

PLAY CONSIDERED AS A STRATEGY FOR KNOWLEDGE
ACQUISITION

Paul D. Scott
Department of Computer and Communication
Sciences
University of Michigan
April 1982

1: INTRODUCTION

If you ask a layman what he means by the term 'play' he will probably reply "activities which are useless but fun" or something very similar. If you ask a developmental psychologist the same question you will probably get much the same answer although he is likely to phrase it differently:-

"Play consists of behaviors and behavioral sequences that are organism dominated rather than stimulus dominated, behaviors that appear to be intrinsically motivated and apparently performed 'for their own sake' and that are conducted with relative relaxation and positive affect."

Weisler and McCall (1976)

Patterns of behavior which appear to have no external purpose but are nevertheless enjoyable for the participant present something of a biological paradox. The majority of activities which are accompanied by positive affect clearly promote, either directly or indirectly, the participant's homeostatic or reproductive goals. An adequate theory of play must resolve this paradox by attributing a function to play.

A number of theories have been advanced which attempt to do this by suggesting what the organism may gain by engaging in play. Space does not permit a discussion of the relative merits of these theories but see Weisler and McCall (1976) and Gilmore (1966) for reviews. Fortunately one particular theory appears to enjoy almost universal support. This we shall call the 'Cognitive Development Hypothesis'. Its basic premise is that the organism learns something through the process of play which is of value in later life. This theory has been advanced in a bewildering variety of forms which largely reflect the enormous range of things which a child learns. Taken together these various theories amount to a claim that play is the fundamental learning strategy by which children acquire mastery of themselves and of the perceptual, motor, cognitive and social skills which they will need throughout life.

The cognitive development hypothesis provides an explanation of the function of play and hence resolves the paradox. It is very widely accepted by developmental psychologists, primatologists, pediatricians and laymen. Strangely it has received little acknowledgment from learning theorists. Thus a large and reputable text on learning theory (Hilgard and Bower 1975) contains no index reference for play. Piaget does assign a relatively minor role to play in his model but regards it as a particular case of assimilation rather than a fundamental learning strategy. Play has been equally ignored by artificial intelligence

researchers interested in machine learning. I am not aware of any program which explicitly incorporates play as a learning activity although I think it would be fair to describe the behavior of AM (Lenat 1976) as playing with numbers. Otherwise AI programs seem to be based on the assumption that learning must be either a classroom experience (learning with a teacher) or an apprenticeship (learning while doing the task). This paper is intended to exhort both learning theorists and AI workers to take play more seriously.

Although the cognitive development hypothesis provides an explanation of the function of play it does not constitute a complete theory. Such a theory must provide an account of how play activity is instigated, motivated and rewarded. It must explain the content and structure of play activities. The cognitive development hypothesis only provides a framework within which more complete theories may be developed. The rest of this paper is devoted to sketching the outlines of one such theory.

2: A THEORY OF PLAY

If play is a method of building a cognitive representation then any theory of play must make some assumptions about the nature of the representation which is built. I therefore begin the development of the theory with the following postulate:-

Play is an activity directed towards building a representation of the world in terms of the organism's abilities to do things to or with the entities which it encounters in the world.

This hypothesis makes a strong claim about what is learned during play. It asserts that the organism is attempting to discover what it can do rather than what it should do. That is, it is not primarily concerned with learning what actions have desirable outcomes. It is of course possible, and indeed probable, that the organism will obtain information about what it should do as a side effect of trying to discover what it can do, but the claim made in the hypothesis is that such information is not the goal of play behavior. Note that this does not imply that the organism will not be trying to determine the consequences of its actions but only that it will not be directly concerned with the values of those consequences.

This form of representation in which the world is modelled in terms of how it relates to the organism's behavioral capabilities has some obvious merits. For example, it is an essential prerequisite for any kind of problem solving behavior since it enables the organism to generate alternative courses of action in a

given situation. However, since it is most readily understood in terms of simple motor responses to a given event there is a serious danger of underestimating its power and generality. It is therefore worth pointing out that it strongly resembles Gibson's notion of 'affordances' (Gibson 1977). It is also closely related to the pragmatic theory of meaning due to Peirce (1878) and subsequently elaborated by James, Dewey and Mead among others. For this reason we shall refer to it as a 'pragmatic representation'. Object-based programming languages such as SIMULA and SMALLTALK represent entities using what is essentially a pragmatic representation.

