A NEW PARADIGM FOR ARTIFICIAL INTELLIGENCE

H. Joel Jeffrey

ABSTRACT

A new paradigm for artificial intelligence is presented that involves treating the computer as a behaving person. Descriptive Psychology provides the formulation of Persons, Behavior, and the Real World in a systematic, interrelated way that makes possible such an approach to the field of artificial intelligence. In the reformulation of the general problem of artificial intelligence, the mechanistic model is replaced by one in which the computational process becomes an instance of the Performance parameter of Intentional Action. The central task of getting the computer to recognize instances of concepts that cannot be reduced to computations is accomplished by a judgment space technology invented by Ossorio. The technology is described, and its use in research on automated information retrieval and several other topics in artificial intelligence is illustrated.

Advances in Descriptive Psychology, Volume 1, pages 177–194

ISBN: 0-89232-179-2

T. S. Kuhn, in *The Structure of Scientific Revolutions* (1970), applied the concepts of paradigm and paradigm shift to revolution in scientific fields. A field's paradigm is basically the world view that defines the field. It includes the set of standards accepted by members of the community about the world, what place their endeavor has in that world, what constitutes legitimate techniques and answers to questions, and perhaps most importantly, what constitutes a legitimate question. Sharing the paradign is the *sine qua non* for being a member of the scientific community whose paradigm it is.

The concept of phlogiston, for example, has literally no place in the practices of the modern chemistry community. Similarly, modern astronomers do not concern themselves with questions involving epicycles in the courses of the planets.

The paradigm for artificial intelligence (AI) has been that both a human and a computer are information processing devices: A person receives "sense impressions" from "the world"; these sense impressions are interpreted to produce what we "see," "think," etc., and sometimes further processing results in an output from "the system," which may or may not alter the environment. The information processing is of course extremely complex, but is basically describable in terms of powerful heuristics for handling such processes as tree-searching. The paradigm is well stated by Newell and Simon (1972), Minsky (1968), Uhr (1973), and (in rather a different context) Ossorio (1971/1978).

PARADIGM FAILURE

A paradigm can fail. Paradigm failure means that the members of the community involved are unable in some significant ways to treat the world as being what the paradigm says it is. When a field's paradigm has failed, the field is (by definition) in a state of crisis. The resolution of the crisis is the shift to a new paradigm. The rise of quantum mechanics early in this century is an excellent example. A new paradigm must be adopted, or the field ceases to be a scientific endeavor, for the existence of a paradigm is a key difference between science and other human activities.

In artificial intelligence (AI), the mechanism paradigm has been the only one that AI researchers have been able to see as providing any basis for scientific work. While non-mechanistic descriptions have at times been proposed, they have not been seen by the scientific community as scientific accounts of human behavior (Dreyfus, 1972). On a wider scale, the mechanism paradigm is the view held by almost all of the scientific community in the Western world (Dreyfus, 1972; Ossorio, 1971).

In recent years there has been considerable debate over the legitimacy of AI as a field of scientific endeavor. The field has been attacked as having had no significant successes, and being based on a fundamentally deficient view of human nature (Dreyfus, 1972). Of course, practitioners in the field have responded vehemently to these attacks.

This paper presents the view that the debate over AI may appropriately and fruitfully be treated as a dispute over the viability of the mechanism paradigm as a basis for AI. From this perspective, the attacks on AI may be seen as claims that the paradigm for AI has failed. Since critics have offered no alternative paradigm that AI researchers have been able to see as a viable foundation for their work, it is not surprising that the attacks have failed: The AI community continues to do research, publish papers, hold conferences, attract Ph.D. students, obtain funding, etc., with no significant change in the way it conducts itself.

As Ossorio (1971/1978a, 1971/1978b) has discussed extensively, the mechanistic paradigm functions adequately in the "hard" sciences, but has some serious conceptual inadequacies as a basis for a science of human behavior. AI is directly concerned with the world of persons and human behavior. Thus, seen from this perspective, AI's paucity of significant results is not surprising, and, more importantly, does not appear to be simply a matter of practical difficulties that can be expected to be solved. The point of this paper is to present a new paradigm for AI, which makes possible a science of AI without having to try to treat humans as mechanisms.

