The Strange Case of Bi-Photo-ism:

Paul Churchland as the 8th Annual Selfridge Lecturer in Philosophy

Mark H. Bickhard
Henry R. Luce Professor of Cognitive Robotics
and the Philosophy of Knowledge
Lehigh University

There were some odd antics on the Lehigh campus the last week of February. That's in addition to the usual ones. In particular, there was a man taking pictures in a rather peculiar way: he would take a picture, then move off to the left or right—to the right seemed most common—and take another of exactly the same thing. Every picture was duplicated, but not identically.

Setting aside for the moment any attributive (or referential) questions concerning the mental health of this person of the "everything in twos, but not identical" compulsion (which is heretofore unknown in the literature), we are in a position to identify him: Dr. Paul Churchland, from the University of California, San Diego. Dr. Churchland was visiting Lehigh for the week of February 24 through February 28 as the Eighth Annual Selfridge Lecturer in Philosophy.

Dr. Churchland visited classes, gave lectures, and argued many points with the Philosophy Faculty Seminar. He also participated in a deep exploration of the finer points of horror movies—truly a Renaissance mind. We will review some of the activities, lectures, and discussions of that week, and, perhaps, find some clues concerning the deeper issues underlying the strange "bi-photo" behavior.

Dr. Churchland's philosophical positions are highlighted by one major antagonist and one (or two?) major protagonist. The antagonist—the position of the Forces of Darkness—is the sentential approach to mind and belief: the theory, inherent in standard ways of talking about such things, that beliefs are somehow actual sentences in the mind. Dr. Churchland points out that there are a great many difficulties with such theories. For example, we seem to have an extremely large number of beliefs, such as that "The moon is not made of Pumpernickel bread," or "It is dangerous to be run over by red trucks," and also "It is dangerous to be run over by green trucks," and so on, and it is difficult to figure out where all of those sentences would fit in your standard brain.

Another difficulty is that such an approach seems to put us in a position that is committed to one of two subsidiary alternatives, each equally unacceptable, having to do with the "minds" and "beliefs" of animals. We could conclude, for example, that animals really do have beliefs, as it would appear, and that, therefore, they really do have language-type sentences in their brains, all the way down to very primitive animals, and that there is some additional mystery reason why they cannot communicate those sentences as human beings do. Or, we could conclude that animals don't have any mentality or beliefs, contrary to appearances, and that human beings constitute, therefore, a remarkable, gigantic, saltatory, evolutionary leap away from—and massively different from—all other animals. There do, in fact, seem to be people who hold each of these positions, and charity (and space) precludes our exploring their mental health.

The difficulties of the sentential approach multiply: There are many mental phenomena, such as dreams, skills, perceptions, emotional feelings, and so on, that don't seem to have much to do with sentences; and there is no evidence from brain studies of anything like little sentences (or big sentences) in the neural-ware. Churchland argues that these difficulties, and others, make the sentential approach a very bad bet, and that this point holds even more strongly when consideration is taken of the fact that this way of theorizing about the mind, this way of talking about beliefs, has been around for millennia, and doesn't seem to have made much progress in all that time. It looks like a bankrupt theory.

Churchland's protagonist is contemporary neuroscience and connectionism. His position is, in effect, that these two fields are each addressing the same sort of phenomena, but are not quite identical. Furthermore, and most importantly, they offer an alternative to the sentential approach for understanding the mind.

Connectionism, or Parallel Distributed Processing, is an architecture for processing that is, among other things, massively parallel in its processing, and massively different from sentential models. It is an approach that was motivated by considerations of the way the brain might actually function, and it arguably captures some properties of that functioning. And, there are no sentences in it.

A connectionist network is an organization of simple processors, called nodes—usually organized in layers, usually in three layers—connected by weighted lines of transmission from the nodes of one layer to the nodes of the next layer. Each node computes its own level of activation from the activation levels communicated to it on its input lines, in accordance with the weights on those lines, and in accordance with some generally simple function of those input activation levels—AND in parallel with all the other nodes computing their own updated levels of activation. The first layer receives its activation levels directly from the environment.

So, activation comes into the first layer from the environment, is distributed to the second layer by weighted connections from the nodes of the first layer to those of the second, and then the activation levels of these second layer nodes are computed. The activation levels of the second layer nodes are, in turn, distributed over weighted connections, now to the third layer, whose nodes' activation levels are determined. The third layer is usually the final and output layer.

There are several properties of such networks that have aroused a great deal of interest and controversy. First, the third layer of nodes—the output layer—will

(continued on page 6)

Bi-Photo-ism (cont.)

generally settle into some stable pattern of activations among those nodes, and such output patterns will classify together all the possible input activation patterns that would yield the particular output patterns. That is, a Parallel Distributed Processing network can be a classifier of input patterns.

Second, there are techniques by which the weights on the transmission connections in such a network can be adjusted so as to train the network to classify input patterns in desired ways—such as to distinguish for a submarine sonar reflections indicating sea floor rocks from those reflections indicating mines. This ability to engage in something like learning has been enormously exciting.

Third, it has been demonstrated that such systems can in fact engage in types of classification that had once been thought impossible for such nets. This demonstration opened up the possibility that such nets were capable of human level pattern recognition.