3: IMPLICATIONS OF THE HYPOTHESIS

We now explore some of the implications of the hypothesis that play is a strategy for building a pragmatic representation of the world. In executing an ordinary goal oriented task the organism is attempting to effect some change of state in its world. In doing this it uses knowledge of the properties of the world. In play the organism is attempting to effect some change of state within its own representation of that world. In doing this it will use knowledge regarding that representation. Thus it can be seen that the goals of play are metagoals and hence that play involves access to metaknowledge.

What sort of metaknowledge would be relevant for the development of a pragmatic representation? If the organism is to discover what it can do then it presumably needs to have some representation of what it does not know it can do. That is the metaknowledge must represent the organism's ignorance. Such a representation could be used to determine the course of play behavior. Thus the organism would in effect conduct experiments whose purpose is to reduce its own ignorance of its capabilities in a manner loosely analogous with scientific research.

The introduction of the concept of metaknowledge raises the spectre of an infinite regress. Where does the metaknowledge come from? Is it necessary to play at playing in order to discover how to play? The threat of an endless regress can be avoided if the same activities which provide information for the pragmatic representation of the world also provide the information needed to build a model of the organism's ignorance. This constraint is not only satisfiable but also explains one of the basic empirical findings regarding play and exploratory behavior: the probability that a child will play or explore is related by an inverted-U curve to the novelty of a situation. In a highly familiar situation the child will have a detailed pragmatic representation and correspondingly low ignorance and thus there is little to be gained by play. Conversely in a totally unfamiliar situation the child will have virtually no pragmatic representation and hence have no knowledge of its own ignorance. In such circumstances he or she would essentially not know how to play. Only in the intermediate case in which a partial pragmatic representation exists is the child able to construct potentially useful play activities.

4: A SIMPLE IMPLEMENTATION

In order to clarify the ideas discussed in the preceding section by providing a concrete example and to demonstrate that such hypotheses can indeed lead to a successful learning program, I shall now describe a very simple concept learning program which learns by playing.

The organism in this case is a LISP program called PAN. PAN operates in a simple blockworld type of environment. In this world are numerous objects which each have the properties of color, size, texture and shape. Each of these properties may take one of several discrete values. PAN is able to apply three types of action to these objects. It can push them, kick them and pick them up. However these actions will only result in the object moving in certain cases. For example the operation of picking up might only result in the object moving if the object were small. Initially PAN does not know what classes of objects its three kinds of action will succeed on. Thus PAN's task is to discover the equivalence classes of kickable objects, pushable objects and objects which may be picked up. It must experiment entirely without external guidance until it is confident that it can predict the applicability of an action to an object.

PAN does this by developing a class hierarchy. Initially it possesses only one class - the class 'Things'. All objects are instances of this class and all actions are initially attached to this class. As part of the attachment of an action to a class PAN stores an estimate of the probability that the action can be applied to a member of that class. This probability estimate is revised every time PAN tries to apply that action to an instance of that class. If this probability estimate is very large or very small then PAN is relatively certain about the applicability of the action to instances of the class and hence has no need to conduct further experiments. If however the probability is in the region of 0.5 then PAN is highly uncertain and further development of the class hierarchy is needed. The actual measure of uncertainty used in the system is Shannon's information function (Shannon and Weaver, 1949). In fact any function of the probability which was unimodal in the interval 0 to 1 with a maximum at 0.5 and a value of 0 at 0 and 1 would serve. The use of the Shannon function has the advantage of allowing one to interpret it as the informational value of a new subclass rather than being a meaningless number. The initial probability estimate assigned to each action for its attachment to the class Things is 0.5 and hence they each have an uncertainty of 1.

A cycle of the system is called an experiment. In each experiment the system finds the action which is attached to a class with highest uncertainty. An instance of that class is selected and the action applied. If the action is successful then, apart from the increase in the estimated probability, nothing else happens and the system begins a new experiment.