THE BEHAVIORAL PARADIGM

Redescription, Not Reduction

It is possible to argue that the mechanistic paradigm is necessary to have a science of artificial intelligence at all. That argument, in very brief form, is roughly as follows: A computer is a mechanism. Therefore, if we have a computer which behaves as a human (though perhaps one with certain physical handicaps), then we would seem to have reduced human behavior to machine processes, for any behavior would have its equivalent machine process. The only alternative would seem to be some form of "ghost outside the machine." If one sees people as machines (as the mechanism paradigm holds) this argument seems compelling.

The key to resolving this apparent dilemma is not to start out attempting to treat persons as nothing but mechanisms. Let us take a common place (in the human world) event, and examine the logic of describing that event. Consider a person making the opening kickoff of a football game. A full description of that behavior includes a specification of all of the eight parameters of Intentional Action. One of those parameters is the Performance parameter. An observer may redescribe the Performance in a number of ways. Some of those descriptions may include physiological processes, objects, events, and states of affairs. In particular, the observer could include processes, objects, events, and states of affairs in the person's brain (neurons firing, signals crossing synapses, serotonin levels, etc.) The observer could, in principle, give a description of what is happening, physiologically, when any part of the *behavioral* process of making the opening kickoff is taking place.

This certainly does not imply that an opening kickoff *reduces* to a set of physiological processes; it means only that when a person engages in this action certain other states of affairs are also the case. In particular, it does not mean that the physiology is what is "really" happening. (This is discussed extensively in Ossorio [1971/1978b].) In other words, this is an example of the (logical) fact that physiological processes that occur when a Person engages in some Intentional Action are exactly that: processes that take place when a Person behaves.

Now let us consider a different example: a program to do medical diagnosis, such as the thyroid-diagnosis program described in Johannes (1977). A list of a patient's characteristics are the input to the program, and a diagnosis is the output. The program's diagnoses have been judged, by a panel of qualified physicians, to be competent diagnoses. Thus, the program may appropriately be said to map sets of characteristics in relation to diagnoses. However, notice that the program may also be described as (a) a sequence of changes in the numerical values of variables in the program; or (b) a sequence of changes in the physical state of various of the components of the machine on which the program is running. None of these descriptions is incorrect. *Neither do any of these descriptions disagree with any of the others.*

Notice that there is no feature of the physical states of the machine which makes these states represent numbers and instructions, and there is no feature of the values of the variables which makes them represent characteristics and diagnoses. The descriptions given are different parameters of the Intentional Action which we can successfully treat the machine as engaging in. In other words, the programmer has designed a process such that we may successfully treat the results of the numerical process as a case of medical diagnosis. The same principle exemplified in the kickoff example may be seen here: There is no implication that diagnosis reduces to, or "is really" numerical calculation. Examples of this principle are common in everyday applications. Consider a program calculating checking account balances for bank customers. There is nothing about the calculations to imply that the program is "really" doing this. Rather, the programmer has written the program so that the numerical process may successfully be treated as a case of computing the balance. That it *is* the balance (or is not) is a fact (state of affairs) in the human world, not a fact about numbers in the program.

Treating a Computer as a Person

Let's look now at a third example: Stipulate a computer that passes "Turing's Test"—i.e., a person who did not know in advance would not be able to distinguish between the program and a human being (Jackson, 1974). This means that we may successfully treat the computer as engaging in those Social Practices that we expect a paradigm case Person to engage in (language, problem solving, negotiation, etc.), although again perhaps with some physical handicaps.

While engaging in these Practices, a variety of physical, electronic, and numerical processes and states of affairs will occur. A number of those might take place within what we could appropriately call the "brain." Thus, while the machine talked, laughed, argued, passed the time with friends, wrote letters to the editor, etc., some number of "brain" processes would be taking place. And just as with medical diagnosis and checking account balancing, having these processes go on while the computer is engaging in these Practices does not mean that any of the social practices have been reduced to electronics, physics, or numerical calculations. When ordinary persons engage in various social practices, a variety of physiological things may happen (recall the kickoff); when this stipulated computer engages in various social practices, different "physiological" things happen. What counts for us about a Person is the Social Practices he engages in, not the concomitant physiology.

