to Friction-Force. The shading indicates the generalized relationship between what the man in **Broken Knife** knows about knife and can strength.

The attribute-property value relationship combined with the bi-level structural isomorphism provides a continuous inference path between intentional and behavioral levels of abstraction. If a situation exists in memory where a carving knife has successfully been used as a *strong* object, say to cut meat, then Edison will likely try to use it again when the need for a *strong* object arises (e.g. in **Broken Knife**).

Reasoning About objects in context

The **Open:varnish-can** situation shown in figure 5 illustrates the effect that attributes have at the intentional level. The associated property value and behavioral effects have also been depicted in figure 5, but a detailed description can be found in [Hodges, 1989]. At the intentional level varnishing a chair entails a number of preparatory steps, one of which is to get the varnish onto a paintbrush (a D-Cont goal). In figure 5 this step leads to a plan for opening the varnish can by prying the lid with the tip of a screwdriver.
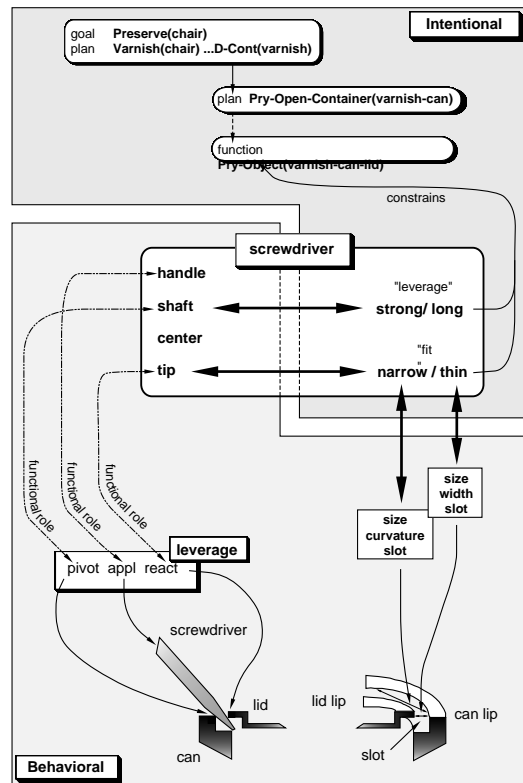
Figure 5: The attributes associated with "leverage" and "fitting" are instrumental in representing how a screwdriver is used to pry a varnish-can lid in **Open:varnish-can** by constraining its application. Attributes map both to object regions (such as handle, shaft and tip) and property values, which constrain functional interactions.

At the functional level the Pry-Object function is governed by two attribute groups, one for leverage and one for fit. The "leverage" requirement states that the screwdriver be *long*, so that sufficient mechanical advantage can be gained to overcome the friction holding the can and lid together, and *strong* so that it won't break under this force. The "fit" requirement states that the tip of the screwdriver must be *thin* and *narrow* compared to the slot between the canlip and lidlip, and constrains the Contact and Magnify processes in Pry-Object at the behavioral level. The leverage box in Figure 5 states that any object with regions of force application, pivot, and force reaction can be used to apply leverage. The screwdriver has these regions bound to its handle, shaft, and tip. In terms of prying these screwdriver regions are the only locations of interest. There are similar regions associated with the can (i.e. the lid, lidlip, can, and canlip). The regions on both objects are also used to define the attribute reference points for prying.

Bi-Level Representation and Situation In-

terpretation

The primary reason for describing object use at varying abstraction levels is to support different object interpretations depending on context. We want a representation model which describes how a screwdriver or knife is used as a utensil in one circumstance, a weapon in another, and a paperweight in a third. Each of these situations calls upon the same object property (weight), but with different required property values. Models that are context independent bar behavioral descriptions from addressing an actor's perspective in the same way that models that are context sensitive bar a functional description from making predictions about behavior. However, even when an object has only been used in a single context (such as using a carving knife for slicing), the attributes which enabled its functionality might enable its use in other contexts.

Knife breaking-strength provides a good example of this. Objects used to cut must be *strong* enough that they do not bend or break before the cut is completed. Of course, knife strength is only meaningful in the dimension of the intended cut. However, a person who naively uses a knife might generalize the extent of strength to all of its dimensions.

