ASME Winter Annual Meeting

C. R. Liu and A. Requicha and S. Chandrasekar
American Society of Mechanical Engineers
Boston, MA

January 6, 2013

Chapter 12

KNOWLEDGE REPRESENTATION FOR DESIGN CREATIVITY

Jack Hodges
Department of Computer Science
San Francisco State University
1600 Holloway
San Francisco, CA 94312

Margot Flowers, Michael Dyer
Artificial Intelligence Laboratory
3531 Boelter Hall
University of California, Los Angeles
Los Angeles, CA 90024

ABSTRACT

This section briefly reports on the representational strategy used in EDISON, a program currently being designed to (1) invent novel mechanical devices through heuristic strategies of mutation, combination and analogy, and (2) to comprehend descriptions of invented device representations. The representational constructs required to support these tasks include: (a) intentional structures such as goals, plans and settings, which organize relationships between device use and context, (b) physical entities such as regions and materials, (c) behavioral process relationships, such as object motion, connection and deformation, which relate objects to their physical states, (d) function relationships, which relate primitive devices to expected applications, and (e) mechanical dependencies and inferences. Invented and comprehended device representations are indexed and generalized into a memory of design episodes. The organization

Example 1: Swinging Door

> Joe Pizzamaker finds himself repeatedly having to carry pizzas through a doorway in both directions. In one direction he merely pushes the door while in the other he must open the door. At some point of discomfort Joe might say *”surely there must be a better way!”*. He already knows the ease of door use in one direction and so he might have the idea to redesign the door into a swinging door by modifying the existing door to *”close”* in both directions. The problem-solving for this scenario utilizes memory retrieval and combinational strategies.

of such a memory supports the use of cross-contextual reminding and analogy during problem solving.

Keywords: mechanical design, mutation, combination, analogy, creative design, design indexing and generalization

12.1 INTRODUCTION

EDISON is a project created to explore the processes of comprehension [Dyer et al., 1987b] and creativity [Dyer et al., 1986a, Dyer et al., 1987a] in naive mechanics [Dyer and Flowers, 1984]. These tasks require basic research in: physical knowledge representation, memory organization, inference and dependency structures, planning, problem-solving, and learning. The overall approach has been to build a prototype process model and to test the limitations of various comprehension and invention heuristics, along with the representational constructs over which they operate.

The situations we are interested in are those relating to the development of a preliminary design, resulting from an idea or goal and the associated context, rather than design optimization or performance. This approach is exemplified by the following scenario:

Swinging Door is an example of naive invention, a design methodology which uses naive, or commonsense mechanical reasoning to solve problems and generate novel devices. Commonsense reasoning is particularly suited to the representation and processing of Swinging Door for three reasons. The first is *motivation.* Joe is motivated to invent, and his idea originates from a need to reduce his discomfort. The second is *feasibility.* Joe is first interested in whether the idea will work in general, rather than how well it works. His understanding of door use, function, and behavior need only be detailed enough to associate the door with the context of its use, recognize the conditions which will enable and disable its functionality, and predict resulting door behavior. The third is *naive evaluation.* Joe is interested in a simple solution, and evaluates the new door by comparison to other (known) devices.

Commonsense reasoning supports invention in situations such as Swinging Door through the application of experiential knowledge, which requires the integration of intentional and physical knowledge constructs organized into a memory of design episodes. A process model for naive invention is comprised of two major components: a representation and memory which support commonsense reasoning, and a creative component which both recognizes serendipitous situations for change and can follow through with a first-cut design approach.

12.2 SYSTEM ARCHITECTURE

The EDISON system is composed of eleven elements (Figure 12.1). In this figure thin lines with arrows indicate flow of information through the system; thin dotted lines without arrows indicate semantic links between knowledge structures; thick lines indicate knowledge access between knowledge bases (squares) and interpretation subsystems (squares with rounded corners). EDISON accepts three types of natural language input: (a) a device description, (b) a question, or (c) a goal specification and context. A detailed discussion of natural language (NL) comprehension in the EDISON system can be found in [Dyer et al., 1987b].

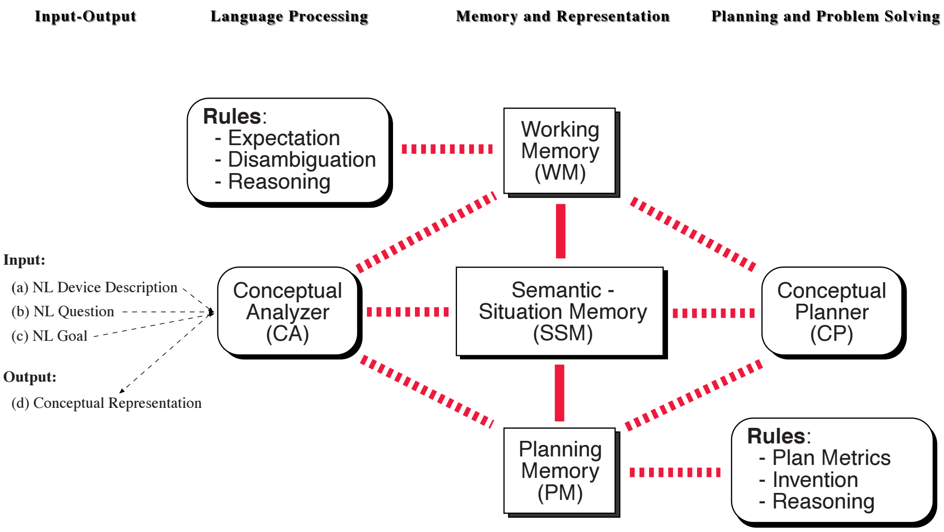

Figure 12.1: EDISON Process Model

Briefly, a goal specification given as input to EDISON is passed to the conceptual analyzer ((1) in Figure 12.1). The CA coordinates the analysis of input text and generates a conceptual representation (c-REP in Figure 12.1) of the goal statement. The c-REP is then utilized by the invention management subsystem to interpret the goal and invent a device.

If the goal is to create a novel device of a given type, then the c-REP is handed directly to the brainstorming component ((10) in Figure 12.1). Brainstorming consists of heuristics which attempt to create novel devices by four general strategies: (1) interpretation of setting and actor intentions to generate design constraints, (2) retrieval and combination of known devices which satisfy, or partially satisfy design constraints, (3) analogy, where some attribute of the device representation is generalized and a device is retrieved (from another episode and/or context) which shares features with the given device at the abstract level, and (4) mutation, where a given device representation is altered along some device property. The door redesign in Swinging Door exemplifies the use of mutation in EDISON.

If the goal specification already includes design constraints, the c-REP is passed first to the problem-solving component of the invention management subsystem ((9) in Figure 12.1). The problem-solver attempts to apply rules and principles of mechanics to satisfy physical constraints. When the problem-solver cannot recall a solution from memory, it calls upon the brainstorming heuristics to improvise a solution to the planning failure.

12.3 NAIVE MECHANICS REPRESENTATION

A naive mechanics representation (NMR) must support comprehension, problem-solving, learning and invention. The general approach of the EDISON representation is to represent physical, relational, behavioral, and functional device attributes as conceptual dependencies, focusing on how device characteristics support device function in the different contexts in which devices are used.