Ossorio has amply demonstrated that the fact that humans have brains which are physiological mechanisms in no way implies that humans are mechanisms, or that behavior is physiology (Ossorio, 1971/1978b). The point of this example is that the same relation holds for computers: Having a person with a computer for a brain does not imply that that person's behavior is physiology, or that that person is a mechanism. A human with a computer for a brain is exactly that: a human, with an unusual brain. Just as there is no *logical* problem in having humans with protoplasmic brains, there is none in having humans with electronic brains, and there is no "ghost outside the machine" in either case.

The apparent dilemma of AI has been resolved, by moving from the machine concept to the Person concept, and examining the Intentional Action formulation of the Behavior of a Person. What we have been doing here can be seen as a case of treating the computer as a behaving person. The concept is the new paradigm for AI: Treat the computer as person behaving in the world.

Research within this paradigm—i.e., acting on this concept—is endeavoring to create programs that we may successfully treat as Persons behaving in the world; the ultimate standard is the degree to which a given program may be so treated. There is no question of trying to reduce persons to mechanisms or behavior to computation; there are many questions involving how to build programs whose behavior we may appropriately describe as the behavior of a person. The issue of how to build such programs, and in particular the question of how to get the computer to do things that do not reduce to computations, is the subject of the next section.

A TECHNIQUE FOR JUDGMENTS

If treating the computer as a behaving person is to be successful as a paradigm for AI, we must answer the question posed at the end of the last section. If we cannot, then the whole enterprise is legitimately subject to the criticism that, while it may make sense to describe the machine as acting on concepts, if the only behaviors actually available to the machine are equivalent (to us) to computations, then there is little point in talking that way.

A concept will not, in general, reduce to some other concept. Instances of a concept may have nothing in common, other than being instances of the concept. How then will we program the computer (which after all can only calculate) to do things that we can appropriately describe as acting on concepts, and not just manipulating numbers?

Ossorio developed a technology that we can use to meet this need. In the original study (Ossorio, 1965) he dealt with the problem of having the computer make judgments of subject-matter relevance. The technique was called a Classification Space. In later work he presented Category Space for judging the category a thing fits into, Property Space for judging properties, Functor Space for judging significant dimensions of variation, and Means-End Space for judging how well a given means is suitable for achieving a given end (Ossorio, 1965, 1966, 1971a/1978).

The original publications (Ossorio, 1965, 1966) present the technique in detail. Rather than repeat that detail, my presentation is designed to provide a preliminary grasp of the procedure and to give those with relevant problems some reasons for trying to use the technology. Perhaps the most important reason for using the new procedures is that they are consistent with and were in fact derived from the behavioral paradigm of Descriptive Psychology. For didactic reasons, the first illustration will be based on judgments of subject matter relevance, but, as we shall see, such a starting place in no way limits the implications of the presentation for the general problem of AI.

In order to simulate human judgment, one must develop a Judgment Space. The first step in such a procedure is to notice that, while in general real-world knowledge is not deductively or mathematically related, human users are able to act on such knowledge. Further, this knowledge, whether factual and certain or fuzzy, vague, and tentative can be represented using numbers. We have a commonly used set of locutions to indicate clarity, degree of applicability of a concept, etc. We introduce an alternate set of locutions by using a numerical scale (e.g., 0–10), and use the highest rating to represent certainty, lowest to represent uncertainty, and intermediate values for intermediate uncertainty. For example, 1 out of 10 represents a case where some description is not totally false, but is extremely far-fetched.

Having represented enough knowledge in some area this way, it is in general possible to make new judgments by combining the values representing the original knowledge. Let us go through the derivation of a Judgment Space (or J-space) for making subject matter relevance judgments.

Step 1. Select the fields of interest. (If this were an attribute-judgment space, one would select the attributes of interest; if this were a concept-recognition space, one would select the concepts of interest.)

Step 2. Select a set of words or phrases from the subject matter fields of interest. (In the case of concepts, select examplars of each concept.)

Step 3. Putting the fields F[l], . . . F[n] across the top, and the vocabulary v[l], . . . , v[n] down the side, we have a (empty) matrix. This is the judgment matrix. Fill it, with judgments of the degree to which each v[i] is relevant to each F[j]. These judgments are obtained from human judges competent to make them. We ask the judges to express their judgments numerically (i.e., using the numerical locutions), as follows:

1. Irrelevant. This term really has nothing to do with this field. Rate 0.
2. Marginal. This term could be said to be relevant, but only in a tangential or farfetched way. Rate 1 or 2.
3. Peripheral. The term has some relevance to the field, but is basically peripheral to it. Rate 3 or 4.
4. Relevant. This term is definitely relevant to this activity. Rate 5 or 6.
5. Highly significant. This term is highly relevant to the field; it is a key concept in the field, or relates directly to several critical concepts. Rate 7 or 8.

