

DTIC FILE COPY

④

AD-A200 696

Knowledge Representation for Design Creativity

Jack Hodges
Margot Flowers
Michael G. Dyer

January 1988

Technical Report UCLA-AI-88-4

Artificial Intelligence Laboratory

Computer Science Department
University of California, Los Angeles
Los Angeles, California 90024

DISTRIBUTION STATEMENT A
Approved for public release
Distribution Unlimited

U
UC
UCLA
UCLA I
UCLA I LAB
AI LAB
LAB
AB
B

DTIC
SELECTED
OCT 13 1988
D

88 1012 059

4

Knowledge Representation for Design Creativity

**Jack Hodges
Margot Flowers
Michael G. Dyer**

January 1988

Technical Report UCLA-AI-88-4

1 3 1988
D

DISTRIBUTION STATEMENT
Approved for public release
Distribution is unlimited

Knowledge Representation for Design Creativity *†‡

Jack Hodges
Margot Flowers
Michael Dyer

Artificial Intelligence Laboratory
3531 Boelter Hall
University of California, Los Angeles
Los Angeles, CA 90024

Accession No.	
NTIS	DTIC
Unannounced	Justified
By <i>pr HP</i>	
Distribution	
Availability	
Dist	Avail
A-1	

Abstract

This paper 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) functional relationships, such as connection and separation, which relate objects to their physical behavior and (d) mechanical dependencies and inferences. Invented and comprehended device representations are indexed and generalized into a memory of design episodes. The organization of such a memory supports the use of cross-contextual reminding and analogy during problem solving.



1 Introduction

EDISON is a project created to explore the processes of comprehension [1] and creativity [2,3] in naive mechanics [4]. These tasks require basic research in:

*The research reported in this paper was funded in part by a grant from the Office of Naval Research (contract no. N00014-86-0815).

†Updated version of a paper presented to the Winter Annual Meeting of the Society of Mechanical Engineers in Boston, MA, December 1987.

‡To appear in "Artificial Intelligent Approaches to Engineering Design", editors Chris Tong and Duv Sriram, Addison-Wesley, in Press.

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:

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 interesting retrieval and combinational strategies.

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 and function need only be detailed enough to associate the door to 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.

2 System Architecture

The EDISON system is composed of eleven elements (Figure 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 [1].

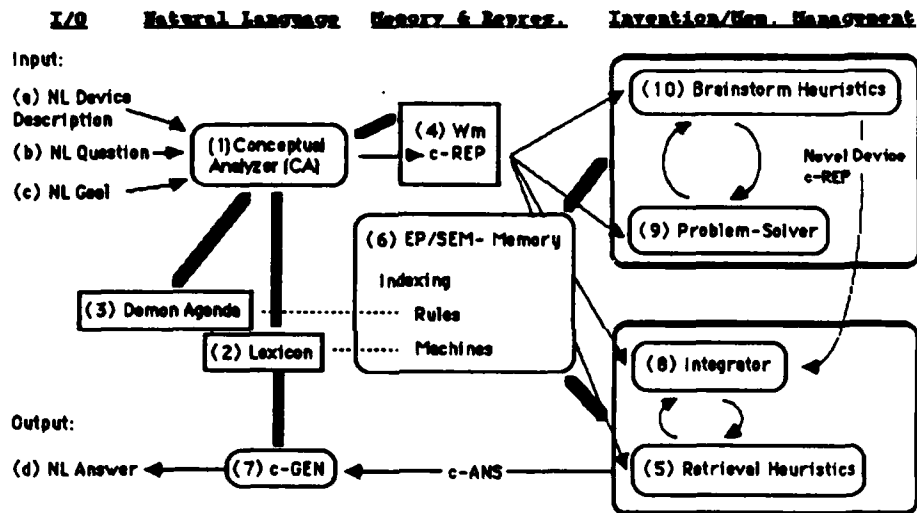


Figure 1: EDISON Process Model

Briefly, a goal specification given as input to EDISON is passed to the conceptual analyzer ((1) in Figure 1). The CA coordinates the analysis of input text and generates a conceptual representation (c-REP in Figure 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 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 sat-

isfy, 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 attribute. 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 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.

