AAAI-93, Working Notes, Reasoning About Function, pp. 137 - 140.

An Ontology of Mechanical Devices

Thomas F. Stahovich
Randall Davis
Howard Shrobe
MIT Artificial Intelligence Laboratory
Cambridge, MA

We are working to develop a large scale ontology for the mechanical engineering world to support a wide range of tasks including analysis and design. Our work is guided by the task of determining the behavior of a mechanical device from a description of its geometry (the shapes of its parts and how they are connected) and its driving inputs. We look for common patterns of behavior and label them with the terms that mechanical engineers use to talk about mechanical devices. We attribute function to the components of a device by relying on the assumption that their intended purpose is to provide the identified behavioral patterns.

We began in the familiar fashion, by examining the language used by engineers to describe how mechanical devices work. We looked at the explanations in Artobolevsky's six volume *Mechanisms in Modern Engineering Design*, with its 5,000 mechanical devices, noted the "mechanical elements" in Shigley's text, *Mechanical Engineering Design*, and, in the spirit of great tinkerers everywhere, took apart several mechanical devices, including cameras and circuit breakers, to see "how they worked" and to try to give precise descriptions of their function.

While people do indeed have informal, tacit understanding of terms like clutch and lever, making those definitions sufficiently explicit and accurate for machine use turns out to be interestingly challenging. We have spent a surprising amount of time getting even very basic concepts suitably refined.

Consider, for example, one of the most primitive mechanical devices, a lever. The term is sufficiently familiar that the definition would seem trivial. One obvious answer is "a rigid bar with a pivot." Figure 1 is a simple example showing some of the subtlety of getting the definition right: the pivot is on the left, there is a weight at the right, and a stack of blocks in the middle. By the definition given, this is a lever. But an engineer would tell us that this is in fact an overhanging beam supporting a weight [Popov68].

Deciding that the blocks are the problem, we might repair our definition by saying that a lever is "a rigid bar with a pivot that is not prevented from rotating by another object." Figure 2 is a simple example of our new definition. The bar is being used to amplify the force exerted by the person so that the weight can be lifted more easily.

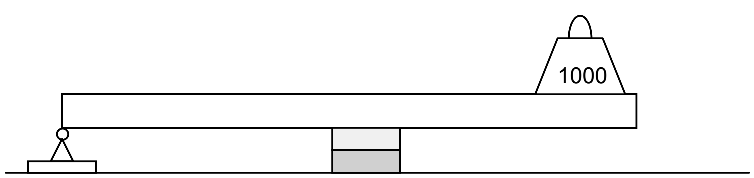

Figure 1

The Ontology

One of the interesting discoveries we made early on was how much understanding of devices is taken for granted in standard sources, even the introductory texts. For example, Shigley's definition of a shaft is "a rotating or stationary member, usually of circular cross section, having mounted on it such elements as gears, pulleys, flywheels, cranks, sprockets, and other power-transmission elements." By this definition, nearly every mechanical component is a shaft. This is a widespread phenomenon: text books uniformly assume that the basic definitions are so obvious that no explanation is needed.

An engineer would certainly agree that this is a lever, so we seem to be making progress.

Now imagine we have a specific mechanical task in mind: we are gluing two blocks and need to squeeze and hold them together with a large force while the glue dries. We could put a weight directly on top of the blocks, but we would get a much larger clamping force if we put the blocks under the bar, as in Figure 1.

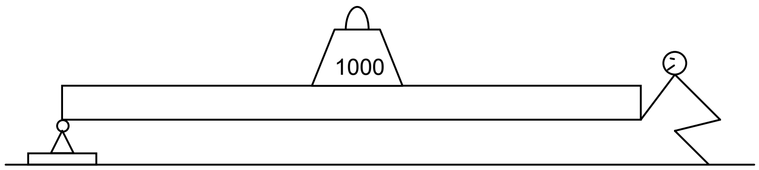

Figure 2

In this circumstance we would be using the bar to amplify the force of the weight and clamp the blocks together. Viewed from this perspective, the device in Figure 1 now seems to be similar to the one in Figure 2; perhaps the device in Figure 1 is a lever after all?

The fundamental difficulty here is that our definition of a lever as "a rigid bar with a pivot" is a structural definition. But a lever is not a *structure*, it is a *behavior*. A rigid bar with a pivot (the structure) can behave as a beam, or it can behave as a lever. If the bar in Figure 1 is being used to support the weight, it is (behaving as) a beam. If, on the other hand, the bar is being used to amplify the force of the weight and clamp the blocks, it is behaving as a lever, just like the device in Figure 2. Hence we ought not say the bar is or is not a lever; but that it is or is not behaving as a lever. To emphasize the distinction, henceforth we refer to lever-behavior.

is in order to decide which behavior is exhibited. If the force of the weight is the input, the bar is exhibiting lever-behavior; if the force of the blocks is the input, the bar is not exhibiting lever-behavior.

A second example will reinforce the point. Consider the three devices in Figure 3 and ask, What kind of devices are they? Clearly they are all ratchets. But examine their structures: there is not a single *physical* component common to all three devices. Clearly the relevant concept is not "a ratchet," rather it is "ratchet-behavior."

