Naive Mechanics Comprehension and Invention in EDISON

Michael G. Dyer
Margot Flowers
Jack Hodges

Artificial Intelligence Lab.
Computer Science Dept.
3532 Boelter Hall
UCLA, Los Angeles, CA 90024

Abstract

This paper briefly reports on the current architecture and status of EDISON, a computer model being designed to (1) understand natural language descriptions of mechanical devices and (2) generate novel device representations through heuristic strategies of mutation and analogy. The representational constructs in EDISON must support both of these tasks and include: goal/plan information, spatial relationships, forces, motion, contact, regions, constraints on (and principles of) device operation, levels of abstraction, and naive mechanics dependencies and inferences.

1. Introduction

The EDISON research project was created to explore processes of comprehension (Dyer, Flowers & Hodges in press) and invention (Dyer. Flowers & Hodges 1986) in the naive mechanics domain (Dyer & Flowers 1984). These tasks require basic research in areas of: memory organization. disambiguation, inference, learning, problem solving and representation of knowledge. Our 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.

2. EDISON Architecture

The current EDISON system is composed of eleven elements, shown in figure 1 below. In this figure, thin lines with arrows indicate the flow of input/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).

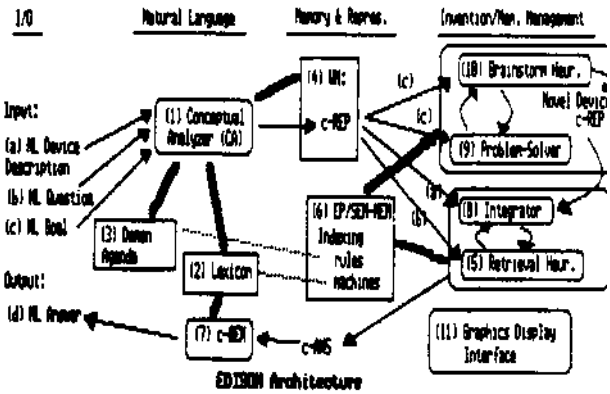

Figure 1

EDISON accepts three types of natural language input: (a) the *description* of a mechanical device, possibly novel to EDISON and possibly including information concerning device principles of operation, structural arrangment, and/or mode of operation, (b) a *question,* concerning causal relationships involved in device structure, function, and operation, or (c) a *goal specification,* which requests the creation of a novel device, possibly satisfying one or more constraints.

If a question is input to EDISON, it is passed to the conceptual analyzer ((1) in figure 1). The conceptual analyzer (CA) is a variant of McDYPAR, a demon-based CA first used in the BORIS story understanding system (Dyer 1983). As each input word (or phrase) is encountered, the CA accesses its corresponding entry in the lexicon (2). The lexicon contains mappings between words/phrases and representational fragments, implemented as frames (Minksy 1977). Attached to each frame are zero or more demon templates, along with parameters supplied to these demon templates by the lexicon. Once supplied, demon instances are spawned . representing delayed, active rules/processes within EDISON. Active demons are maintained in a demon agenda (3) and are polled to simulate concurrent rule application and testing. Demons seek to connect up conceptual fragments into large conceptual structures. Each conceptual fragment anaVor demon may access knowledge of a mechanical device, device component, device motion, device force, device region, etc. This knowledge is stored in episodic/semantic memory (6). Active conceptual fragments are stored and manipulated in a working memory

Once a completed conceptual representation has been formed (c-REP), if it represents the conceptual content of a question ((b) in figure l)» then it is passed to the memory management subsystem, where retrieval heuristics (5) are applied. An answer to the question is then sought by accessing episodic/semantic memory (6). Semantic memory (SEM-MEM) holds general mechanics knowledge while episodic memory (EP-MEM) holds knowledge of instances of specific mechanical devices. These specific device exemplars may exist in EP-MEM for any one of three reasons: (1) the programmer handcoded the device into memory, (2) EDISON read about the device, or (3) the device came about as the result of EDISON's invention heuristics. A question may concern general mechanical relationships (in this case SEM-MEM will be searched) or a specific device EDISON has just read about (then EP-MEM will be searched). If a conceptual answer (c-ANS) is retrieved, then it is passed to a conceptual generator ((7) in figure 1), which accesses lexical information to produce a naturallanguage answer.