3 Naive Mechanics Representation

A naive mechanics representation (NMR) must support comprehension, problem-solving, learning and invention. The EDISON representation is not finalized, but the general approach is to represent physical, relational and functional device attributes as conceptual dependencies, focusing on how device functional characteristics might support the different contexts of device use.

3.1 The Need For Intentional Knowledge in Problem Solving

Consider the doors in Figure 2. Most people easily recognize that the door in Figure 2a simply won't work, and that the door in Figure 2b 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 2 requires that EDISON be able to (1) receive a conceptual representation of a door, (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 2a and 2b illustrate two ways in which motion can be disabled. In Figure 2a motion *capability* is disabled from the placement of hinges. In Figure 2b existing door motion is disabled by a path constraint (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 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

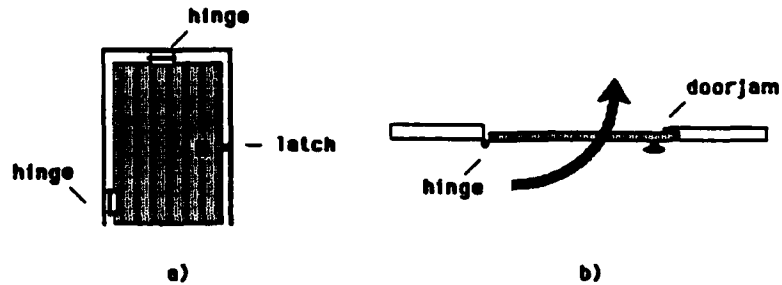


Figure 2: Examples of non-functional doors: a) attribute-based, and b) process-based motion disablements

example, use of the door represented in Figure 3 is instrumental to achieving the intentional goal (D-PROX, [5]) of moving (PTRANSing) between rooms. Door use, and the function with which a use is associated, thus depends on the context of actor goals.

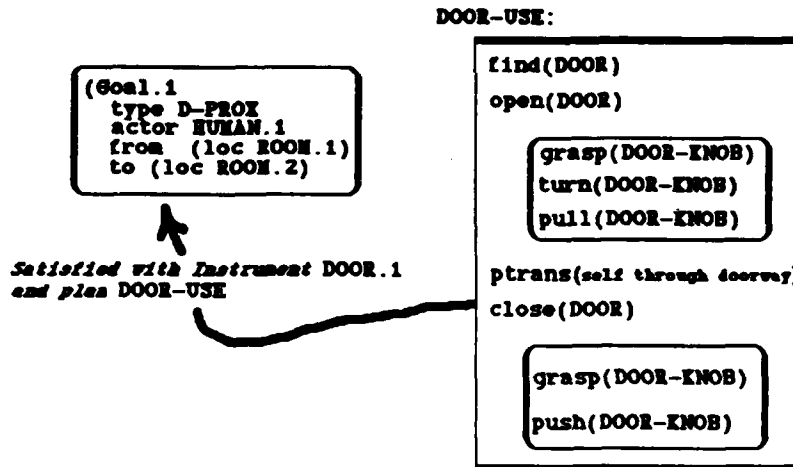


Figure 3: Use of intentional representation in device comprehension

The intentional use of objects is represented as a series of events¹, 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

¹Dyer views an event as an action-state pair, or causal primitive, [6]

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 plans², but here realized through the operation of physical devices. For example, using the door of Figure 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-constraining parts).

3.2 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 analysis and retrieval. In the mechanical domain all basic machines [8,9] manifest the principle of mechanical advantage [10]; and all devices in EDISON decompose to the interaction of simple mechanisms [3] 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 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 either (a) the comprehension of physical behavior which the device exhibits, or (b) device use and interaction.

3.2.1 Functional Comprehension and Representing Physical 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 mechanical processes similar to Forbus' Qualitative Process (QP) theory [7]. Processes represent causal event sequences relating actor goals to physical states, and are used to predict and comprehend device behavior. The difference between process representation in EDISON and QP theory is that EDISON has no relationships or influences that can be used to simulate device behavior. Instead simple process rules are used to effect and

²The application of a plan is equivalent to what Forbus refers to as a history, see [7]

check device behavior under a process. The effect of this difference is that QP theory is better suited to simulating device function so as to recognize and understand new functions. The EDISON methodology, on the other hand, is directed at understanding function through context, and is better suited to integrating a device with the context of its use. 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 the QP model.