12.4 THE NEED FOR INTENTIONAL KNOWLEDGE IN PROBLEM SOLVING

Consider the doors in Figure 12.2. Most people easily recognize that the door in Figure 12.2(a) simply won't work, and that the door in Figure 12.2(b) cannot be opened in the direction shown. It takes a little longer to realize exactly *why* the normal function of these doors is disabled. This comprehension process often requires that they re-examine how a working door actually functions.

Comprehending the bugs in Figure 12.2 requires that EDISON be able to (1) receive a conceptual representation of a door as input, (2) recognize it as a door (either from a label or by comparing its representation to that of a device in memory), and (3) realize that this particular representation disables a door function. Figures 12.2(a) and 12.2(b) illustrate two ways in which motion can be disabled. In Figure 12.2(a) motion *capability* is disabled from the placement of hinges. In Figure 12.2(b) existing door motion is disabled by a path restraint (doorjam).

We believe that the processes of invention and comprehension share high-level, abstract features across a variety of task domains. In order to detect

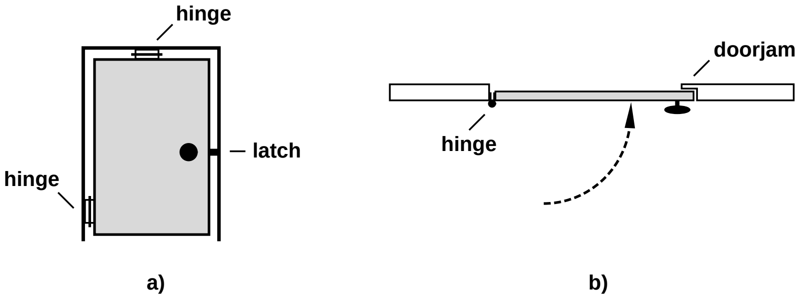

Figure 12.2: Examples of Non-functional doors: a) attribute-based, and b) process-based motion disablements

device errors, EDISON must be able to analyze a device in terms of the goals its use accomplishes. In story understanding and invention domains the relevant goals are those of the characters and include hunger, health, achievement, etc. In the naive mechanics domain, goals involve physical transformations, such as connection and separation. Physical goals are achieved by the *use* of devices. For example, use of the door represented in Figure 12.3 is instrumental to achieving the intentional goal (D-PROX, [Schank and Abelson, 1977]) of moving (PTRANSing) between rooms. Door use, and the function with which a use is associated, thus depends on the context of actor goals.

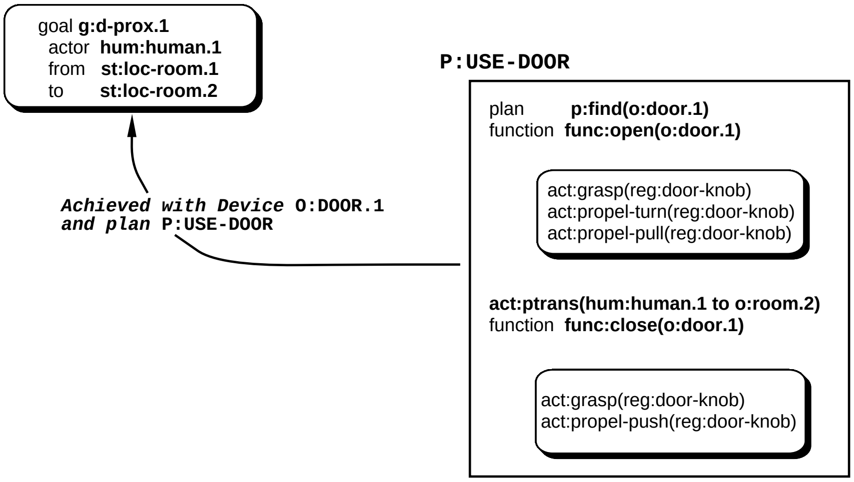

Figure 12.3: Use of intentional representation in device comprehension

The intentional use of objects is represented as a series of events [1], and how those events achieve particular goals. For example, door function (e.g. opening) is initiated by a combination of actions: GRASPing the knob and turning it (a PROPEL resulting in door latch release from the door jam), and pushing the door (a PROPEL resulting in door rotation about its hinges).

In story domains, goals are achieved through the application of plans, and a number of plans may exist which are able to achieve a single goal. Likewise, in naive mechanics, goals are also achieved through the application of abstract plans, but here realized through the operation of physical devices. For example, using the door of Figure 12.3 requires release of a [implied] door latch. Door mobility can be realized by executing the processes used to achieve latch release (e.g. unbolting and untying are acceptable plans for un-restraining parts).

12.5 DEVICE TAXONOMY FOR REPRESENTING FUNCTIONAL COMPREHENSION

A simple door is comprised of many devices (a doorslab, doorway, latch and hinges). Each device is used for different purposes, and functions in different manners. If every device has a unique representational form, EDISON would never be able to distinguish one device from another, nor recognize similarities. On the other hand, if all devices are decomposed to a primitive set of devices, then similarities can easily be traced; supporting both device retrieval and analysis. In the mechanical domain all basic machines [Bramwell and Mostyn, 1984, NAVY, 1971] manifest the principle of mechanical advantage [Weiss, 1983]; and all devices in EDISON decompose to the interaction of simple mechanisms, called machine primitives [Hodges, 1993], which exhibit mechanical advantage.

Notice that one can understand the function of a door and recognize when a door will fail to work (such as those in Figure 12.2) without knowing the exact principles behind leverage. We only need a shallow model of what components do, and not exactly why they do it. In terms of door hinges we need only know that hinges realize mechanical advantage, how their use is enabled and disabled, and how hinges interact with other devices. In EDISON, the representation of device physical and relational properties directly supports the comprehension of (a) physical behavior which the device exhibits, (b) the device function which describes sequenced behavior and produces observable states, and (c) device use and interaction.

Mechanical Comprehension and Representing Behavioral Processes. Each mechanical device interacts with other devices, objects, and the environment. In EDISON mechanical *interactions*, (e.g. motion and connection) are represented as qualitative behavioral processes similar to Forbus' Qualitative Process (QP) theory [Forbus, 1985]. Processes represent causal state sequences relating perturbations to physical state changes, and are used to predict and comprehend

---

[1] Dyer views an event as an action-state pair, or causal primitive, [Dyer, 1983]

device behavior. There are two differences between process representation in EDISON and QP theory.

First, EDISON has no relationships or influences that can be used to explicitly simulate device behavior. Instead, processes are represented as frames: by their behavioral and quantity enablements, and by the states an enabled process results in. A process can be used to predict the resutling state given the proper enabling conditions, or to explain a failed process, but not to simulate spatial behavior. Nor can EDISON processes by used to simulate or predict transient behavior. Second, in EDISON all mechanical behavior can be decomposed to one of five behavioral process primitives: BPP-Motion, BPP-Restrain, BPP-Transform, BPP-Store, or BPP-Deform. Each BPP results in a unique change in state: BPP-Motion to location, BPP-Restrain to restraint, BPP-Transform to force, BPP-Store to energy, and BPP-Deform to size/shape. Moreover, BPPs can be combined to describe arbitrarily complex mechanical behavior, so analysis of mechanical behavior is somewhat simplified.

Despite differences in representational detail, the EDISON methodology is directed at understanding function through context. The approach is best suited to integrating a device with the context of its use; for conceptual or preliminary design, rather than optimization. Clearly both points of view play significant roles in a complete representational model, and one intention of this project has been to maintain predictive continuity with qualitative representation models.

