

WUMPUS PROTOCOL ANALYSIS

by

Barbara Y. White

ABSTRACT

The goal of this research was to assist in the creation of a new, improved Wumpus advisor by taking protocols of ten people learning to play Wumpus with a human coach. It was hoped that by observing these subjects learn Wumpus from a human coach, that insights would be gained into how the computer coach could be modified or extended. In particular, attention was paid to the representations subjects used, the goals they pursued, and the problems they had as well as to the teaching methods used by the human versus the computer coach.

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the Laboratory's artificial intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-75-C-0643.

Working Papers are informal papers intended for internal use.

INTRODUCTION

The ultimate goal of this research is to create a successful and intelligent computer coach for a game of logical and probabalistic reasoning called "Wumpus".

"The theoretical value of constructing intelligent tutoring programs is that such programs provide an experimental medium for systematically investigating alternative cognitive and pedagogical theories. In designing a computerized tutor, it is necessary to represent in explicit form the knowledge and problem solving methods to be taught as well as the pedagogical strategies to be employed. This brings a new level of rigor to educational research. The practical value of such intelligent tutors lies in the personal learning environments they provide, thereby lessening the demand for constant human supervision of students." (Goldstein and Miller, 1976, abstract)


For those unfamiliar with "Wumpus", I include the following quotation from Carr's recent masters thesis (Carr, 1977) which describes the game:

" "Wumpus" is an example of a game which encourages deductive processes and develops a sense of probabilities. The player must seek out and kill a monster, the Wumpus. The player moves in the Wumpus's warren, a network of interconnecting caves containing the Wumpus and other dangers, namely bats and pits. At the start of the game the player is told the number of caves in the warren and the number of bats and pits. In a normal game, there are twenty caves, three pits, and three caves with bats. Before every move the player is told which cave he is in and the caves that he can move to. (Each cave is identified by number.) If any of the neighbouring caves contain a bat he will be informed that he hears squeaking. Likewise, if a neighboring cave contains a pit, the player will be informed of a draft (as pits are bottomless chasms), but in neither case is the player told which of the caves are dangerous. Whenever the

FIGURE 1

AN INTERMEDIATE STATE IN A TYPICAL WUMPUS GAME

(adapted from Goldstein, 1976, figure 2)



UNDERLINED CAVES  
HAVE BEEN VISITED  
BY THE PLAYER

WW = Wumpus Warning  
BW = Bat Warning  
PW = Pit Warning

player is within two caves of the Wumpus he will smell its horrible stench. If the player enters a cave with a pit he loses the game, whereas if he enters a cave containing a bat, he will be carried to a random cave which may contain another bat, a pit, or the Wumpus. The Wumpus eats unwary players who stumble into his lair. The player tries to visit enough caves (avoiding bats and pits) to locate the Wumpus without actually entering its cave. Once the player has found the Wumpus, he can shoot an arrow into the Wumpus' lair from a neighboring cave, killing the beast. If the player shoots an arrow into a cave and the Wumpus is not there, his arrow will ricochet through the warren at random for roughly four caves and may kill either the player or the Wumpus.

Playing the game can involve simple deductions and risk minimization as well as more complex strategies and considerations. At the simplest level, the player can deduce that certain caves are absolutely safe by the absence of warnings. These caves should be explored before any others. At a higher level the player can perceive that certain caves are probably safe, i.e. less likely to contain dangers than other caves. These perceptions are based on the patterns of the warnings and require application of probabilistic heuristics which are commonly used by knowledgeable persons. An advanced player can usually deduce the exact location of the Wumpus through a quite thorough application of logic, but most players develop a general idea of its location without completing all the required deductions (though they are wrong often enough to encourage them to improve their deductive powers). There are also unusual situations which require very advanced considerations of the risks in order to select the best move. The Wumpus is a game that can be enjoyed by the beginner as well as the advanced player." (Carr, 1977, pp.5-7)

Brian Carr has written a tutoring program for this game called "Wusor II". This program builds a model of the player's competence with respect to the reasoning skills required to play Wumpus and adjusts its explanations on the basis of this model. In addition, Carr's (1977) "Wusor II" has elaborate English generation capabilities. However, the current version of the program has limitations. It relies solely on observing the player's moves to build its competence model of the player. It can neither understand nor tutor the representations subjects employ to record information gained whilst playing Wumpus. Therefore, it was the goal of this research to assist in the creation of a new, improved advisor by taking protocols of people learning to play Wumpus with a human coach. It was hoped that by observing subjects in a more


normal setting, that is, without a computer coach, that insights would be gained into how the coach could be modified or extended. In particular, attention was paid to the representations subjects used, the goals they pursued, and the problems they had as well as to the teaching methods used by the human versus the computer coach.