When a device description is input to EDISON, the result of conceptual analysis (c-REP) is also passed to the memory management subsystem. The integrator (8) must determine where in EP-MEM the device representation belongs. Integration into memory involves: (a) determining if the device already exists in memory (i.e. is not novel to EDISON), (b) constructing an instantiation of the correct internal format for integration into memory, and (c) building the necessary indexing structures for future access. This last step may include processes of generalization (so that the device is available at various levels of abstraction).

* This research is supported in part by an ONR grant (N00014-86-K-0615) to the first two authors. Eariy exploratory research in this area was supported in part by an IBM Faculty Development Award to the first author, and by a grant from the Hughes AI Center.

When a design goal specification is input, the c-REP is passed to the invention management sybsystem. 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). Brainstorming consists of heuristics which attempt to create novel devices, by two general strategies: (1) *mutation,* where a given device representation is altered or combined with other device components, or (2) *analogy,* where a device representation is generalized and another mechanism is recalled (from a different context) which shares features at an abstract level with the given device. The recalled device is then adapted to the target context.

If the goal specification includes a number of constraints, the c-REP is passed first to the problem-solving component of the invention management subsystem ((9) figure 1). The problem-solver attempts to apply mechanics rules and principles to satisfy mechanics constraints. When the problem solver cannot recall a solution from memory, it calls upon the brainstorming heuristics to invent a device. For each device invented, constraint satisfaction is applied.

All novel devices (whether from comprehension or invention) are placed in EP-MEM. Although EDISON currently cannot generate natural language descriptions of arbitrary devices, the graphics interface (11) does display a graphical representation of the semantic relationships constructed to represent a device in memory.

3. Naive Mechanics Representation

A *naive mechanics representation* (NMR) must support comprehension, problem solving, learning, and invention. The NMR used by EDISON is not finalized, but the general approach is that of representing mechanical areas, forces, motions, objects, and relations in terms of conceptual dependencies, along with associated inference rules.

3.1. NMR Requires Goals/Plans for Problem Solving

Consider the nail clipper in figure 2. Most people, after looking at this figure for a moment, realize that this particular nail clipper simply will not work. It then takes them a moment longer to realize (in exact detail) why it will fail. This comprehension process often requires that they re-examine' exactly how a bug-free nail clipper actually functions.

EDISON must be able (1) to receive a conceptual representation of a nail clipper, (2) recognize it as a nail clipper, either from a label; or by comparing the input representation with knowledge in memory, and (3) realize that this representation has a 'bug' and understand now the bug interferes with the function of the device.

In order to spot device errors, EDISON must be able to analyze each component 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 preservation, achievement of social status, finding an agent etc. In the naive mechanics domain, goals involve physical transformations, such as connection and separation. For instance, the nail clipper achieves the goal of separating a nail into two objects: the remaining nail and the nail clipping. In story domains, goals are achieved through the execution o? plans and a number of plans may exist which are able to achieve a single goal. Likewise, in the naive mechanics domain, goals are also achieved through abstract plans, but here realized though the operation of physical devices. For example, the goal of separation can be realized, e.g., by shearing, slicing, ripping, or cutting. The cutting performed by a nail clipper involves forcing Two wedge-shaped objects against both sides of the object to be cut. There are many problems to be solved here, including: where separation is to occur and how the cutting edges are to be attached (i.e. alignment), and how the cutting edges are to be removed after cutting has occurred (answer: use a spring). But the main problem is how enough force is to be accessed, i.e. the goal of attaining mechanical advantage (Weiss 1983), which can be realized by a number of plans, involving the use of various devices. In figure 2, a lever is selected. Here, however, the lever has been attached incorrectly, so that force is being applied only to part 1, rather than to both parts 1 and 2. The solution of the modern nail clipper (figure 3) h actually rather elegant, since it involves attaching the lever to a post that is sunk through the center of part 1 and attached to part 2.