To illustrate how a theory of mechanisms and processes can be useful in creative device interpretation (and generation), let us decompose the representation of door-use that was introduced in Figure 12.3. Early intentional [object] models, e.g. Lehnert [Lehnert, 1978], represented device use in context but didn't associate use and function, or use and behavior. The Lehnert representation could infer what the device was used for, but not how or why. We are interested in how the door actually behaves as a result of an intentional act, and how device behavior is interpreted. Figure 12.4 shows how the open and close *functions* of a door-use plan are represented in EDISON. Device function is represented as the observable input to a device, as perturbations, and by the observable states which the device produces. Device function can be described as a sequence of behavioral processes which causally relate user/device input to the function terminating states. The function terminating state is the state associated with the original purpose for which the device was chosen. Each device may have multiple functions, associated with different properties, mechanisms, or combinations therein, and these may be used together or separately in different contexts. A door has two simple functions: open and close. Each door function consists of an initial action, a motion (or motions), and a resulting position (state).

The *close* function shown in Figure 12.4 describes a simplified (black box) version of the key steps in door closing. The contact between latch (reg:linkage-doorknob) and doorway, sliding and compressing of the spring, and the resulting linkage containment in the doorway have been omitted. The *open* function shown, on the other hand, describes enough detail so that all but the most specific relationships are represented. Decomposing door-use representation to this

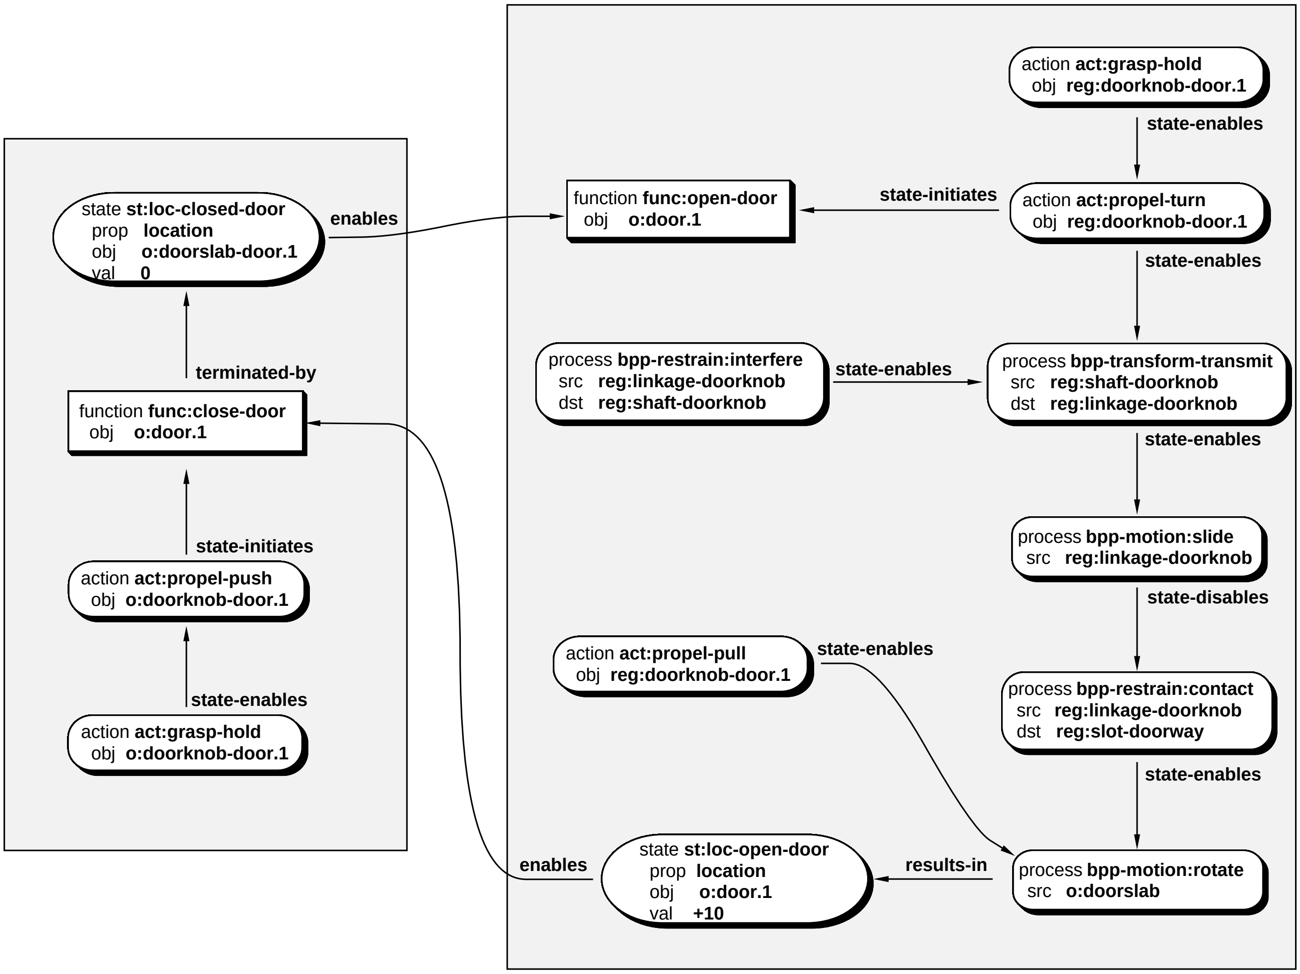

Figure 12.4: Representing door functions: Opening and Closing

level is useful for (a) constraining processing, (b) making inferences and predictions about gross device behavior, (c) integrating the intentional and physical representations, and (d) presenting limiting, or bounding, information for device function. The information obtained from Figure 12.4 enables EDISON to recognize motion of the door toward (direction is not shown in the figure) the doorway as a closing function, and to predict that the door will very likely reach a closed state (processes are scriptal). EDISON can also make the inference that someone, or some thing, was responsible for the motion of the door, and that its closing will satisfy one of their goals (this is only implied in the figure through reference to the actions of actors, and hence to their higher level goals and plans: see [Hodges, 1989, Hodges, 1992] for complete examples and taxonomy).

Although Figure 12.4 shows how processes interact in a device function, nothing specific has been said about what processes do, or how. Bounding the door-use plan enables some inference and prediction for cyclic behavior, however, predicting and explaining door behavior requires some representation at the process level. Figure 12.5 details the process representation level representation in EDISON, and how it supports understanding the *BPP-Restrain* processes in Figure 12.4.