Figure 6 illustrates how **Broken Knife** is represented at the situational level. The upper window illustrates the information given. The lower window illustrates a number of situations where simple devices have been used in standard ways, and the attributes which constrain their functionality. The D-Cont goal to get silver polish onto a rag leads to an Open-Container plan. This information is provided to memory as a retrieval cue. The **Open:varnish-can** situation, where a screwdriver is used for prying, is the best functional match but conflicts with the man's P-Comfort goal. The result is that screwdriver-use, and screwdriver-related experiences, are unavailable for planning (with a screwdriver). This is shown with circle-ended dotted lines. The screwdriver attributes which are pertinent to prying are shared (situationally) with other devices (e.g. carving knife in **Slice:turkey**) which can be applied to the Pry-Object function. The **Flip:pancakes** situation is inappropriate because the spatula has attribute *broad,* which conflicts with the *narrow* attribute instrumental to Pry-Object. The carving-knife is also applicable based on availability, since the carving-knife resides in the kitchen setting of **Broken Knife**. The end result is a plan combining the Pry-Object function with the carving-knife object.
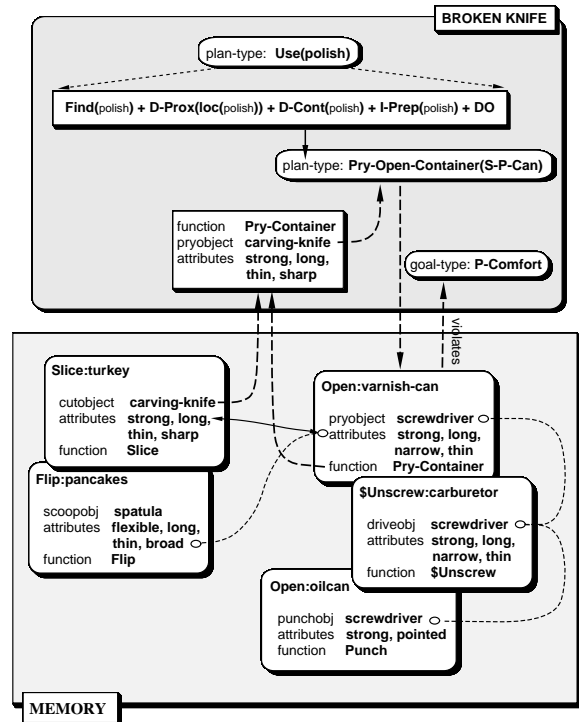

Figure 6: The **Broken Knife** situation illustrates situational interpretation of a carving knife based on its *strong* attribute. The Open-Container(S-P-Can) plan is indexed in memory to situations where objects have been used for opening. **Open:varnish-can** is strongly associated but cannot be applied directly because of a goal conflict. The screwdriver and carving knife share attributes instrumental to prying, so that an alternate Pry-Object plan using the carving knife can be applied to the situation.

a detailed example

The representation for the functional and behavioral inference paths in **Broken Knife** in figure 6 is fleshed out in figure 7. The behavioral description shown in figure 7 represents the process interactions supporting the Pry-Object function with the knife instantiated as the Pry-Object. The representation is shown instantiated with the carving knife after retrieval from memory and combination into the Pry-Object function. The attribute/property-value relationship is shown as it affects the functional/behavioral description under the heading Attribute-Property Mapping. The *fit* requirement affects Pry-Object in two dimensions, so the comparison to size is made in two dimensions. The size values constrain the processes Magnify and Transmit-Force. The darkened two-way arrows between attributes and property values (states) represent a "many-to-one" link. The dashed and darkened two-way arrow between Pry-Object and Open-Container illustrates the inference cross-over between the functional and behavioral level.

Attribute-Property Mapping
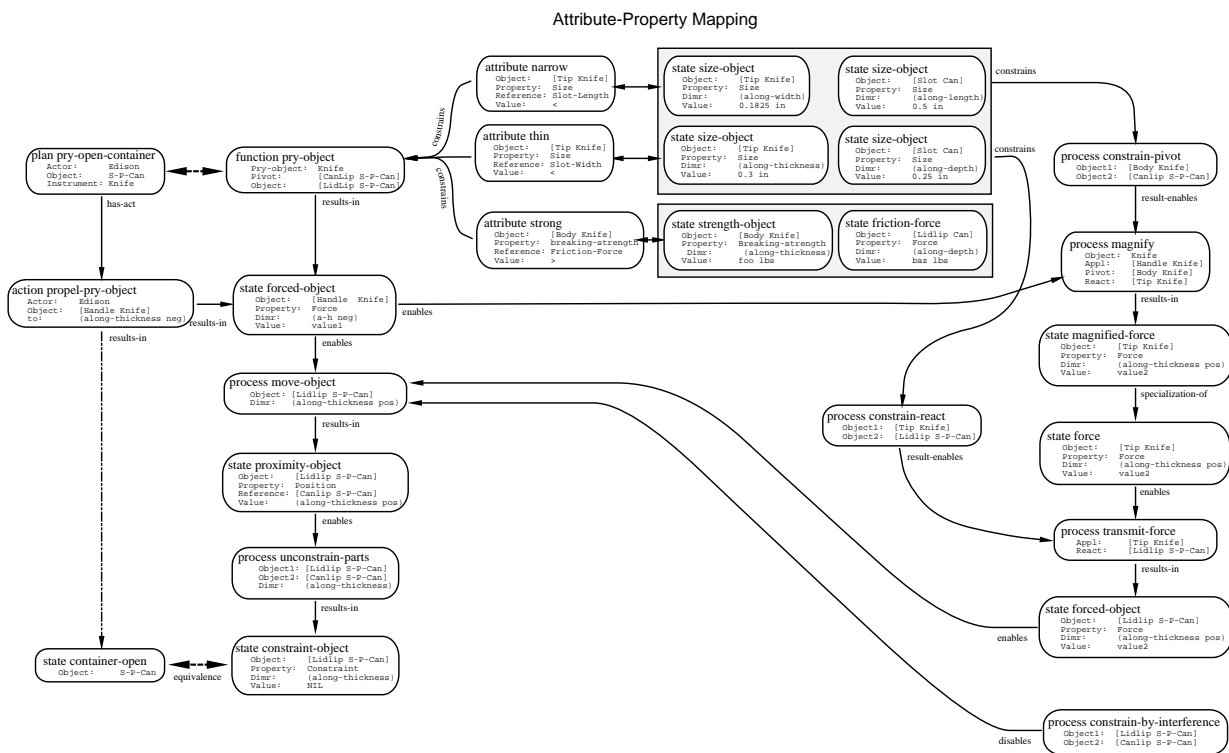
Figure 7: Functional-behavioral representation for **Broken Knife**. The attribute-property relationships constrain the Pry-Object function and the processes which comprise it. Attributes are associated with an object in context so the carving knife *strong*, *narrow* and *thin* attributes are associated with a retrieved situation, **Slice:turkey**. Some of the fillers illustrated (e.g. [Lidlip S-P-Can]) are simplifications of the actual representation.