Within each category, the rating is higher when the relevance is higher.

Step 4. We now have a filled-out judgment matrix. It is very difficult to use this matrix as it stands, because at this point the numerical locutions may not represent the human world well. Consider the following example. Suppose we had three subject matter fields, Computer Software, Computer Hardware, and Zen Buddhism, and three terms with ratings:

	CS	CH	Z
T[1]	8	0	0
T[2]	0	8	0
T[3]	0	0	8

Using only the numerical information here, T[1], T[2], and T[3] are equidistant. But in the real (human) world, software is certainly more closely related to hardware than either is to Zen. So, the numbers are not representative of the real-world situation.

In actual cases of judgment matrices, we have a large sample of terms from each field. Since some fields are more closely related, this means that some columns of the matrix will be more closely correlated. What we would like is to have another representation of the data, which represents the information in the matrix in terms of independent "types of fields" or "types of content." Since a high correlation between two columns represents high subject matter similarity (because of the sampling of the fields), this means that we would like to have a representation of the judgment matrix in terms of groups of columns, such that different groups are independent and the fields within a group are highly correlated. This is precisely the result produced by intercorrelating and factor analyzing the judgment matrix (Comrey, 1973).

Therefore, we get the desired orthogonal basis by intercorrelating and factor analyzing the judgment matrix. The common factors, made up of highly correlated F[i], and unique factors, which are those F[i] having no significant content in common with any other field, are an orthogonal basis for the Judgment Space. (The reader is referred to [Ossorio, 1966] for a detailed description of the methods of factor extraction and rotation used.)

Step 5. The factor analysis produces numbers, called *loadings,* which relate the F[i] to the factors; the loading is the cosine of the angle between the vector F[i] and the factor. The factor may be viewed as a combination of those F[i] with a loading of over 0.7 (approximately the cosine of a 45-degree angle).

This step is done as follows: Suppose we had factor loadings like

	Factor I	Factor II
F[1]	0.9	0.1
F[2]	0.8	0.6
F[3]	0.2	0.9
F[4]	0.1	0.9
F[5]	0.3	0.7

Supposing that term v[1] were rated 6 with respect to F[1] and 7 with respect to F[2], and 0 with respect to all other fields, the rating on Factor I would be $0.9*6 + 0.8*7 / (0.9 + 0.8) = 6.5$. The rating on Factor II would be 0, since F[1] and F[2] are not used to measure the value for Factor II, since their loadings are less than 0.7 and thus they are in a direction more than 45 degrees away from Factor II. (Readers familiar with factor analysis will recognize this as computing the factor scores.)

Up to this point the mathematical procedures have been standard factor-analytic ones (or close variants). We have used factor analysis to produce a vector space with an orthogonal basis, each of whose basis vectors represents a distinct type of content, and populated the space with the vocabulary items. At this point we leave factor analysis and simply use the vector space.

Step 6. The result of the above step is that the set of vocabulary terms v[i] is located in the relevance space. (In the general case, the items, objects, or whatever would be located in the Judgment Space.) The hallmark of judgment, though, is to be able to judge novel cases. This is simulated by using known objects (the terms, in the relevance space case) to simulate judgment of new items: documents. When a document is to be located in the Judgment Space, which is a case of judging its relevance, it is scanned for terms recognized. Suppose we have K terms. The locations of the recognized terms are a set of points in the Space, p[1], . . ., p[K]. To judge the document's relevance we need to calculate its location in the Space. This is done by combining the K locations mathematically. In my work, a log average has been quite successful. For example, in order to calculate the value of the jth axis, one obtains the value, q[j], as given by:

$$q[j] = \log_b \left((b^{p[1,j]} + \ldots + b^{p[K,j]}) / K \right) \qquad (1)$$

(Recall that the axes are the common and unique factors of the judgment matrix.) In other words, the value of axis j is the log average of the values of each of the terms on that axis. This formula gives higher weight to more highly relevant terms, which appears to fit the facts of human judgment.