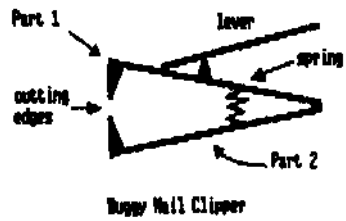

Buggy Nail Clipper

Figure 2

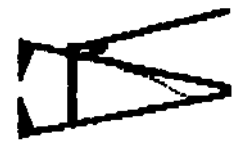
Modern Nail Clipper

Figure 3

In addition, in the modern nail clipper solves one other [4]bug* (i.e. that the lever must be pulled upward in the buggy case) by modifying the lever to have a protrusion on the same side.

Notice that one can understand the function of the modern nail clipper and detect mistakes within the buggy nail clipper, all without having to understand the principle behind the mechanical advantage of the lever. One need only know that levers realize mechanical advantage. The principle behind this advantage, simply stated, is that one moves a greater distance with a constant Force F to produce a greater force F' over a shorter distance (i.e. the same principle of the inclined plane). Although the principle behind a device may not be necessary to understand the device's operation, it can be important during invention. Thus, a complete representation of a nail clipper must include the following information:

goal: separation, plan: cut
device: wedges forced against obi.
goal: mechanical advantage, plan: lever
device: post w/ one-sided lever

3.2. NMR Requires Spatial 'Gestalt' Structures for Comprehension

In addition to the functional relations described above, a NMR must include the spatial orientation of each component in the device and its connectivity and orientation (Lehnert 1978) with respect to other components. Consider the following text:

Turn the screw to apply force and deform the object.

When mechanical engineers read this text, they tend to form a mental image of the spatial relation of the object to the screw. Furthermore, they fill in missing information in the text to make the text coherent. As a result, they can answer questions about information not directly supplied in the text:

Q: Where is the object in relation to the screw?
Q: What holds the screw?
Q: What holds the object as the screw moves?

The supplied information comes from their general knowledge of what a screw looks like and how it is used to apply force or form a connection. In the case of applying force, one spatial gestalt is depicted in figure 4. Notice that the object is placed at one end of the screw (instead, say. along a side of the screw). In addition, there is a framework for holding both the object and the screw, along the direction the screw will move if turned clockwise.

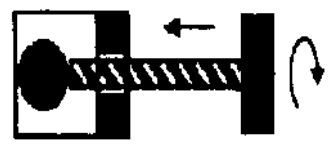
Screw Press

Figure 4

For the nail clipper, this spatial/configurational information consists of the parts of the nail clipper, their relative orientations, regions of parts where connectivity or other constraints hold, and the nature of those constraints.

3.3. NMR Requires Process Information for Prediction

Mechanical devices move and their components move. The movement of each component is a function of the connection of one component to another, the nature of the component materials/shapes, and the nature of the forces between them (Forbus 1983). For example, if a string: S is attached to a free object O and the string is pulled, then D will move in the direction of the pull once the string is taut; however, one cannot push on a string and expect O to move. In contrast if a rigid rod is connected to O, then pushing on the rod will transmit motion to O.

In the nail clipper case, process information consists of a sequence of the motions performed by the nail clipper, as a result of the connectivity of the parts of the nail clipper and forces applied to those parts.

3.4. Limits of NMR in EDISON

Currently EDISON does not have any capability of examining visual or iconic device representations. Thus, any spatial or connectivity information is hand-coded into memory. Probably the most difficult representational and reasoning tasks in mechanics involve kinematics (Forbus 1986), especially the interrelations among forces and 3-dimensional object shapes. In general, we are avoiding this very difficult class of problems. We want to see what comprehension and invention tasks EDISON can perform without a serious theory of kinematics. We believe the remaining, simpler world of basic connectors and motions is still extremely challenging.