Fourth, such classifications of input patterns can constitute performances of many other kinds of tasks as well—connectionist nets are not limited to just pattern classification per se. For example, one net "classifies" written English into those character combinations and contexts appropriate to "this" phoneme, those appropriate to "that" phoneme, and so on—it produces the phonemes of spoken English from the characters of written English.

Fifth, such processing in these nets is similar in several ways to brain processing. Both are massively parallel. Both nets and the brain seem to involve generations and transmissions of levels of activation—levels of neural activity in the brain, perhaps levels of impulse frequencies. Even the errors that such nets make have been shown in several cases to be similar to those that humans make. It is clear why this approach has aroused excitement.

Churchland argues that these models provide a plausible and exciting alternative to sentential models, and that they yield many interesting consequences for other problems in philosophy and cognitive science. Much of the week of his visit was focused on discussion of these further consequences, when he wasn't taking bi-pictures.

Dr. Churchland's Selfridge Lecture, *Perception, Understanding, and Action: A Neural Network Account*, presented the case against sententialism, and introduced the connectionist-neuroscience alternative. He argued the many advantages that neural net "concepts," or classifiers, seem to have as a ground for accounts of mental phenomena.

One problem that emerges with connectionist accounts of cognition is that the process of training a net is very slow, and this makes it apparently difficult for the approach to handle fast conceptual change—such as theoretical insight in science. In *Some Further Thoughts on Learning and Conceptual Change*, a presentation for the Cognitive Science group, Dr. Churchland addressed this problem, suggesting that a solution might be the fast "re-deployment" of concepts already learned in old domains into new domains. This point connects strongly with accounts of the role of analogy and metaphor in scientific development.

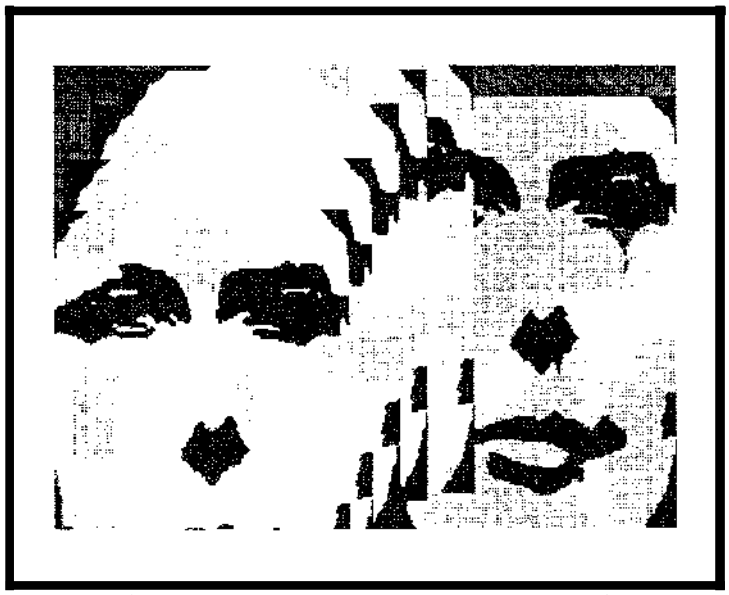

(computerized image by Michael Keough)

In a philosophy of science class, Dr. Churchland argued, in *A Deeper Unity: Some Feyerabendian Themes in Neurocomputational Form*, that the neurocomputational model of cognitive activity and theoretical understanding provides strong support for a number of Feyerabendian themes of the philosophy of science, such as the theory-ladenness of perception, the incommensurability of competing paradigms, the displacibility of folk psychology, and the desirability of the proliferation of theories and methodologies in science.

In *A Neural Network for Fast Stereo Vision*, Dr. Churchland presented a network model that can identify relative planes of the convergence of stereo vision. It was fascinating to see how this network architecture could pick out parts of stereo scenes that were in front of or behind other parts.

There were also a number of props for this talk—a set of hand-held stereo viewers. These devices illuminate two pictures of a scene, taken from some distance apart which gives the stereo baseline—this is the distance between the eyes in normal stereo vision—and provide a viewing of those pictures, one in each eye. The stereo pictures are fused in the mind, providing stereo vision of the original scene.

The most interesting of these devices was one that Dr. Churchland had built himself for stereo viewing of slides that he had taken of many scenes around the world. If the stereo pair of slides are taken with a larger than normal separation—larger than the usual separation between the eyes—then the experience of viewing them is like looking at the scene with eyes that are feet or hundreds of feet apart, depending on how far apart the original slides were taken. An aerial view of San Francisco with a baseline of thousands of feet, and a wide baseline view of Yosemite valley, were two of the fascinating examples. The stereo views obtained in this way provide visual experiences that would not be possible otherwise.

Meanwhile, Dr. Churchland met on almost every day with the Philosophy Faculty Seminar where, in spite of all the nice things we have to say about Dr. Churchland's position above, we endeavored mightily to get him to see the errors of some of his ways. Unfortunately, it seems that all these hours and efforts at philosophico-therapy were to no avail: at last report, Dr. Churchland was still furtively taking double slides of things, not quite identical, whenever he thought he could get away with it.