For example, a document in which ten terms were recognized but five were judged to be as highly relevant as possible, say eight on zero through eight rating scale, and the other five were judged to be entirely irrelevant, say zero on the scale, would receive a higher relevance scale than the simple average of the rating of these ten terms.

The function being componentwise (i.e., for the new value on axis j, we use only other values *on axis j*) illustrates the correspondence between the mathematics and the non-mathematical use of it: We have an orthogonal basis for the space. Mathematically, this means that the values on each axis are independent of values on the other axes. As noted above, each axis represents an independent type of content. One would not judge relevance to one type of content by examining relevance to entirely independent, unrelated content.

Step 7. Documents (and terms) may now be compared for conceptual content similarity by calculating their distance in the Judgment Space. A variety of metrics is of course possible; I have had good results using the standard Euclidean metric (Jeffrey, 1975). By using a metric, the Space may be used to retrieve documents by treating a retrieval request as a document, locating it in the Space, and then retrieving those documents, in order of closest document first. Since the axes of the space represent types of content, and the value on each axis represents the degree to which a document has that type of content, a document is mathematically close to another (or to the request) precisely when it is close in conceptual content.

A cautionary reminder may be useful here. It is tempting, if one is still operating in the mechanism paradigm, to view Judgment Space technology as probabilistic reasoning, number-based inference, etc. To a certain extent it can be seen that way, but doing so misses the point: A program using a Judgment Space is doing something that we may successfully treat as a case of making judgments.

The question raised at the beginning of this section was how we could have the computer act on concepts that do not reduce to computational processes. By gathering numbers which are instances of acting on concepts (numerical locutions), and manipulating the numbers so that that relationship is maintained, we arrive at a mathematical object (the vector space together with the combining function) such that we may appropriately treat the results of calculating that function, in those cases, as acting on the concept. The computer (viewed as a machine) still only calculates; but when it calculates with these numbers, in this way, we can view it as acting on concepts.

RESEARCH IN THE BEHAVIORAL PARADIGM

This section discusses some examples of research in the field of AI as defined by the Behavioral Paradigm. Some of the work presented has been done, and some is proposed work. The examples cited are illustrations of what appear to be interesting and fruitful ways of acting on the concept of treating the computer as a behaving person.

Subject Matter Relevance

When a person retrieves information in response to a request, he makes a relevance judgment of the request—i.e., the relative importance, within various domains of human activity, of the information. This judgment is one criterion which can be, and is, used in practice.

Ossorio (1966) constructed a Judgment Space with the ability to make subject matter relvance judgments of documents and requests. He began with 24 fields from science and engineering and 288 technical terms from these fields. The factor analysis yielded 6 common factors. When the space was used to index and retrieve documents, the correlations of system ranking with human judge ranking ranges from .896 to .984 (Ossorio, 1966).

I implemented a complete document retrieval system based on this approach (Jeffrey, 1975). A relevance space covering 62 fields and specialties within computer science was constructed, using 800 technical terms. The system behaves just like a competent human librarian in a computer science library. When responding to a request which is within the range of content covered and that uses vocabulary for which judgments are present, the system achieves an average recall of 75–85% of the relevant documents *simultaneously with* an average precision of 80–90%. These results are a very significant improvement over results obtained by the usual techniques of word-matching, in which in almost every case the recall percentage plus the precision precentage total 100%—i.e., 30% recall—70% precision, 80% recall–20% precision, etc.

It is interesting to see that the retrieval system has the same limitations a human librarian does. If a user states his request in terms a librarian does not know, the librarian will not understand it; the same holds for this system. Subdivision of fields was not implemented. Thus, requests about game playing programs would result in retrieval of documents on natural language understanding, since both are topics within artificial intelligence. Again, this is exactly what a human librarian who had no knowledge of the subdivisions of AI would do—the only judgment available would be relevance to artificial intelligence.

Significant Feature Selection

It is an accepted fact, within the AI community today, that one ability that humans have and machines currently do not is the ability to judge what features of an object or situation should be examined. This is what Dreyfus (1972) terms "zeroing-in."