4. Status of NL Comprehension in EDISON

Currently, the natural language subsystem of EDISON can handle just a few. single-sentence length texts. A sample text (Dyer, Flowers, Hodges in press) reaefby EDISON is:

> TOY GUN: An object is pushed into a barrel, against a spring, compressing the spring until it catches on a trigger.

The EDISON lexicon contains mappings from words, such as "push" and "pull", to expectations (demons) for the direction and source of the force, along with the object to be moved. Objects, such as "barrel", are represented in terms of containment and constraints on motion. The term "against" is represented in terms of both spatial proximity (e.g. "the painting against the wall") and force (e.g. X "PROPEL O against" Y). Some objects, such as springs, are represented as primitive devices,, with attached inference rules. E. g. if one pushes on a spring, the spring will push back, Trigger" serves as an enabling/disabling device for release of a force. Note that "it" in TOY GUN could refer to the spring, barrel, object, or entire gun. However, syntactic constraints prefer "spring" while semantic constraints disallow both the barrel ana the entire device as the referent of "it". In some toy guns the trigger catches on the spring; in others, on the object (e.g. where the object is a plastic dart with a notch at the end).

The ability for EDISON to build a complete representation of the text depends on (1) what EDISON already knows in memory about toy guns, (2) what EDISON already knows about barrels, springs, triggers and objects in general and (3) what EDISON already knows the top-level goal of the device to be (in this case, to propel the object in a given direction). In general, mechanical device descriptions are difficult for people to read unless they already know something about the device under discussion, or know how to read (and have access to) a visual drawing containing the gestalt configuration of known iconic elements making up the device. In EDISON's case, the drawing is representee as a hand-coded conceptual representation already present in memory.

5. Status of Invention in EDISON

Currently, EDISON is capable of (re)inventing the swinging bar room door (figure 3A) through the process of applying operations to, and altering features of, a standard door. In this case, one way of arriving at a bar room door is to apply the CUT operation to the door slab to alter the number of slabs and then apply problem solving to attach the free-standing slab via hinges to the other side of the door frame.

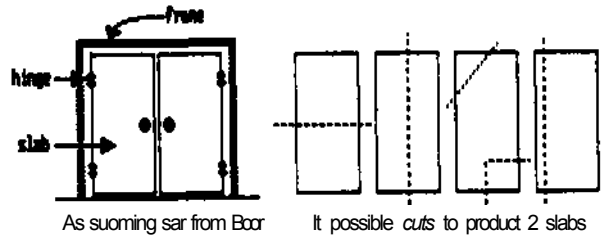

Figure 5

Even in this simple invention scenario the possibility space is huge; e.g. the slab can be cut in numerous ways and the position of a hinge can conceivably be anywhere on the surface of the slab (figure 5B). However, ninge constraints reduce this space to positions along the edge of the slab. Still, various strange bar room doors result if the hinges are placed at the top or bottom of the free-standing slab (figure 6). In addition, an 'accordian' door can result if ninges are used to attach the free-standing slab to the other slab (figure 7: left).

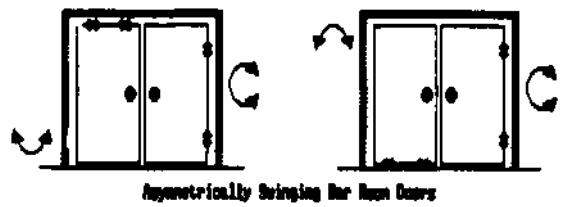

At this point EDISON is incapable of making use of this serendipitous invention to realize that (1) the cut operation can be used several times to reduce the size of each slab while increasing the number of slabs, and thus cover the same open area while reducing the width of the door when open (figure 7: right) and (2) the slabs can be attached to a runner so that the door effectively slides open rather than swings open.