The planning and interpretation involved in **Broken Knife** and the other situations illustrated in figure 6 are currently being implemented in Robin, a localist spreading-activation model of high-level inferencing [Lange & Dyer, 1989], which uses the Descartes connectionist simulator [Lange, Hodges, Fuenmayor, & Belyaev 1989]. In this implementation there will be equivalent inference paths for other devices which could be used as the Pry-Object filler, such as the candlestick itself. These inferences compete with the use of the knife through the spread of activation. The carving knife inference path will win out and be chosen as the plan for prying open the container, however, since its *strength* and *fit* attributes match the constraints on the Pry-Object role better than the other available objects.

Conclusions

Designing a knowledge representation model which supports the invention process requires an integration between intentional and behavioral object descriptions. The model must address how the environment and people's higher-level goals and intentions affect object choice during problem-solving, and how objects' properties support that functionality at the behavioral level. The bi-level representation used in the Edison model provides the necessary integration and maintains the inferences from each abstraction level. The concept of attributes is introduced, and their relation to property values is discussed with respect to how they affect inferences between intentional and behavioral levels of abstraction.

Acknowledgements

References

DeKleer, J. & Seely-Brown, J.S. (1985): Qualitative Reasoning About Physical Systems, edited by Daniel G. Bobrow, MIT Press, pages 7-84.

Doyle, R.J. (1988): *Hypothesizing Device Mechanisms: Opening Up the Black Box*, MIT Artificial Intelligence Laboratory TR 1047.

Dyer, M., Hodges, J.B., & Flowers, M. (1986): EDISON: Engineering Design Invention System Operating Naively, *Journal of Artificial Intelligence in Engineering, Vol. 1 No. 1*, p. 36-44.

Forbus, K. (1985): Qualitative Reasoning About Physical Systems, edited by Daniel G. Bobrow, MIT Press.

Hodges, J.B. (1989): *Foundations for Creativity: Integrating Functional and Behavioral Object Representations for Problem-Solving*, Ph.D. Dissertation, Computer Science Department, University of California at Los Angeles (forthcoming).

Hodges, J.B., Dyer, M. G., & Flowers, M. (in press): Knowledge Representation for Design Creativity. In D. Sriram and C. Tong, editors, *Artificial Intelligence Approaches to Engineering Design*, Addison-Wesley, (in press).

Lange, T. & Dyer, M. G. (1989): Frame Selection in a Connectionist Model of High-Level Inferencing. *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society (CogSci-89)*, Ann Arbor, MI, August 1989.

Lange, T., Hodges J., Fuenmayor, M., & Belyaev, L. (1989): DESCARTES: Development Environment For Simulating Hybrid Connectionist Architectures. *Proceedings of the Eleventh Annual Conference of the Cognitive Science Society (CogSci-89)*, Ann Arbor, MI, August 1989.

Lehnert, Wendy (1978): The Process of Question Answering. Lawrence Erlbaum Associates. Chapter 10.

Rieger, Chuck (1975): In *An Organization of Knowledge for Problem Solving and Language Comprehension*, Morgan Kaufmann, p. 487-508.

Schank, R. C. & Abelson, R. (1977): *Scripts, plans, goals and understanding*. Hillsdale, NJ: Lawrence Erlbaum Associates.