Lack of this ability leads to the necessity for searching a tree of possibilities. In some cases this has produced some reasonable results; the cases are those in which the possible states of affairs are simple enough, or have enough mathematical structure, that techniques such as tree pruning and alpha-beta searching are not too poor a substitute for judgment. (Jackson [1974] discusses tree searching techniques.) In checkers and, to a degree, chess, this has been so. In other cases, such as the game of Go, the results have been dismal.

It appears that we could obtain some very interesting results in the area of intelligent game-playing by using Judgment Spaces to provide a program with the ability to (a) judge what features are significant about a situation, and (b) judge the degree to which a situation has the features or properties of interest.

These abilities would be provided by a Functor Space and a Property Space, respectively. To construct a Property Space, the columns of the matrix represent properties of interest, and the rows represent objects. The judgments are the degree to which each object has each of the properties. A new object (for example, a new board position in a chess game) is located in the space by identifying subobjects, or related objects (for example, already-recognizable features or other board positions) and combining the positions in the space of those objects. A Functor Space is the result of starting with a list of significant features or dimensions of variation, which are represented by the columns. The judgment is the degree to which each dimension D is a significant dimension of variation of each object—i.e., the degree to which it is important to know D about object X. This directly attacks the zeroing-in problem.

Ossorio (1965) constructed both of these spaces, and reported that there were no difficulties in doing so. There has not, to my knowledge, been a game playing program constructed using this approach.

Medical Diagnosis

Johannes (1977) addressed the problem of thyroid disorder diagnosis by simulating the judgment of qualified physicians. The system takes in a set of patient characteristics. It uses a Diagnosis Space to make an initial diagnosis. It then uses a Test Space to make recommendations of tests to be done. The Diagnosis Space is then used to revise the initial diagnosis in light of the known test results. As an attempt to simulate competent physicians, Johannes' program is highly successful: A panel of 7 physicians reviewed the system's diagnoses on 15 cases. The panel agreed with the system's initial diagnosis in 98.1% of the cases, the test recommendations in 91.4%, and the final diagnoses in 92.4% of the cases.

Chess

Chess has long been recognized as a paradigm case of human behavior. Playing good chess (Master level and above) requires the ability to recognize patterns, judge what is important in a position, whether a position should be examined further, pick appropriate goals, pick appropriate courses of action, and in general act on a great variety of chess concepts that do not reduce to any physically definable set of characteristics of boards and pieces. Currently, chess programs are limited in just that way—they can only deal with reducible attributes, not concepts. By giving the computer the ability to recognize instances of, and act on, chess concepts, we can construct a program that plays chess like a human does, i.e., by recognizing and acting on concepts. (How well it plays is a separate issue, just as it is for humans.)

Such a system would operate as follows:

1. When a position is presented, the Strategy Space returns the name of the strategy to use. (A strategy is a Process, and thus is described by a Process Description. The reader is referred to Ossorio [1971/1978a, 1971/1978b] for detailed discussions of Process Descriptions and how they may be used.)
2. The Attribute Space and Tactics Space are used to recognize instances of non-computable concepts and select the tactic(s) best suited to the strategy in this case.
3. Using common board position analysis techniques (Jackson, 1974), and probably a Move Space, a move is selected.

The strategy is what is being done—not an abstraction of reality as strategies have usually been viewed. Selecting a particular tactic is an instance of engaging in the strategy, for the strategy is a process that is made up of stages and options such that the selection of a move is the exercise of a particular tactic.

Problem Solving

The monkey-and-bananas problem (Jackson, 1974) is a standard toy problem for illustrating reasoning: A monkey is in a room where a bunch of bananas is hanging from the ceiling, too high to reach. In the corner of the room is a box, which is not under the bananas. How can the monkey get the bananas?

This is a "toy" problem because of its size—the number of facts about the situation, objects, and available actions is very small. A predicate calculus formulation of the problem requires 11 axioms, and a proof that the monkey can get the bananas can be given in 13 lines, starting from the axioms (Jackson, 1974).

A human is in a room where a bunch of bananas is hanging from a 20-foot-high ceiling. The room is much like an ordinary living room; it has a couch, a straight wooden chair, a wooden table, a table lamp, a pole lamp, and a 4-foot-square rug on the floor. On the table are a box of 20 drinking straws, a pile of 100 rubber bands, 5 toy balloons, 3 pencils, and a roll of wire. On the rug are a toy truck, 6 paperback books, a cardboard box for toys, and a floor lamp. How can the human get the bananas?