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 3. Early intentional [object] models, e.g. Lehnert [11], represented device use in context but didn't associate use and function. The Lehnert representation could infer what the device was used for, but not how or why. We are also interested in how the door actually behaves as a result of an intentional act, and how device behavior is interpreted. Figure 4 shows how the open and close functions of door-use are decomposed in EDISON. Functions are a series of linked processes which relate user/device input to the purpose (end state) 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: opening and closing. Each door function consists of an initial action, a motion (or motions), and a resulting position (event).

The *close* function shown in Figure 4 describes a simplified version of the key steps in door closing. The contact between latch (linkage) 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 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 4 enables EDISON to recognize motion of the door toward the doorway as a closing function, and to predict that the door will very likely reach a closed state. EDISON can also make the inference that someone, or something, was responsible for the motion of the door, and that its closing will satisfy one of their goals.

Although Figure 4 shows how processes interact in a device function, nothing specific has been said about what processes do, or how. Bounding door-use enables some inference and prediction for cyclic behavior, however, predicting and explaining door behavior requires some representation at the process level. Figure 5 details the lowest level of [process] representation in EDISON, and how it supports understanding the *constrain* processes in Figure 4.

Figure 5 shows the standard representational form for all processes and how

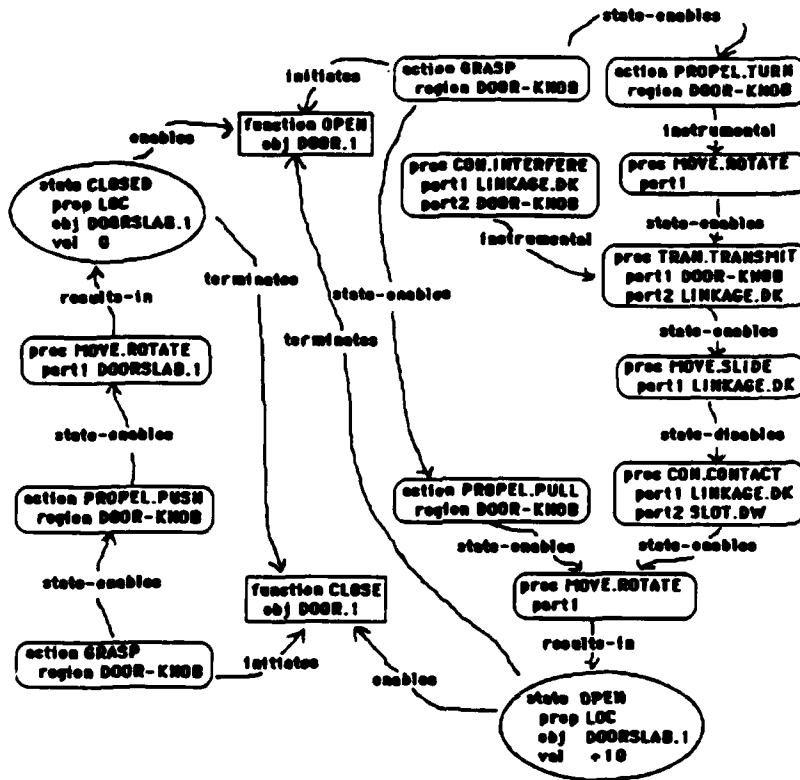


Figure 4: Representing door functions: opening and closing

different constrain processes are realized by different role bindings. The representation of processes is very similar to that of Schank's actions [5], but there are three differences: (a) processes have no actor, (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 constraints 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 actor because