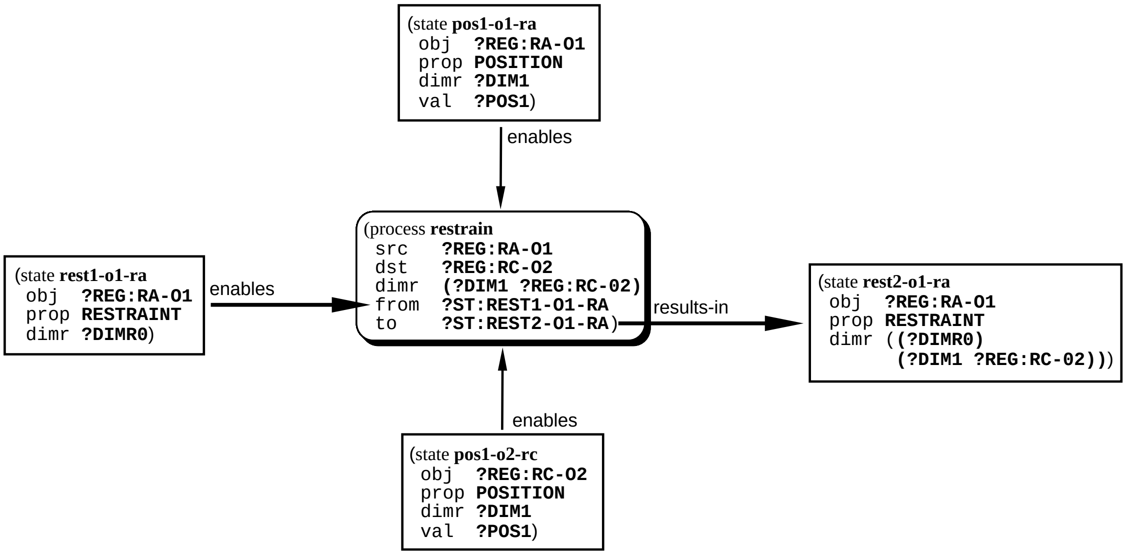

Figure 12.5: Representing BPP-Restrain process in EDISON

Figure 12.5 shows the representational form for EDISON processes and how different BPP-Restrain processes are realized by different role bindings. The representation of processes is very similar to that of Schank's actions [Schank and Abelson, 1977], but there are three differences: (a) processes have no agent, (b) processes are context-free, and (c) processes are more predictive. The rationale for introducing processes over new actions is that processes occur in a physical world which parallels the intentional world. To illustrate, consider an action such as push (propel) as applied by an actor to a ball. The action may result, at the intentional level, in the ball flying through the air (ptrans) from one location (the actor) to another. People generally do not think of the lower level processes of how the impulse is transmitted from the actor to the ball, the storage of energy in the ball, the restraints on the ball, whether or not the ball can move, or what path the ball will take. However, these processes all occur as the object is propelled. Processes have been introduced to maintain the ability to address both representational levels independently. Processes do not have an agent because the forcing function can be supplied by another mechanism (such as a device, or gravity). Processes are context free because they have specific conditions which, when met, result in their expected behavior. These conditions are situation independent, and do not index directly to any intentional knowledge structures. Finally, processes are more predictive because the physical world (process dependencies) is well defined. That is, states resulting from enabled behavioral processes are true physical states.

In EDISON, all mechanical behavior is represented with five behavioral process primitives: *BPP-Motion, BPP-Restrain, BPP-Transform, BPP-Store*, and *BPP-Deform*. The process BPP-Restrain describes object interactions which produce mutual restraint states, thus disabling motion, which is represented with the process BPP-Motion. BPP-Motion and BPP-Restrain are sufficient to enable the transmission and transformation of force between objects, represented with the process BPP-TRANSFORM, the storage of elastic energy in objects, represented with the process BPP-STORE, and the plastic deformation of objects, represented with the process BPP-DEFORM. From Figure 12.5, BPP-Restrain can be seen to require two parts, a dimension and direction, and potentially some medium (e.g. a connector) for holding the objects together. All processes have enablements, and BPP-Restrain requires that the parts be in physical contact to one another. Processes, like actions, cause state changes. Once enabled, BPP-Restrain results in a restraint state on each object, in equal dimensions but opposite directions. BPP-Restrain:Interfere describes object contact in which object motion is disabled along an entire dimension axis.[2] The meaning of *BPP-Restrain:Interfere* can now be interpreted. O:Linkage-DoorKnob and Reg:Slot-DoorWay instantiate the process roles *src* (the source or reference object) and *dst* (the destination object). The object which fills the *src* role determines the process dimension. The dimension (ALONG-RADIUS) refers to the O:Linkage-DoorKnob radial dimension. The process *from* and *to* roles refer to the state change produced by the enabled process. BPP-Restrain processes describe restraint states, which are defined by the process dimension and direction, so the *from* and *to* roles are uninstantiated. The interference between O:Linkage-DoorKnob and Reg:Slot-DoorWay causes a set of restraint

---

[2] As compared to BPP-Restrain:Contact or BPP-Restain:Support, which act on specific directions along a dimension.

states for each: along the O:Linkage-DoorKnob radial dimension.

Two basic process assumptions are made in the EDISON representation approach: (a) parts are free to move unless specifically restrained, and (b) enabled processes will continue unless otherwise acted upon. These assumptions, and other basic knowledge for processes and process interactions, are formulated as process enablements, and take the place of more formalized relations and influences in QP theory, the intention being to make a reasonable accounting for a depth of representation which is beyond the scope of the EDISON project. The assumptions do, however, enable similar types of reasoning, and support limited process prediction, diagnosis, and explanation.

Machine Primitives and Function Comprehension. Behavioral Process Primitives underlie the representation of complex device behavior and device function. Nevertheless, devices, as physical objects, play the central representational role in EDISON, because they index directly to both why the device is used (intentional representation), and how it produces the desired effect (function and behavior representation). The more compact the device representation, the easier it is to associate device use and behavior, and less computational effort will be required to do so. Because we are indexing devices by their use, it is inappropriate to decompose devices to the most primitive known physical mechanisms [Alonso and Finn, 1970]. Instead, we decompose all devices to a set of eleven commonly accepted basic machines [NAVY, 1971], called Machine Primtives: MP-Linkage, MP-Lever, MP-Wheel-Axle, MP-Gear, MP-Pulley, MP-Bearing, MP-Spring, MP-Container, MP-Plane, MP-Blade, and MP-Screw. Machine primitives represent simple devices which have a single expected function. For example, MP-Linkage is associated with objects which are used to extend force over some distance by transmission. The objects which can be involved in this function are those which can transmit force in at least one dimension and direction. The roles of the primitive are those regions where applied forces are applied, called *appl*, and reacted, called *react*. All mechanical devices can be decomposed to combinations of machine primitives, and by understanding them EDISON has the capacity to understand, reason about, and generate, more complex devices.

Figure 12.6(a) presents the EDISON representation for BPP-Lever, which is instantiated by simple lever-objects. A lever-object is a linkage-object with the addition of a pivot location. Thus MP-Lever specializes MP-Linkage with the addition of a *pivot* role (i.e., MP-Lever has three roles; appl, pivot, and react). The pivot location, as with the locations associated with the appl and react MP roles, respresents a generalized location directly associated with device function. Generalized locations are represented with a physical characteristic called a *region* [Dyer et al., 1986b]. Whereas MP-Linkage is used to transmit or translate forces and velocities, the function of MP-Lever (Figure 12.6(b)) is to magnify force or speed; both of which enable specializations of the process BPP-Transform.[3] MP-Lever is realized in different ways depending on how the

---

[3] All EDISON machine primitives, except MP-Spring and MP-Container, enable BPP-Transform. MP-Spring enables BPP-Store, and MP-Container enables BPP-Restrain.

remaining MP-Lever roles are instantiated: (a) type of applied input, (b) relative locations (represented as *relations*) of the input, fulcrum, and reaction regions, and (c) relative magnitudes of input and reaction (whether velocity or force). The resulting state change is effected through the representation of BPP-Transform, and BPP-Transform:Magnify in particular. The bindings for *door-hinge* in Figure 12.6(c) are shown as they apply to the function representational form. The doorhinge is really two simple lever-objects *pinned* together. However, the effect of the MP-Lever instantiated by O:Plate1-DH is nullified because BPP-Motion enables MP-Lever use, and the doorway is grounded.

