Experiential Reasoning

Brian V. Funt
School of Computing Science, Simon Fraser University,
Vancouver, B.C., Canada V5A 1S6,
Phone: (604) 291-3126 e-mail: funt@cs.sfu.ca

At the time I developed Whisper in 1976 I encountered extremely skeptical—often downright hostile—audiences and as a result I (mistakenly) published only an expurgated account of my thoughts about analog representations.[1] I am delighted that times have changed a bit and this symposium seems like an appropriate forum for a little speculation and re-evaluation.

One of the main ways Whisper would be different if implemented today is that it could be made to interact with the real world instead of a simulated diagram. Certainly now nothing stands in the way of using a real "retina" with true parallel processing along with a robot arm holding a pencil to draw diagrams. To the extent that such a system can interact physically with the environment in the course of problem solving, let me define it as being an *experiential* system. Does literally experimenting with and experiencing the world make a difference?

When a computer system solves a problem by observing the world, changing the world and then observing the new state of the world—as would an experiential Whisper—how much credit should be given to the computer system for solving the problem? Is it "cheating" to use the world? To what extent does the world act as part of the computational structure of the computer system? And if we are trying to build computer systems possessing intelligence, does the "intelligence" created by the world count as part of the computer's thought process? In other words, what defines the boundary between the computer and the world?

For most computer systems the boundary is clear. A program receives input data and generates output data; the input-output devices are at the boundary. Most computer systems do not, however, operate on diagrams or other models. In such systems the input-output loop cycles more than once: input results in output that is then taken as further input resulting in further output and so forth.

The sensors and effectors of an experiential system allow it to "read" and "write" to the world in much the same way as a disk head enables reading and writing data on a disk drive. To write, a disk generates an appropriate magnetic field at the particular instant the right part of the disk is under the head. The disk head acts as an "effector" to make a physical change on the disk surface which can then be "sensed" at a later time. As the disk rotates away from the disk head, it moves outside the range of the head, but we still do not think of it as being external.

Generally, we would say that a hard disk is a part of the computer system and that an algorithm which uses the disk to store intermediate results during a complex calculation, nonetheless, has done all its computation without reference to anything external. The crucial feature that makes the disk drive part of the computer system and not something external to it lies in the limited capabilities of the disk head as sensor and effector. The disk head can only alter the physical world in a very constrained way.

We tend to think that we think with our heads given input from our senses. This view is typified by closed-book exams. Even during an exam, however, one has access to much more than one's head that can be useful in problem solving. We bring with us our bodies and a pen and are usually given paper; other things such as space, time, mass, gravity, rigid bodies, flexible materials (clothes), colors, sound, friction, air, light, and so on are also available.

I suggest that since all these aspects of the world are always present we should consider them to be part of the computational resource available for problem solving. A problem-solving method which uses these resources—counting on one's fingers, say—is not to be denigrated as an invalid form of reasoning. An experiential system having sensor and effector capabilities advanced enough to support experimentation with the world gains many aspects of the world as computational resources. The complete computational resource of the system combines its portable structure (e.g. our bodies or a machine's cpu, memory, robot arm, camera, etcetera) and those aspects of the world mentioned above that are present everywhere.

To see how this sort of computational structure of the world can be used, consider the following Geography exam question: What countries will an airplane flying the great circle route from Vancouver to Bombay pass over? Clearly, to solve this problem you need some knowledge of the shape and location of various countries. This knowledge can be amplified, however, by re-representing whatever stored knowledge you do have in terms of a globe. In the exam, you're not given a globe, or even a sphere, but you can make one. One hand can be curled into a rough sphere and then you can draw on it with the other. If you don't have a pen, then you can locate the fingers of the other hand on the "sphere" to represent the locations of up to five countries.

The point of this exercise is that the properties of three-dimensional space and in particular the properties of possible paths between two points on a sphere are available as part of the world. These properties of the world are everywhere and so the can legitimately be used in solving the problem. A great many more spatial relationships can be derived from the stored geographical facts by re-representing them in the context provided by the world because some of the interesting geographical relationships depend upon relevant properties of space and spheres and these properties need not be part of the stored knowledge. The geographer need not also be a topologist.

Being able to use the world in this experiential way was the main point of Whisper, but to a certain extent this point was obscured by the need to simulate the world (i.e. Whisper's diagrams). To simulate the world, aspects of it must be formalized and in the process much of the spirit of using properties of the world without the need for prior formalization is lost.

[1] "Problem-Solving with Diagrammatic Representations," *Artificial Intelligence* (13), 201-230, 1980.

1992 AAAI Spring Symposium on Reasoning with Diagrammatic Representations, AAAI, Mar. 1992.