These examples support our belief that the appropriate fundamental ontology for mechanical engineering ought to be organized around behavior, not structure. To date our ontology contains behavioral definitions of: lever, ratchet, cam, shaft, gear, bearing, clutch, brake, latch, catch, stop, trip, and spring.

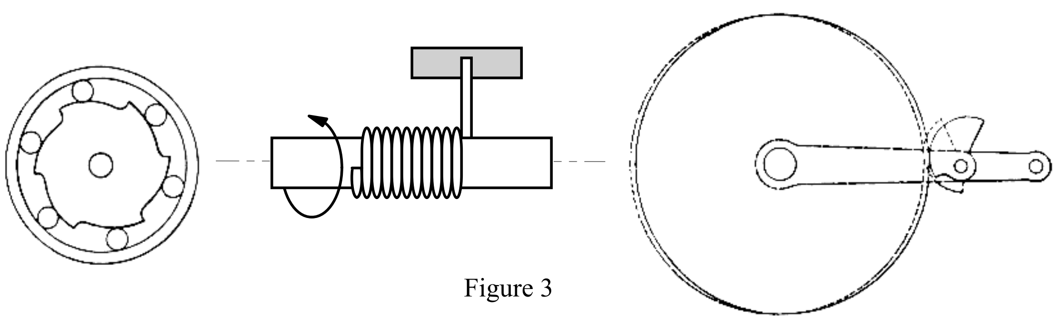

Figure 3

A slightly simplified version of our definition of lever-behavior is, "the transformation of a force, which is not a reaction force, by means of a balance of moments (torque) about a pivot." A reaction force is the constraint force imposed by a position constraint. If the bar is Figure 1 is being used to hold the weight up, then the blocks are providing a position constraint and the force of the blocks on the bar is a reaction force. In this case the behavior is not lever-behavior. If the force of the weight is being transformed to clamp the blocks, then the bar is behaving as a lever.

The notion of the "driving input" is clearly crucial to our definitions of behavior: we need to know what the input

Causal Explanations

We have found that causal explanations are a particularly useful tool in determining what behavior a device is exhibiting. Consider, for example, the device in Figure 4, consisting of a shaft constrained to rotate about an axis in the page, and a rigid link, constrained to rotate about the same axis (and hence constrained to move perpendicular to the page).

Imagine that the link is driven by an external source of periodic motion (not shown); it oscillates about the fixed axis. The shaft is also driven by an external motion source (not shown), so that its angular velocity matches that of the link as the link rotates in the one direction, but is stationary as the link rotates in the other.

AAAI-93, Working Notes, Reasoning About Function, pp. 137 - 140.

The link is oscillating, while the shaft is undergoing intermittent, unidirectional rotation. This would clearly appear to an external observer to be ratchet behavior. But the connection between the two motions is only coincidental; there is no causal relationship between the motion of the link and the motion of the shaft. As a result we do not think it appropriate to term this ratchet-behavior. More generally, a causal explanation is necessary to establish the correct classification of a behavior.

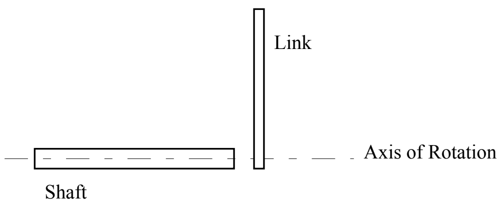

Figure 4

We can push this point one step further by considering the situation in which the shaft and link in Figure 4 are connected using the first device in Figure 3. In this situation we have the required components (a link and shaft) connected via a mechanism designed to produce ratchet-behavior. But if the link and shaft are still being driven externally, as in the previous example, we suggest that the resulting behavior is still not ratchet-behavior, because there is still no causal story that explains how the motion of the link is *causing* the motion of the shaft.

A device connected in this way has ratchet-behavior as one of its *possible* behaviors, but in the situation just described, that is not in fact what it is doing. Once again, a causal explanation is essential for distinguishing between the actual behavior and the possible behavior of a device.

We have been exploring the use of energy flow as a means of finding causal explanations. Consider the third device in Figure 3, consisting of a wheel, an arm, and a semicircular pawl. Imagine the arm is oscillated by an external motion source and the wheel is connected to a rotating load which has both inertia and friction. On the driving stroke the motion source will supply energy to the arm, the arm will supply energy to the pawl, the pawl will supply energy to the wheel, and finally the wheel will supply energy to the load. On the driving stroke there is an energy flow path from the arm to the wheel. On the return stroke the energy flow path will be broken because there is no energy flow between the pawl and the wheel. In this example, it is clear that wheel is caused to rotate on the driving stroke because of energy that flows from the arm. Here the energy flow path generates the appropriate causal path. We believe that we will be able to generalize this notion, and find a variety of paths by which causality is transmitted.

RELATED WORK

There is a large and growing body of literature in the area of representing and using functional knowledge. Representations for function and purpose can be found in [Keuneke91, Franke91, Pegah93]. Functional knowledge has been used, for example, in diagnosis [Abu-Hannah91], debugging [Allemang91], design improvisation [Hodges92], and design evaluation [Bradshaw 91].