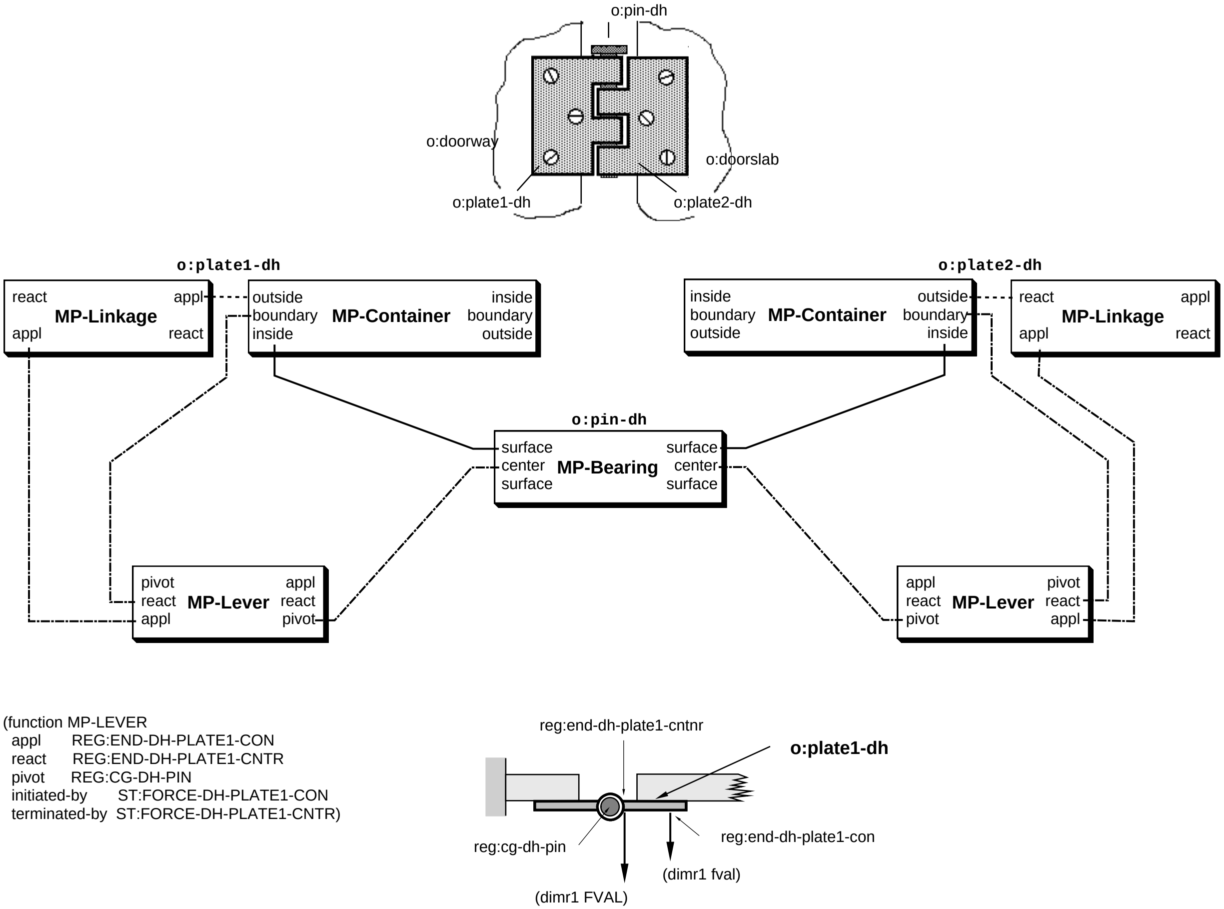

```
(function MP-LEVER
  appl      REG:END-DH-PLATE1-CON
  react     REG:END-DH-PLATE1-CNTR
  pivot     REG:CG-DH-PIN
  initiated-by      ST:FORCE-DH-PLATE1-CON
  terminated-by  ST:FORCE-DH-PLATE1-CNTR)
```

Figure 12.6: Doorhinge MP-Lever representation. Simple mechanisms can be used to reason about mechanical interactions and device Use: (a) MP-Lever representation, (b) MP-Lever function,(c) Doorhinge Representation

The significance of physical and relational characteristics is that all device-related knowledge structures index directly to device use, device function, or to a process which the characteristic enables. The representation for a device thus indexes into both intentional (e.g plans) and physical (e.g processes) knowledge types. EDISON will always be able to say which device characteristic is responsible for a particular use, or why an intended use failed. For example, regions describe generalized locations on a device, and instantiate the roles of a machine primitive. By representing only those device regions which directly affect a particular device function, the complexity of spatial descriptions is reduced, thereby aiding in differentiating uses and processes. How do we recognize the futility of trying to cut a metal rod with a rolling pin? People recognize that cutting requires an object with a sharp edge, where edge is a region, and sharp is called a *property attribute*. An attribute describes a simple comparison between the property values of objects used in a particular context. Both the edge and its sharpness are associated with the cutting of objects by the machine primitive *MP-Blade* and the process *BPP-Deform*. A rolling pin simply doesn't have a sharp edge, so most people do not consider it in the light of cutting. The door-hinge fulcrum (instantiated by the object O:Pin-DH) is a pivot region which allows the hinge plates to rotate relative to one another. The fulcrum location and implementation are actually unimportant in relation to the knowledge that either plate can carry the door weight.

The combination of process and device knowledge, with primitives, enables a broad view of physical interactions. EDISON can now make predictions and explanations of device behavior given only limited knowledge. For example, when a door is mentioned in text we *expect* some reference to dooor open or door close. Given an event in either the open or close function of door-use, we can *predict* the processes, and events within the processes, which are temporally local to the known event. EDISON can also *explain* behavior which deviates from that expected either at the device or process level. This kind of behavioral, and functional, analysis is used during comprehension of text describing mechanical situations. Consider the inferences required to understand the text of Broken Foot(figure 12.7).

The inferences required in building a conceptual representation of Broken Foot utilize knowledge in the door-closing function not explicitly mentioned in the text. The lexical entry for *”door”* sets up expectations for the functions associated with door use[Dyer et al., 1987b]. The phrase *”would have...but”* indicates a failure to achieve a given state, followed by an explanation. An explanation for the failure leads to a consideration of how the door-closing function is disabled. Closing is disabled either by restraining door motion or by eliminating the propelling force (see Figure 12.4). The conjunction *”but”* is a causal indicator linking foot placement with the disabled closing function. *”Would have”* and *”closed”* enable the inference that the door was being closed. Foot placement is thus assumed to restrain door motion, since motion once enabled can only be disabled by direct behavioral interaction. Thus the foot must be positioned somewhere along the door's path of motion.