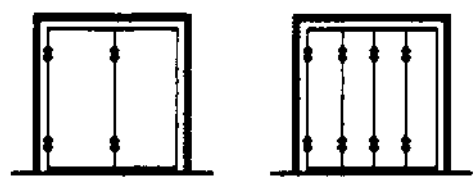

Figure 7

5.1. Ingenuity, Problem Solving and Experience

Two issues of importance for invention are (1) recognition of ingenuity in a vast search space of uninteresting and/or useless devices and (2) tradeoffs between efficient problem solving through effective use of memory and role of experience in 'brainstorming' for novel (or possibly overlooked) patterns.

Change in the fundamental motion of an object is one heuristic in EDISON for recognizing that a truly novel design has been discovered. Thus, the invention of a sliding door through mutation of a swinging door should be recognized by EDISON as an event of significance.

EDISON maintains devices in episodic memory. If the problem solver solves each new problem by simply recalling a past solution, then inventiveness will dimmish as the number of devices grows. However, with human inventors, such as Thomas Edison, the acquisition of a novel device serves as a platform for coming up with more devices. Such inventors use processes of analogy and adaptation to apply knowledge in one domain to create a device in another domain. In this way. growth in episodic memory increases the potential of inventiveness rather than diminishes it A simple example is that of inventing a new nut cracker by adapting the mechanism of the nail clipper (figure 8A).

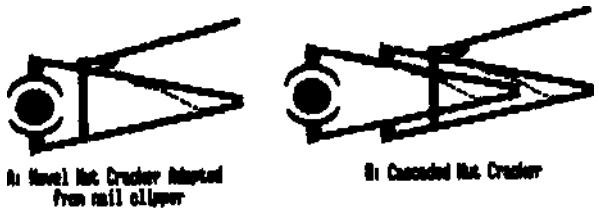

Figure 8

The issue of ingenuity is illustrated in figure 8B, where the nail clipper mechanism has been pointlessly complicated. However, cascaded operations can result in novel and useful devices. Consider the rose clipper in figure 9B. This design allows one to cut at a greater distance without having to greatly increase the movement of the handles (figure 9C).

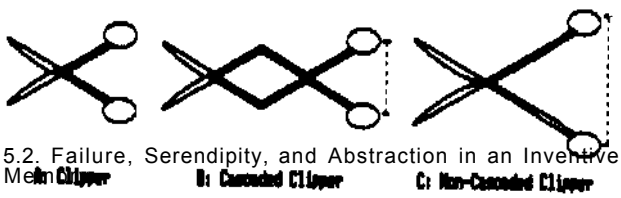

Figure 9

5.2. Failure, Serendipity, and Abstraction in an Inventive Memory

What is the role of failure in memory? Schank (1982) has argued that failures are important because learning occurs at failure points. Dyer (1983) has shown that plan failures represented at an abstract level serve as an indexing structure to cross-contextual memories. If every trivially baa design is stored in EDISON's episodic memory, then problem-soiving efficiency may suffer, as a result of recall of 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. This means that EDISON must generalize specific failures wherever possible and store the abstracted negative design experiences in episodic memory. However, sometimes it is incorrect (from a brainstorming point of view) to avoid exploring a possibility space because of past failure. Why would an area blocked by past failure be worth re-exploration? Often, new mechanical devices, principles, materials, etc. become available since the time that design area or approach was abandoned. So a robust invention system must know when to re-explore an area because of new invention in potentially related area. These tradeoffs, between efficient problem solving and invention, are topics of current research.

As with the cascaded clippers, a mutation which fails to satisfy a goal in current focus may end up serving another design goal. Consider the last cut snown in figure 5B, which produces a lopsided bar room door and appears to be a design of little value. However, such doors do exist In refrigerators, the freezer door is often of less width in comparison to the food door. Thus, an invention may not achieve an active design goal, but turn out to be of use for another design context An inventor can only make use of serendipitous effects if a large number of design goals are concurrently active in memory.