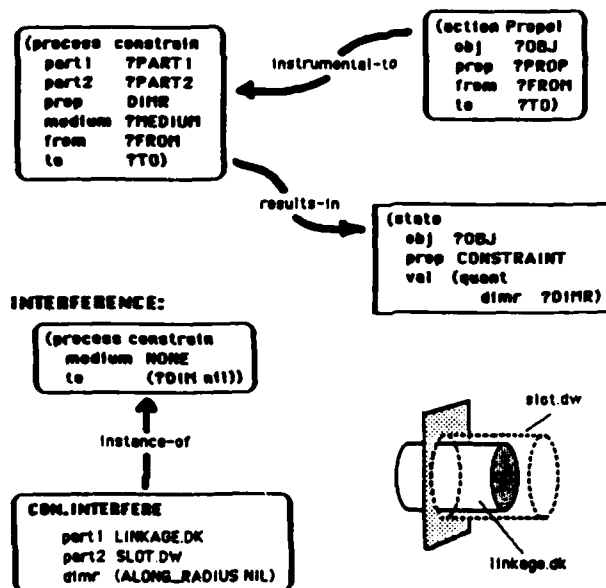


Figure 5: Representing the constrain (connection) process in EDISON

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.

There are five kinds of processes which the EDISON representation is used to describe: *move*, *constrain*, *deform*, *transform*, and *store*. Move and constrain are the most simple, and are instrumental to deform, transform, and store. From Figure 5 constrain³ can be seen to require two parts, a dimension and direction, and some connector (or medium) for maintaining the applied force which is instrumental to the process. All processes have preconditions, and constrain requires that the parts be in physical contact to one another before a connection can be made (not shown). Processes, like actions, cause physical state changes. Constrain processes always cause constraint states, on both parts. Interference is a constrain process with no medium, and causes a constraint state along an entire dimension.⁴ The meaning of *con.interfere* can now be interpreted. Linkage.dk and slot.dw instantiate the process parts. Part1 is always used to

³EDISON is based on mechanical devices. Constrain in the mechanical domain is connection

⁴As compared to contact or support, which act on specific directions along a dimension.

define the dimension used for EDISON processes. The dimension (dimr along-radius) thus refers to the linkage.dk radial dimension. The process (from and to) roles refer to the process prop role, so the interference between linkage.dk and slot.dw causes a set of constraint states for each along the linkage.dk radial dimension.

Two basic process assumptions are made in the EDISON representational model: (a) parts are free to move unless specifically constrained, and (b) initiated processes will continue unless otherwise acted upon. These assumptions, and other basic knowledge for processes and process interactions, are formulated as rules associated with a process. These rules take the place of 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 rules do, however, enable similar types of reasoning, and support process explanation. Some rules associated with comprehending the parent *constrain* process are presented and discussed on Page 14.

3.2.2 Primitive Mechanisms and Functional Comprehension

Physical processes underlie the representation of device behavior, function, and interaction. Nevertheless, devices play the central representational role in EDISON because they index directly to both intentional and physical representations. The more compact the device representation the easier it is to associate device use and behavior, and less computational effort is 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 [12]. Instead, a minimal set of mechanisms has been selected consisting of the commonly accepted basic machines [9], as well as a few mechanisms (non-machines) arising frequently in naive descriptions of device function⁵. Eleven primitive mechanisms are represented in EDISON: linkages, levers, gears, pulleys, wheel-axles, planes, screws, blades, springs, bearings, and containers. All mechanical devices can be decomposed to combinations of these mechanisms, and by understanding these mechanisms EDISON has the capacity to understand, and generate, more complex devices.

Figure 6a presents the EDISON representation for simple levers. Simply put a lever is any linkage with a fulcrum, where a linkage is itself a mechanism, and a fulcrum is a physical attribute called a *region* [13]. Whereas linkages are used to transmit or translate forces and velocities, the function of levers (Figure 6b) is to magnify force or speed; both of which are transform-related.⁶ Of course, lever function is realized in different ways depending on how the remaining

⁵From ad-hoc experiments run during the period 1985-1987 at the UCLA Artificial Intelligence Laboratory.

⁶All EDISON primitive mechanisms except springs, bearings, and containers are functionally transform-related. Springs and bearings are store-related, and containers are constrain-related

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 rules associated with the transform process, and levers in particular. One such rule, associated with simple hinges, is described and discussed on Page 14. The bindings for *door-hinge* in Figure 6c are shown as they apply to the standard device representational form. The doorhinge is really two simple levers *pinned* together. However, the effect of lever1 is nullified because motion enables lever use, and the doorway is grounded.