The integration of process and machine knowledge from the last two sections

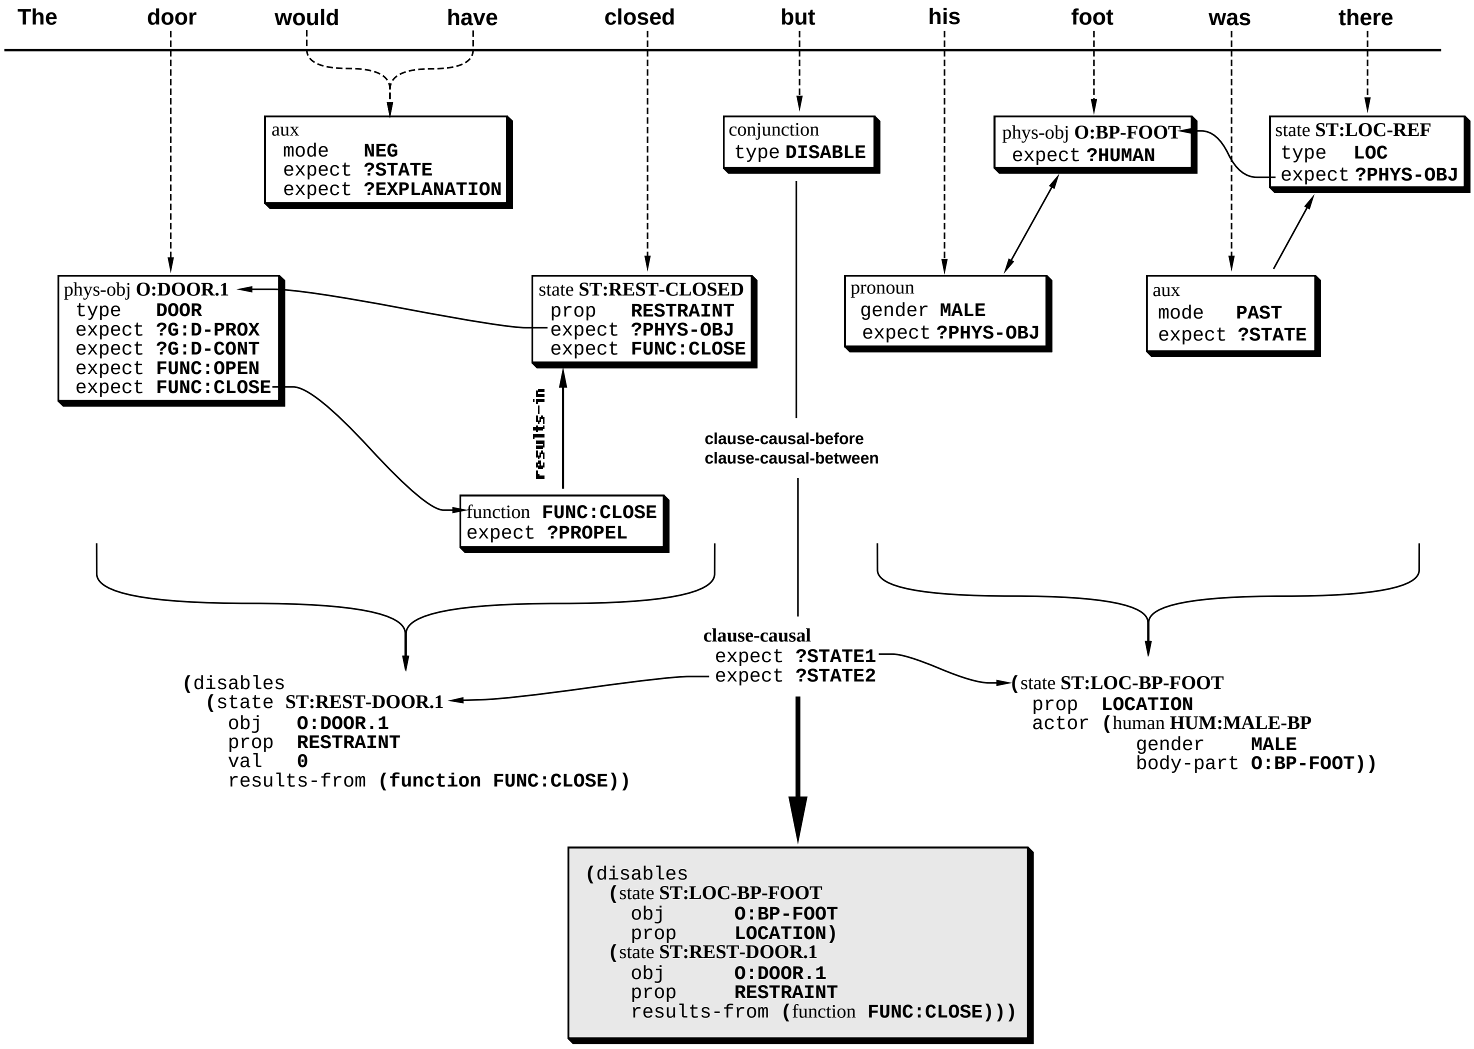

Figure 12.7: Comprehending Broken Foot using Process Theory

enables an explanation to be constructed for the buggy doors in Figure 12.2. Behavioral process primitives and machine primitives are instantiated to describe the configurations depicted. The knowledge captured by these representations can be formulated as rules such as H1 and C1-C4 below:

> **H1:** If object O1 is a hinge, then the plates of O1 can rotate relative to each other about the long axis of O1.
>
> **C1:** If two objects O1 and O2 are connected along direction D, then if one moves in D the other moves in D.
>
> **C2:** If two objects O1 and O2 are in contact, then if either moves toward the other the other will also move.
>
> **C3:** If two objects O1 and O2 are connected in multiple points, then the global restraint on the objects is the union of restraints along each dimension.
>
> **C4:** If two objects O1 and O2 are connected in more than one location but do not share a common axis, then the connection is rigid.

H1 is a simple statement that hinges transmit forces in all dimensions except about their longitudinal axis. That is, relative rotation between the plates is the only motion that a hinge is capable of. H1 is loaded onto a rule agenda when a hinge is recognized and retrieved from memory. When the agenda is cycled the rule is applied to knowledge in working memory. C1-C4 can all be derived from the simple relationship that two objects connected along a dimension share the restraints of the connection type, minimally along that dimension. Process rules are applied in the same manner as device rules. The result of applying these rules to the devices in Figure 12.2 is a global (device) restraint state which disables motion.

Device Representation and Episodic Comprehension. Naive mechanics reasoning in EDISON is experience based. The potential for making interesting device comparisons and combinations is directly related to (a) the amount of experience, and (b) the number of possible connections between representational constructs. However, representational complexity, which is directly related to the number of possible connections, is inversely related to comprehension, and to the ease of comparison. EDISON organizes device knowledge behaviorally, functionally and intentionally to account for this contrast. Behaviorally, device characteristics, represented as states, index to behavioral processes. Functionally, device behavioral sequences index to the observed behavior associated with device use. Intentionally, device functions *must* index to the context which motivates device use. The relatively small number of machine and behavioral primitives, combined with the use/functional nature of the model, provide an environment where comprehension and diverse comparisons can coexist.

People tend to learn about, remember, and retrieve devices in terms of *attributes* associated with a situation. A device attribute is a comparison between a device property value and its boundary values, or with property values of other devices. For example, we may consider a faucet *leaky* if it won't close all the way. The comparative property is position, and the bounding values are open and closed. Were we to make the same kind of comparison, only w.r.t. the open position, then we might say that the faucet is clogged or restricted. The attribute thus tells us the point of view whereby device function is evaluated. Property attributes can index to any contextual component, and so device use can be interpreted in context. Also, because the physical property is directly associated with a behavioral process, EDISON can infer which function the situational context refers to.