This example is far out of the range of toy problems. The number of different objects (not counting the 100 identical rubber bands, etc.), the number of properties of each, and the number of actions that each is suitable for would result in an enormous number of axioms if the problem were formalized. Even more important, this problem has a whole range of problems not even present in the toy version: Which facts should be represented? For example, a table has a certain size, shape, and weight. It is suitable for a place to put objects, work at, etc. Less commonly, it could be climbed on, sat on, etc. Straightforward, so far. However, it is also a physical object, and so may be decomposed in various ways—legs, top, etc. Further, depending on its composition, it might be that the top could be broken into long sticks. The same situation holds for many of the objects in the room. Trying to represent the facts and the redescriptions leads to a hopeless combinatorial explosion.

But a human does not face these problems; he reasons with the facts he sees, and (depending on ability) acts on redescriptions if necessary. The following is one way a person might act in the given situation:

1. Decide to try climbing.
2. Stack chair on table, and climb on top.
3. Notice bananas are closer, but not yet in reach.
4. Decide to hit bananas from top of stack.
5. Take apart floor lamp, getting 6-foot-long center pole.
6. Notice this is not long enough to reach the bananas.
7. Wire the lamp pole to the table lamp.
8. Climb up, hit bananas with extended pole.

Let us examine this sequence, using a question and answer format to pinpoint the judgments being made:

1. Q. What known actions look good for getting the object out of reach?
   A. Climbing.
2. Q. Does climbing require any props?
   A. Yes—an object tall enough to help, and which can be climbed on.
3. Q. Any such objects present?
   A. No.
4. Q. Are there any known methods for creating tall climbable objects?
   A. Yes—stacking objects.
5. Q. Does stacking require any props?
   A. Yes—at least 2 objects that can be lifted, one of which must have a flat top.
6. Q. Are such objects available?
   A. Yes—table and chair.
7. Q. Can bananas now be reached?
   A. No.
8. Q. Does this approach look reasonable, or should you start over?
   A. Reasonable, keep stack for now.
9. Q. From the top of the stack, what known actions look suitable for getting the object out of reach?
   A. Hitting object.
10. Q. Does hitting require any props?
    A. Yes—a stick long, strong, and light enough to be lifted.
11. Q. Such a stick available?
    A. No.
12. Q. Any known methods for creating objects from other objects?
    A. Yes—putting objects together, and taking them apart.
13. Q. What objects have long, strong, light parts?
    A. Floor lamp has long center pole. Table has legs. Chair has legs.
14. Q. Center pole of lamp long enough?
    A. No.
15. Q. (Repeat 12.)
    A. Yes—putting objects together and taking them apart.
16. Q. Any objects suitable for putting together with lamp pole?
    A. Yes (marginally)—the table lamp.

17. Q. Any known methods for putting objects together?
    A. Yes—tieing, gluing, nailing, screwing, bolting.
18. Q. Does tieing require props?
    A. Yes—string.
19. Q. Any string present?
    A. No.
20. Q. Any objects with similar relevant properties?
    A. Yes—wire.
21. Q. Is new object (table-lamp-and-lamp-pole) long enough?
    A. Yes.

Now notice that *every one* of the above steps in this complicated piece of problem-solving behavior can be implemented by either simple lookup in Object or Process Descriptions, or via one or more Judgment Spaces. (Object Descriptions are also discussed in [Ossorio, 1971/1978a, 1971/1978b].) Further, subobjects and combinations of objects need not have any location in the Spaces in advance. Step 13, for example, involves Property and Functor Spaces; Step 19 uses Relevance and Property Spaces; Steps 1, 4, 9, and 12 use a Means-End Space. A Means-End Space is a Judgment Space in which the columns represent means, the rows represent goals, and the judgment is the degree to which each means is suitable as a means to each end. This is discussed in Ossorio (1965).

Finally, it is of interest to see how the problem of combinatorial explosion, which has long been recognized as *the* primary problem in AI, simply does not arise here. In Step 13, for example, the floor lamp was selected for dismantling by the (hypothetical) system. It was selected on the basis of being the most highly rated object in the Judgment Spaces, at each stage which required a judgment. Since a system operating with Judgment Space is reproducing human judgments, the system will make several attempts, or have several alternatives to consider in some stage, just when a human does: when the knowledge does not indicate a clear choice. In terms of the Judgment Space operation itself, this would be the case, for example, if several alternative methods were rated 4 (indicating "could be suitable, but you wouldn't normally think of it for this goal'").