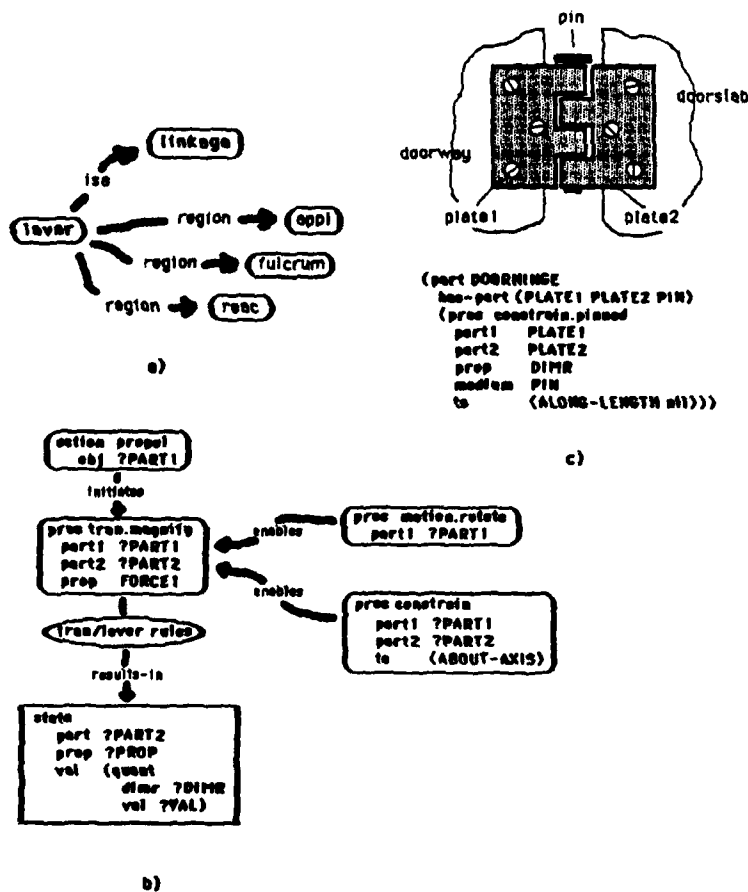


Figure 6: A doorhinge exemplifies lever representation and how simple mechanisms can be used to reason about mechanical interactions and device use; a) lever representation, b) lever function, c) doorhinge representation.

The significance of physical attributes and relations is that all device-related knowledge structures should index directly to device use, or to a process in which the attribute (or relation) is required. The device representation thus becomes a structure which always indexes into both intentional and physical representations. EDISON will always be able to say which device attribute is responsible for a particular use, or why an intended use failed. Physical regions exemplify this point by reducing the complexity of spatial descriptions, and by differentiating uses and processes; they describe functional areas of a device. How do we recognize the futility of trying to cut a metal rod with a rolling pin? People recognize that cutting requires a part with a sharp edge (where sharp is a regional descriptor associated with the *deform* process of cutting), and that a rolling pin simply doesn't have one. The door-hinge fulcrum 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 enables a broad view of physical interaction. Specifically, 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 the functions opening or closing. 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 functional analysis is used during comprehension of text describing mechanical situations. Consider the inferences required to understand the text of **Broken Foot**.

Example 2: Broken Foot

"The door would have closed but his foot was there"