Design episodes in EDISON are comprised of four components: (1) an environmental context, represented as states, (2) a problem solver's goals, motivated by the environmental context, (3) the problem solver's planning, related to the goal, which includes the devices applied, and (4) the observable states resulting from the executed plans. Each component adds a contextual element to the episode and serves as a point of view for episodic interpretation. To illustrate this concept consider the doors in Figure12.8. One door may be used in a bank vault as security, while the other door is used in a flood for flotation.

The environmental state of flooding motivates a not-drown (G:Preserve-Health) goal. One way to avoid drowning is to stay-afloat, and staying afloat is associated with devices which float, and to materials capable of floating. Because the door is wooden, it may well be used to stay afloat. In contrast, a $banking script [4] builds expectations for money containment (G:D-Cont). This goal suggests a default (prototypical) door use with emphasis on material strength (for security), which is also met with a material (metal) property.

12.6 NAIVE INVENTION IN EDISON

In EDISON the point of view is taken that the creative process requires the ability to (a) address and interpret a situation from multiple perspectives, (b) select an interpretation among many, and (c) visualize the environmental effect of the interpretation. If a problem-solver resolves each new problem by simply recalling a past solution, then inventiveness should diminish as the number of devices grows. However, with human inventors the acquisition of a novel device serves as a platform for coming up with more devices. Debono found, in his research with children [DeBono, 1980], extensive use of analogy and combination when the task given to the children was to create novel devices. Making device

---

[4] The use of $ follows the convention used by Schank and Abelson [Schank and Abelson, 1977] for scripts.

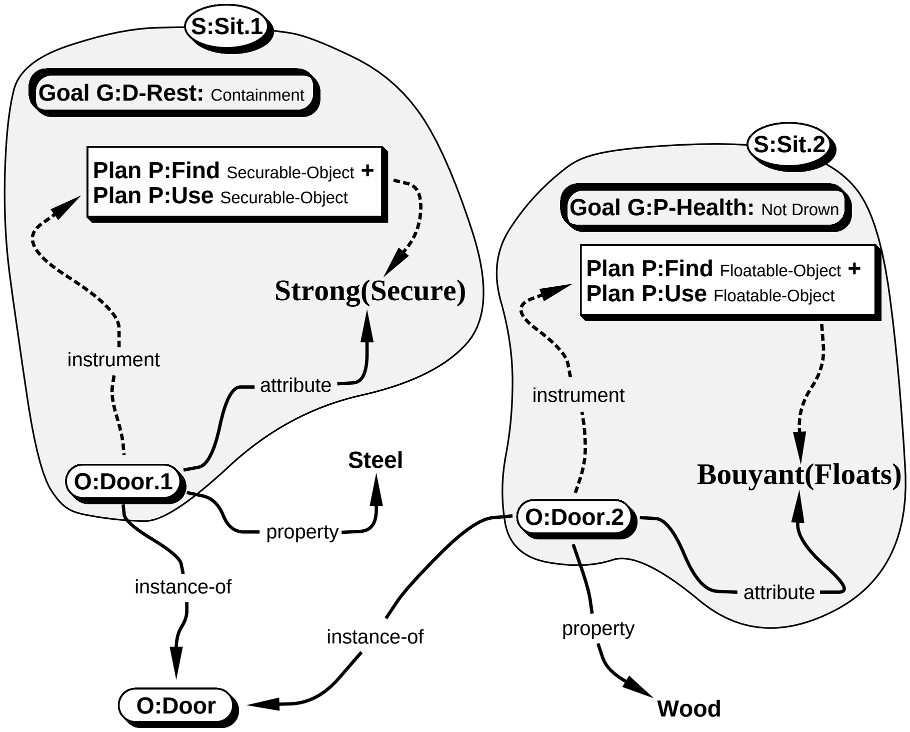

Figure 12.8: Contextual Determination of Door Use in Flooding and Bank Vault

comparisons this way is supportive of the idea that growth in episodic memory increases the potential of inventiveness rather than diminishing it.

The representation presented uses design episodes to support the ability to make and comprehend comparisons. The creative utilization of design episodes introduces four issues important to the study of naive invention: (1) the motivation for invention, (2) preliminary design and invention, (3) methods for generating new designs, and (4) assessing the ingenuity and worth of new devices.

Failure motivates invention. The quote *"necessity is the mother of invention"* has popularized a basic tenet in recognizing the potential for invention: goals are significant motivators for change. Goal successes rarely lead to inventions, but goal failures point out planning limitations, conflict, and/or competition between goals. These are good indicators that an invention process will be useful. When invention is initiated, past design failures can be reviewed in the light of new knowledge, and may result in a successful design. Likewise knowledge generated from remindings [5] may result in more goals being achieved by a single design.

Invention and conceptual designs. Invention is customarily associated with the early, conceptual, stages of design; inventors identify factors which are instrumental to a successful design, and build prototypes to demonstrate the concept. EDISON is a model of conceptual design. We seek contextual interpretations which lead to the understanding, and development, of design constraints. The invention itself results from the interaction of constraint and relaxation based methods applied to the design constraints. The device representation is fundamental for interpreting context and developing constraints, and thus fits into the creative strategy of this model.

Design generation. Devices can be generated by the application of three simple invention heuristics, (1) combining known devices, each of which partially satisfy a design constraint, (2) analogically mapping a known device (and source domain) to a new device and target domain, and (3) mutating known devices. Mindless generation of devices, however, is anything but creative. Each invention heuristic has its place, and the inventor knows when best to apply them. An example illustrating an appropriate use of analogy for invention is the door redesign in Swinging Door. Once Joe has decided to make a door which opens both ways he runs into the problem that standard door hinges only open in one direction. If Joe analogizes swinging horizontally to swinging in any dimension he can be *reminded* of a clock radio with numbers on flash cards which flap as their axis is turned. The cards use an axial hinge to enable swinging in both directions. Making the comparison between the two doors Joe can now consider whether the axial hinge will work on a door in the vertical dimension.

Design ingenuity and uselessness. Two kinds of knowledge constrain EDISON's processing. First, physical knowledge constrains the generation of novel but useless devices. A good example is the use of physical orientations between objects. In Figure 12.2 the door wouldn't secure were the linkage and slot not

---

[5] Remindings are spontaneous similarity-based retrievals, see Schank [Schank, 1982].

coaxial, a state which would render the device useless for door restraint. Second, the interaction of planning *metrics* constrains the design process.

Many problems arise in designing a door, including the selection of hinge type and placement, latch type and placement, even the material out of which the door is made. Each of these details is significant in arriving at an overall door design. Achieving the intended use, however, will generally have priority over satisfying more detailed design constraints. In EDISON new designs are created using simple heuristics such as mutation and combination. Similarly, the design process is both constrained and evaluated using invention planning *metrics*. EDISON has six invention metrics: (1) functional cost, (2) elegance (physical and functional simplicity), (3) utility, (4) performance, (5) novelty, and (6) efficiency. Invention metrics oversee the invention process and compete for priority in the design. A device is considered ingenious if multiple invention metrics are satisfied in its design.

In some cases only one planning metric may be activated, resulting in a natural focus. One such case arises in improvisation, in which the only metric involved is utility (i.e. will the device work). In such cases any invention heuristic resulting in a design contradicting the desired use will be avoided. In other cases competition between metrics forces the design process. Swinging Door is a good example of competition between planning metrics. Joe has a goal to get Pizzas from one room to the next; this involves utility. Simultaneously, Joe has a personal goal to maximize personal comfort; this involves ease and simplicity. The two goals conflict, the result of which is a conflict between the design metrics. Depending on the strength of Joe's goals the door design will vary.