Our task is to determine the behavior of a mechanical device from a description of its geometry (the shapes of its parts and how they are connected) and its driving inputs. There has been similar research in other domains. Rich and Shrobe developed a system for understanding computer programs [Rich76]. They reasoned from the structure of the program to determine the purpose of its parts. deKleer worked in the domain of electric circuits and developed a system which starts with a structural model of a circuit, generates causal explanations for its behavior, and parses the behavior into behavioral features used by electrical engineers [deKleer79].

Our work is most closely related to the work on device understanding by Shrobe [Shrobe93] and Joskowicz and Sacks [Joskowicz90]. Shrobe parses a numerical simulation of a linkage to identify the function of its parts. Joskowicz and Sacks use a region diagram (configuration space) to produce a description of the behavior of mechanical devices. Shrobe's system is limited to fixed topology mechanisms, and Sacks and Joskowicz do not address forces, such as friction. Our work attempts to extend these methods to variable topology mechanisms with forces.

FUTURE DIRECTIONS

We believe that the generation of causal explanations and identification of behavior will be useful for a variety of tasks. It will, for instance, support the analysis of mechanical devices. If a device is being used to provide a specific behavior, there is often a particular set of questions that will be asked about the device. If the device is a clutch (*i.e.*, exhibits clutch behavior), the questions that might be asked are: "How much torque will it transmit?" "What will the temperature rise be?" "How much actuating force is required?" "How much energy loss is there?" A software system that could recognize behaviors might guide the analysis by suggesting what questions should be asked. Even more interesting is the idea that the causal explanations might be useful for setting up the types of equations that an engineer would use during analysis. An engineer can write down small sets of relatively simple equations, rather than the large sparse matrices of brute force numerical simulators, because he has a qualitative understanding of the device and its behavior. Causal explanations appear to be one good way to capture the engineer's qualitative understanding.

The program might also be used to "look over the designer's shoulder" and assist in the documentation of an evolving design. The program would identify each behavior and produce a causal explanation of it. If the program couldn't determine the behavior of a component, it would query the designer. In this way the designer would have to document only the more subtle (and hence more interesting) parts of the design.

As a redesign tool, the program would identify behaviors and query a database for alternative implementations of the behavior.

Behavior recognition software could also be used to mine suitable geometric models of mechanical devices for new implementations of a behavior, accumulating a database of design alternatives.

SUMMARY

By means of catalogs of mechanical devices, mechanical design textbooks, and dismantling of actual devices, we identified an initial set of terms that appear to be at the core of an ontology of mechanical engineering. We have been led to the interesting observation that the ontology should be based on behavior, rather than structure. To date we have identified suitable definitions for about a dozen terms and have been verifying our definitions by using them to recognize behaviors.

BIBLIOGRAPHY

Abu-Hanna, A., Benjamins, R., and Jansweijer, W., "Device Understanding and Modeling for Diagnosis," *IEEE Expert,* Vol. 6, No. 2, 1991, pp. 26-32.

Allemang, D., "Using Functional Models in Automatic Debugging," *IEEE Expert,* Vol. 6, No. 6, 1991, pp. 13-18.

Artobolevsky, I.I., *Mechanisms in Modern Engineering Design*, MIR Publishers, Moscow, 1976.

Bradshaw, J. A. and Young, R. M., "Evaluating Design Using Knowledge of Purpose and Knowledge of Structure," *IEEE Expert,* Vol. 6, No. 2, 1991, pp. 33-40.

deKleer, J., *Causal and Teleological Reasoning in Circuit Recognition,* MIT/AI/TR-529, September, 1979.

Franke, D. W., "Deriving and Using Descriptions of Purpose," *IEEE Expert,* Vol. 6, No. 2, 1991, pp. 41-47.

Hodges, J., "Naive Mechanics, A Computational Model of Device Use and Function in Design Improvisation," *IEEE Expert,* Vol. 7, No. 1, 1992, pp. 14-27.

Joskowicz, L. and Sacks, E. P., "Computational Kinematics," Tech. Report CS-TR-300-90, Princeton University, April, 1990.

Keuneke, A. M., "Device Representation, The Significance of Functional Knowledge," *IEEE Expert,* Vol. 6, No. 2, 1991, pp. 22-25.

Pegah, M., Sticklen, J., and Bond, W., "Functional Representation and Reasoning About the F/A-18 Aircraft Fuel System," *IEEE Expert,* Vol. 8, No. 2, 1993, pp. 65-71.

Popov, E. P., *Introduction to Mechanics of Solids*, Prentice-Hall, Inc., Englewood Cliffs, N.J., 1968.

Rich, C. and Shrobe, H. E., *An initial Report on a LISP Programmer's Apprentice,* MIT/AI/TR-354, December, 1976.

Shigley, J. E., *Mechanical Engineering Design*, 4th Edition, McGraw-Hill Book Company, New York, 1983.

Shrobe, H. E., Understanding Linkages, AAAI-93, to appear.