If the action is unsuccessful then one of two processes may occur: a new subclass may be created or the action may be detached from the class. If the uncertainty exceeds a certain

threshold then PAN will attempt to construct a subclass of the class in which the action has just failed and then attach the action to the new subclass. (The action remains attached to the original class). It does this by repeatedly selecting instances of the original class until it finds one on which the action succeeds. It then selects a random attribute of that instance, for example its color or its size, and uses that as the criterion for membership of the new subclass. The action is then attached to the new subclass with an initial probability estimate which is identical to the current probability estimate for the attachment of the action to the original class.

This newly created subclass may or may not contain a higher percentage of objects to which the action may be applied. If it does then the subclass is clearly useful and hence will be retained. If it does not then the probability estimate will eventually fall below that of the corresponding estimate in the parent class. When this happens the action is detached from the subclass. If any class has no actions attached then it is removed. Hence only useful classes are retained. In this way the system develops a hierarchy of classes as it attempts to reduce its uncertainty. The system is able to learn both conjunctive and disjunctive concepts and will eventually reach a stage when all uncertainties are below threshold. In this situation PAN announces that it is bored and halts. Note the system does not necessarily find a minimal set of classes to represent the concepts it is discovering. This could be done at the expense of more elaborate rules for modifying the hierarchy. It does however achieve a correct if redundant representation. In this respect its behavior resembles that of human beings.

The above account is simplified in one respect. Once subclasses have been constructed any given object may be an instance not only of a given class but also of one or more of its subclasses. Thus, if PAN is doing an experiment which involves applying an action to an instance of some class, the particular instance selected may also be a member of a subclass to which the action is also attached. In these circumstances the experiment is effectively transferred to the subclass which has the highest probability estimate. The result of the experiment modifies the probability estimates of both the subclass and the parent class. However a second probability estimate is also kept for each attachment which is a measure of the proportion of attempted applications which were not passed down to a subclass. This second probability estimate, called usage, is multiplied by the Shannon information function in determining the uncertainty. This is analogous to Shannon's measure of the entropy of an information source.

The reason for this modification is that if it were omitted the uncertainty of parent classes would remain high even when the appropriate subclasses had been constructed. This would lead to endless redundant experimentation. The modification described ensures that a class with successful subclasses will have low uncertainty values despite not having probability estimates close to 1.

As indicated earlier PAN is only intended as a demonstration that the play theory can be used as the basis of a learning system which works without the assistance of a teacher. We

are developing a much larger version of the system in which objects may possess relational attributes and actions may change those relations. PAN is however only a simple instantiation of the use of a play based learning strategy.

The pragmatic representation takes the form of a class hierarchy with actions attached to classes. The metaknowledge of its own ignorance takes the form of the associated uncertainties. The same experiments which lead to alterations in the pragmatic representation also change the representation of ignorance.

Because PAN operates in a very restricted universe it eventually learns all that can be learned. Generally we should not expect this to happen. As the pragmatic representation becomes richer the organism has more things to be uncertain about. Hence the process of building the representation becomes a never ending search for something even better while retaining the best that has been achieved so far.

5: REFERENCES

- Bruner, J.S., A. Jolly and K. Sylva 1976
 "Play - Its Role in Development and Evolution"
 Penguin Books Ltd., Harmondsworth, England
- Gibson, J.J. 1977
 "The Theory of Affordances"
 In "Perceiving, Acting and Knowing: Toward an Ecological Psychology" Eds. R. Shaw and J. Bransford, Erlbaum, pp 67-82
- Gilmore, J.B. 1966
 "Play: A Special Behaviour"
 In "Current Research in Motivation", Ed. R.N. Haber, Holt, Rinehart and Winston, pp 343-354
- Lenat, D. 1976
 "AM: An Artificial Intelligence Approach to Discovery in Mathematics as Heuristic Search"
 Doctoral dissertation, Stanford University, July 1976.
- Peirce, C.S. 1878
 "How to Make Our Ideas Clear"
 Popular Science Monthly, January 1878, pp 286-302
 Reprinted in "Charles S. Peirce: Selected Writings"
 Ed. P. P. Wiener, Dover, 1966
- Shannon, C.E. and W. Weaver 1949
 "The Mathematical Theory of Communication"
 University of Illinois Press, Urbana, Chicago, London
- Weisler, A. and R.B. McCall 1976
 "Exploration and Play - Resume and Redirection"
 American Psychologist, 31, pp 492-508