12.7 FUTURE WORK IN EDISON

The EDISON representation is designed to support the creative process, but the creative capacity suggested by this model leaves many issues unanswered. Some of these issues have been addressed to some extent but remain unimplemented, others are just too difficult to consider at the present stage of model development. We present here a few interesting concepts which we would like to pursue further.

Throwing in the towel. Designers and inventors alike tend to get an idea and milk it to death, oftentimes ignoring simple and more elegant solutions. The issue of competing models, the importance which a creator gives to a partially-successful invention, and what the creator does with a partial invention when the evidence points against it (in terms of processing) is interesting. The same comments can be made of device interpretation. Often times there may be many mechanisms in a device, and understanding one may be requisite to understanding another. Perhaps some processing stack exists and invention (and comprehension) processes can be shuttled to and from the stack, depending on the context and available information.

Interpreting failure in an inventive memory. We have seen, above, that failures motivate invention scenarios. But what is the role of failure in memory?

Schank [Schank, 1982] has argued that failures are important because learning occurs at failure points. Dyer [Dyer, 1983] has shown that plan failures represented at an abstract level serve as an indexing structure to cross-contextual memories. If every trivially bad design is stored in EDISON's episodic memory, then problem-solving efficiency may suffer, as a result of recalling bad designs. However, if failures are never stored in memory, then EDISON will be doomed to repeat its mistakes. Therefore, along with design successes EDISON must store design failures. The generalization of specific instances, whether success or failure, leads to abstract experiences in memory. Situations which are not generalized remain salient as episodes. The overall effect is that EDISON will later be able to apply a bad design to resolve a different problem, or will be able to re-explore the bad design in lieu of new knowledge, in the same ways that successful designs are used.

Interference and invention. A conflict exists between the use of reminded experiences during invention and the interference [**?**] of reminded experiences upon invention. Creative people use their broad experience as a platform for creating new designs *because* their experience can be applied across domain boundaries when the context is similar. In this respect remindings aid invention. During invention, however, continual reminding of old solutions can detract from being creative. The inventor must be able to override reminded memory *in order* to invent. Inventors don't seem to block remindings but, rather, make decisions as to what knowledge is pertinent. The EDISON model is being designed to address this fundamental issue in design creativity. The current approach is to consider the active goals being processed. When an active goal is associated with device use, remindings are not used as direct solutions. Thus if EDISON is trying to invent a better bicycle, a bicycle may be retrieved for comparison purposes, or to generate new indices into memory, but won't be used as a solution. Nominally, if the bicycle is the only item retrieved, then mutation of some bicycle attribute would be applied. When remindings are associated with non-primary design goals direct use is acceptable. One example are the screws used to connect a hinge to a door/doorway. Why reinvent a screw unless the mode of connection is of interest. We hope that this initial approach will lead to further insight into the problem of interference in creative design.

12.8 CONCLUSIONS

Naive mechanics comprehension and invention can be modeled in terms of symbolic manipulations on representational constructs. Invention and creative design can be motivated from an interpretation of situational context in terms of actor goals and plans. Interpreting design episodes results in the development of conceptual design constraints. Invention heuristics then enable us to combine, analogize and/or mutate representations so as to achieve constraint driven goals; resulting in a preliminary design. The representational approach stresses the interaction of intentional and physical knowledge structures in memory, as applied to the creative process. The resulting designs are indexed into memory by features common across domains, increasing the amount of knowledge potentially applicable to future design goal achievement.

The model emphasizes the role of episodic memory in creativity, and lacks the same ability to simulate device behavior as some qualitative, and all quantitative, representation models. The difference lies in the approach. EDISON is directed at reasoning about multiple device uses, and emphasizes a simple representation for behavior and function through the introduction of knowledge primitives associated with each. This limits the ability of EDISON to simulate device behavior, but enables us to describe entire problem-solving scenarios and to express similarities between devices used for different purposes, in different contexts. We believe that this representational outlook is a necessary component to an overall representational scheme which can support creativity.

12.9 ACKNOWLEDGMENTS

The research reported in this section was funded in part by a grant from the Office of Naval Research (contract no. N00014-86-0615)

Bibliography

[Alonso and Finn, 1970] Alonso, M. and Finn, E. J. (1970). *Physics*. Addison-Wesley, Reading, MA.

[Bramwell and Mostyn, 1984] Bramwell, M. and Mostyn, D. (1984). *How Things Work*. Usborne Publishing.

[DeBono, 1980] DeBono, E. (1980). *Children Solve Problems*. Penguin.

[Dyer, 1983] Dyer, M. G. (1983). *In-Depth Understanding*. MIT Press, Cambridge, Massachusetts.

[Dyer and Flowers, 1984] Dyer, M. G. and Flowers, M. (1984). Automating design invention. In *Proceedings of Autofact 6*, Anaheim, CA.

[Dyer et al., 1986a] Dyer, M. G., Flowers, M., and Hodges, J. (1986a). Edison: An engineering design invention system operating naively. *International Journal of Artificial Intelligence in Engineering*, 1(1):36–44.

[Dyer et al., 1986b] Dyer, M. G., Flowers, M., and Hodges, J. (1986b). Edison: An engineering design invention system operating naively. In Sriram, D. and Adey, R., editors, *Applications of Artificial Intelligence in Engineering Problems*, volume 1, pages 327–342, Southampton, United Kingdom. 1st International Conference, Springer Verlag.

[Dyer et al., 1987a] Dyer, M. G., Flowers, M., and Hodges, J. (1987a). Naive mechanics comprehension and invention in edison. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence*, volume 2, pages 696–699. Morgan-Kaufman.

[Dyer et al., 1987b] Dyer, M. G., Hodges, J., and Flowers, M. (1987b). Computer comprehension of mechanical device descriptions. Technical report, University of California at Los Angeles. UCLA-AI-87-9.

[Forbus, 1985] Forbus, K. D. (1985). Qualitative process theory. In Hobbs, J. R. and Moore, R. C., editors, *Formal Theories of the Commonsense World*, chapter 5, pages 185–226. Ablex Publishing Corporation.

[Hodges, 1989] Hodges, J. (1989). The ucla cognet: A tutorial introduction. Technical Report UCLA-COGSCI-89-1, Cognitive Science Program, University of California, Los Angeles.

[Hodges, 1992] Hodges, J. (1992). Naive mechanics: A computational model of device use and function in design improvisation. *IEEE Expert*, 7(1):14–27. cover article.

[Hodges, 1993] Hodges, J. (1993). *Naive Mechanics: A Computational Model for Representing and Reasoning about Mechanical Devices.* PhD thesis, University of California at Los Angeles.

[Lehnert, 1978] Lehnert, W. G. (1978). Representing physical objects in memory. Research Report 131, Department of Computer Science, Yale University.

[NAVY, 1971] NAVY, U. (1971). *Basic Machines and How They Work.* Dover Publications, Inc.

[Schank, 1982] Schank, R. (1982). *Dynamic Memory: A Theory of Learning in Computers and People.* Cambridge University Press, New York.

[Schank and Abelson, 1977] Schank, R. and Abelson, R. (1977). *Scripts, Plans, Goals, and Understanding.* The Artificial Intelligence Series. Lawrence Erlbaum, Hillsdale, NJ.

[Weiss, 1983] Weiss, H. (1983). *Machines and How They Work.* Thomas Y. Crowell.