6. Relation to Work of Others

The overall approach toward naive physics in EDISON is inspired by Hayes,e.g. (Hayes 1985). The need and utility of an episodic memory or device exemplars is taken from (Schank 1982) and (Kolodner 1984) and their general work on episodic memory organization. The use of heuristics of invention and heuristics to assess interestingness of concepts (here, device ingenuity) are inspired from Lenat's work on invention in mathematics and geometry (Lenat 1976,1983).

The representation constructs in EDISON share features with those of Lehnert's object primitives for comprehension of stories involving the use of objects, described in (Lehnert 1978); Rieger's CSA representational scheme to represent such objects as flush toillettes and light bulbs (Rieger 1975) and the work of Forbus on qualitative processes (Forbus 1983, 1985). The natural language comprehension component shares representational similarities with the representations of physical objects read by the patent abstract conceptual analyzer of Wasserman and Lebowitz (1983).

7. Conclusions

Naive mechanics comprehension and invention can be modeled in terms of symbolic manipulations on representational constructs. Device comprehension consists of accessing conceptually dependent representations and combining them to form larger coherent structures. Device invention consists of altering device representations through goal/plan analysis, constraint satisfaction, feature mutation, and processes of abstraction and analogy. While the resulting approach lacks the detailed numerical/simulation capabilities or mathematical models typically used in mechanical engineering, it provides the potential capability of modeling the engineer's cognitive processes of comprehension and invention.

References

Dyer, M. G. *In-Depth Understanding.* MIT Press, Cambridge, MA. 1983.

Dyer, M. and Flowers, M. Toward Automating Design Invention. *Proceedings ofAUTOFACT 6 Conference,* Anaheim, CA. 1984.

Dyer, M. G„ Flowers, M. and Hodges, J. EDISON: An Engineering Design Invention System Operating Naively. *International Journal of Artificial Intelligence in Engineering* Vol. 1, No. 1, pp. 36-44,1986.

Dyer, M. G., Flowers, M. and Hodges, J. Computer Comprehension of Mechanical Device Descriptions. In J. Gero (ed) *Knowledge-Based Systems in Engineering and Architecture,* Addison-Wesley, Reading, MA (in press).

Forbus, K. Qualitative Process Theory. *Artificial Intelligence.* 24, 85-168,1983.

Forbus, K. Chapter 5. *Formal Theories of the Commonsense World.* Abkx. pp. 185-226,1985.

Forbus, K. Invited Talk. *AAA1-86,* Phila. PA, 1986.

Hayes, P. The Second Naive Physics Manifesto. In Brachman & Levetque (eds). *Readings in Knowledge Representation.* Morgan Kaufmann, Los Altos, CA. 1985.

Kolodner, J. *Retrieval and Organisational Strategies in Conceptual Memory: A Computer Model.* LEA, Hillsdale, NJ. 1984.

Lehnert, W. *The Process of Question Answering.* LEA, Hillsdale, NJ. 1978.

Lenat, D. *AM: An Artificial Intelligence Approach to Discovery in Mathematics as Heuristic Search.* Stanford Univ. PhD. Palo Alto, CA. 1976.

Lenat, D. *EURISKO: A Program that Learns New Heuristics and Domain Concepts.* Artificial Intelligence. 21(2,1), 61-98,1983.

Minsky, M. A Framework for Representing Knowledge. In Johnson-Laird and Wason (eds). *Thinking: Readings in Cognitive Science.* Cambridge Univ. Press, MA. 1977.

Rieger, C. An Organization of Knowledge for Problem Solving and Language Comprehension. 1975. Reprinted in Brachman and Levesque (eds). *Readings in Knowledge Representation.* Morgan Kaufmann, Los Altos, CA. 1985.

Schank, R. *Dynamic Memory.* Cambridge Univ. Press, Cambridge, England, 1982.

Wasserman, K. and Lebowitz, M. Representing Complex Physical Objects. *Cognition and Brain Theory.* Vol. 6, No. 3,259-285,1983.

Weiss, H. *Machines and How They Work.* T. Y. Crowell. 1983.