Automatic Fact Analysis

The automatic fact analysis problem is the problem of producing an automatic system for analyzing the implications of facts. A paradigm case is the problem of analyzing military intelligence. It is in some sense the supreme AI problem. All of the difficulties of traditional AI must be faced in attacking it, the worst being the problem of how to handle real world knowledge (Jackson, 1974). Certainly the best research with the Behavioral Paradigm is the State of Affairs Information System (SAIS) designed by Ossorio (1971/1978a). The SAIS forms a complete package for operating with Object and Process Descriptions, including Judgment Spaces for the places where human fact analyzers exercise judgment. What Ossorio did was to analyze, in terms of the Person Concept, what it is to do fact analysis, and then use that analysis to design a system to reproduce those achievements. This system has not yet been built. In my judgment, some of the most fascinating and significant research in the near future will be the implementation of a State of Affairs Information System.

CONCLUSION

A new paradigm for artificial intelligence has been presented: the Behavioral Paradigm. Whereas with the mechanistic paradigm one attempts to treat a human as an information processing mechanism, and tries to describe behavior by computational processes, with the Behavioral Paradigm one treats the computer as a behaving person, and constructs behavioral models for computational processes. In order for this approach to be viable as a scientific paradigm, one must have a precise, systematic formulation of the concepts of Persons, Behavior, and the Real World. Descriptive Psychology is that formulation. It is also necessary to have a technique by which the computer can deal with descriptions of parts of the real world, without having to replace them with others of a computable form. The technique for having the computer do non-computable things is the Judgment Space. The Behavioral Paradigm is thus a new concept of the computer, which is scientifically useful. As such, it constitutes a new paradigm for the science of artificial intelligence.

ACKNOWLEDGMENTS

I wish to acknowledge the helpful comments of Thomas Mitchell, Anthony Putman, and Paul Zeiger on earlier versions of this paper. The material contained in the paper was presented at the meeting of the Society for Descriptive Psychology, Boulder, Colorado, August, 1979. Author's address: Bell Laboratories, Naperville, Ill. 60540.

REFERENCES

Comrey, A. L. *A first course in factor analysis*. New York: Academic Press, 1973.
Dreyfus, H. *What computers can't do*. New York: Harper and Row, 1972.
Jackson, P. *Introduction to artificial intelligence*. New York: Petrocelli/Charter, 1974.
Jeffrey, H. J. *Information retrieval by conceptual content analysis* (Technical Report 75–6). Nashville, Tenn.: Computer Science Department, Vanderbilt University, 1975.

Johannes, J. D. Automatic thyroid diagnosis via simulation of physician judgment. (Doctoral dissertation, Vanderbilt University, 1977). *Dissertations Abstracts International,* 1978, *38,* 3773B. (University Microfilms No. 77–30, 364).

Kuhn, T. S. *The structure of scientific revolutions*. Chicago: University of Illinois Press, 1970.

Minsky, M. (Ed.), *Semantic information processing*. Cambridge, Mass.: MIT Press, 1968.

Newell, A. & Simon, H. A. *Human problem solving*. Englewood Cliffs, N.J.: Prentice-Hall, 1972.

Ossorio, P. G. *Dissemination research* (RADC-TR-65-314). Rome Air Development Center, New York, 1965.

Ossorio, P. G. Classification space. *Multivariate Behavioral Research,* 1966, *1,* 479–524.

Ossorio, P. G. *State of affairs systems* (LRI Report No. 14). Whittier and Boulder: Linguistic Research Institute, 1978. (a) (Originally published in 1971 as (RADC-TR-71-102) Rome Air Development Center, New York.)

Ossorio, P. G. *"What actually happens"*. Columbia: University of South Carolina Press, 1978.

(b) (Originally published in an earlier version in 1971 as LRI Report No. 10a. Whittier and Boulder: Linguistic Research Institute.)

Uhr, L. *Pattern recognition, learning, and thought*. Englewood Cliffs, N. J.: Prentice Hall, 1973.