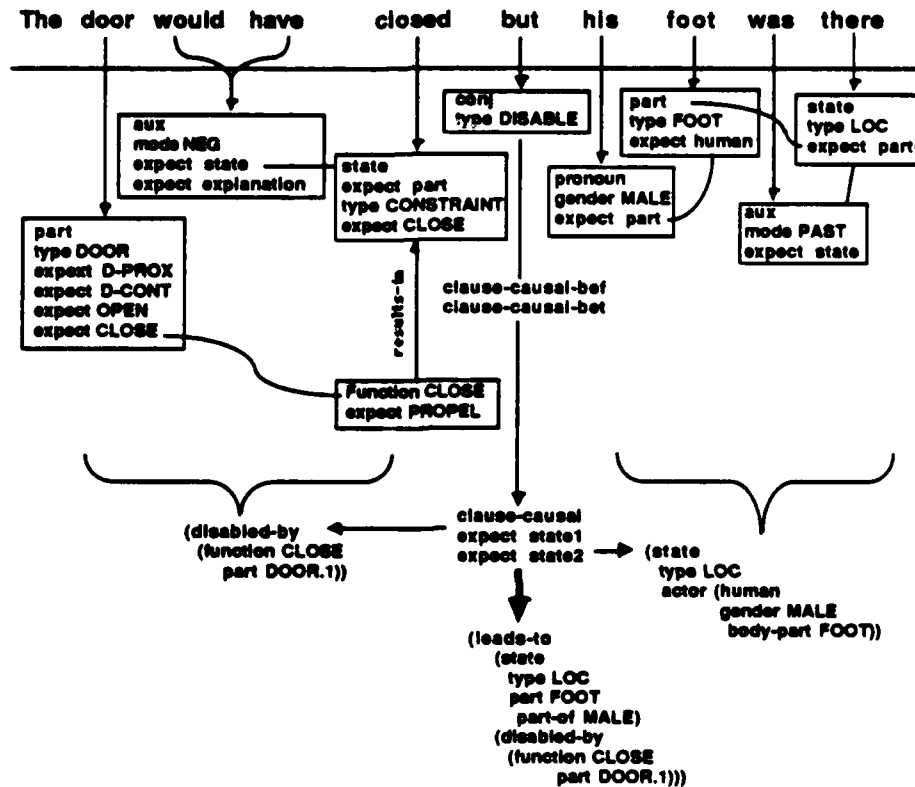


Figure 7: Comprehending Broken Foot using process theory

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 door use [1]. 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 constraining door motion or by eliminating the propelling force (see Figure 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 constrain door motion, requiring that it be along the door path of motion according to a rule associated

with motion (M1, below).

M1: If a part O1 disables motion of part O2, then the O1 lies along the path of O2's motion.

The integration of process and machine knowledge from the last two sections now enables an explanation to be constructed for the buggy doors in Figure 2. Process and machine-specific rules, such as H1 and C1-C4 in Figure 5, are generalized from the application of process relationships to specific parts and situations.

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 constraint on the objects is the union of constraints 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 constraints 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 2 is a global (device) constraint which disables motion.

3.2.3 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 both functionally and intentionally to account for this contrast. Functionally, device physical characteristics and relations index to physical processes. Intentionally, device physical characteristics and relations *must* index to the context which motivates device use. The relatively small number of primitive mechanisms and device-related representational constructs, 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 *properties* associated with a situation. A device property is a comparison between a device attribute value and its boundary values. For example, we may consider a faucet *leaky* if it won't close all the way. The comparative attribute 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 property thus tells us the point of view whereby device function is evaluated. Device properties can index to any contextual component, and so device use can be interpreted in context. Also, because the physical attribute is directly associated with a physical process EDISON can infer which function the situational context refers to.

Design episodes in EDISON are comprised of four components: (1) settings, (2) states, (3) devices, and (4) intentional structures such as actions, goals, and 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 Figure 8. One door may be used in a bank vault as security, while the other door is used in a flood for flotation.

Context (possible drowning) helps to reinterpret door use as flotation despite the prototypical use for containment.

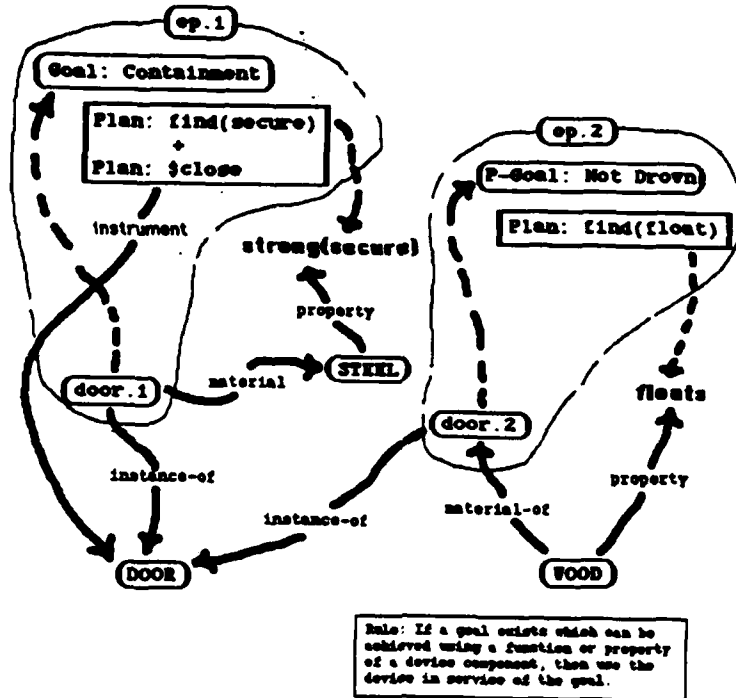


Figure 8: Contextual determination of door use in flooding and bank vault

The environmental state of flooding motivates a not-drown (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 script⁷ builds expectations for money containment (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.

4 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)

⁷The use of \$ follows the convention used by Schank and Abelson [5] for scripts.

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 [14], extensive use of analogy and combination when the task given to the children was to create novel devices. Making device 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 reminders⁸ 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

⁸Reminders are spontaneous similarity-based retrievals, see Schank [15].

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 2 the door wouldn't secure were the linkage and slot not coaxial, a state which would render the device useless for door constraint. 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.

5 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 [15] has argued that failures are important because learning occurs at failure points. Dyer [6] 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 [16] 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 reminders 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 reminders 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, reminders 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 reminders 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.

6 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, representations. The difference lies in the approach. EDISON is directed at reasoning about multiple device uses, and emphasizes a simple representation for behavior. This limits the ability of EDISON to simulate device behavior. We believe that this representational outlook is a necessary component to an overall representational scheme which can support creativity.

References

- [1] Michael G. Dyer, Jack Hodges, and Margot Flowers. Computer Comprehension of Mechanical Device Descriptions. In John Gero, editor, *Knowledge-Based Systems in Engineering and Architecture*, Addison-Wesley, 1987.
- [2] Michael G. Dyer, Margot Flowers, and Jack Hodges. Naive Mechanics Comprehension and Invention in EDISON. Tenth International Joint Conference on Artificial Intelligence, Morgan-Kaufmann, Milan, Italy, August 1987.
- [3] Michael G. Dyer, Margot Flowers, and Jack Hodges. Edison: An Engineering Design Invention System Operating Naively. *International Journal of Artificial Intelligence in Engineering*, 1(1), July 1986.
- [4] Michael G. Dyer and Margot Flowers. Automating Design Invention. *Autofact 6*, Anaheim, CA, 1984.
- [5] Roger Schank and Robert Abelson. *Scripts, Plans, Goals, and Understanding. The Artificial Intelligence Series*, Lawrence Erlbaum, Hillsdale, NJ, 1977.
- [6] Michael G. Dyer. *In-Depth Understanding: A Computer Model of Integrated Processing For Narrative Comprehension. Artificial Intelligence Series*, MIT Press, Cambridge, MA, 1983.
- [7] Kenneth Forbus. Qualitative Process Theory. *Artificial Intelligence*, 24:85-168, July 1983.
- [8] Martyn Bramwell and David Mostyn. *How Things Work*. Usborne Publishing, 1984.
- [9] Bureau of Naval Personnel, editor. *Basic Machines and How They Work*. Dover Publications, Inc., 1971.
- [10] Harvey Weiss. *Machines and How They Work*. Thomas Y. Crowell, 1983.
- [11] Wendy G. Lehnert. *The Process of Question Answering*. Lawrence Erlbaum Associates, 1978.
- [12] Marcello Alonso and Edward J. Finn. *Physics. Series in Physics*, Addison-Wesley Publishing Company, Inc., Reading, MA, 1970.
- [13] Michael G. Dyer, Margot Flowers, and Jack Hodges. Edison: An Engineering Design Invention System Operating Naively. In D. Sriram and R. Adey, editors, *Applications of Artificial Intelligence in Engineering Problems*, pages 327-342, First International Conference, Springer Verlag, Southampton, United Kingdom, April 1986.

- [14] Edward DeBono. *Children Solve Problems*. Penguin, 1980.
- [15] Roger Schank. *Dynamic Memory*. Cambridge University Press, 1982.
- [16] R.G. Crowder. *Principles of Learning and Memory*. Lawrence Erlbaum Associates, Hillsdale, N.J., 1976.