For the purposes of this study, ten protocols were taken in individual sessions. All subjects had the same teacher and were presented with a sequence of twelve standard (i.e. 20 caves with 3 bats, 3 pits and 1 Wumpus) Wumpus games which had been constructed by a random procedure. There was a variation in the number and lengths of sessions for each subject depending on the subject's availability, interest, and energy. Further, subjects were chosen on the basis of their willingness to participate rather than on any attempt at a randomly or systematically selected sample.

### REPRESENTATION DATA

There were three primary data representations that subjects used to record information gained whilst playing Wumpus:-

(1) graphs, e.g.




NOTE: BW = bat warning  
 PW = pit warning  
 WW = Wumpus warning

(2) tables, e.g.

CAVE NO.	NEIGHBORS	WARNINGS
3	16,1,12,5	B
5	14,12,3	W,B
14	5,1,19	W,P

(3) and a list structure used by two players which attempted to look like a map.



The most common data structure was the graph with five out of the ten subjects using the graph alone to represent all the information. The remaining five utilized lists, tables, and graphs, often simultaneously, at some time during their Wumpus career. Subjects' representations evolved as they continued to play not only in terms of the form of the representation but also in terms of what they chose to represent.

The following contains a list of the things people chose to represent:

- cave connections
- warnings
- certain dangers
- the route through the warren
- caves that could contain bats
- caves that could contain pits

- caves that could contain the Wumpus
- where bats are NOT
- where pits are NOT
- where the Wumpus in NOT
- what caves they have heard of
- hypothetical cave structures

- cave connections
- warnings

Most people initially started by recording these two pieces of information since this is what they are given after each move. Some people, however, did not. An eleven year old boy did not initially record the connections between caves until it was discovered that he was failing to make the assumption that saying two caves are neighbours meant the two caves were connected and that you could go from one to the other. Other subjects initially failed to record warnings indicating perhaps that they were failing to define the game adequately. That is, to appreciate that playing Wumpus involves exploring a warren of connected caves gaining information and avoiding dangers until you can deduce where the Wumpus is. The majority of subjects, graduate students, were generally adept at defining the game appropriately and this seemed to help them draw upon the appropriate skills.

- certain dangers

Often people learn that a given cave contains a danger either by landing in the cave or by deduction. Some subjects initially failed to record these certain dangers but this was remedied when they realized it was useful to record this information.

- where they have been

This information is inherent in a list or table structure, however, it is a necessary piece of information and some subjects who used graphs recorded it by some notation such as underlining the cave. Many subjects who used graphs did not initially record where they had been and one had to wait until they discovered it was important before they could be converted to consistently recording it.

- the route through the warren

Again this information is inherent in a list or table structure. It is unnecessary to record this but one subject who used a graph recorded this with arrows. This is not only irrelevant but really makes a mess of the graph when you start back tracking.


- caves that could contain bats
- caves that could contain pits
- caves that could contain the Wumpus

All but two of the subjects who recorded this information did so in the form of tables rather than by putting it directly on their graph. Furthermore, instead of making three separate tables, most subjects put such potential danger locations in one common table which was difficult to work from, e.g.

CAVE NO.	DANGER
7	pit?
1	safe
16	pit?,bat?
3	Wumpus?



I found it impossible to convince subjects who were using a graph for other purposes to put such information on their graph. That is, instead of recording warnings at their point of origin, the dangers could be recorded at their potential location, i.e.



I even tried to convince them by playing a version of Wumpus, which I called Wumpus Junior, where the best data representation is more obviously to record dangers at their potential location. The game involves nine connected squares and one Wumpus whose warning propagated one square, e.g.

START		
	W	

One can move horizontally, vertically, and diagonally. After playing this game for a short time, subjects would concede that one should use warnings, or lack of them, to label the neighbouring squares as "safe" or "Wumpus?". However, they would not generalize this to Wumpus arguing that it would make the graph too messy and hard to read. Some argued that it was therefore better to keep such hypotheses in a separate table and some argued that you didn't need to record such things at all as it was easy enough to compute them from the warnings. Unfortunately, those people who used tables found it hard to coordinate the table with the graph and those people who didn't record this information at all often failed to compute it when it would have helped them.

- where Bats are NOT
- where Pits are NOT
- where the Wumpus is NOT

A few subjects liked to record negative information as well as or instead of positive. In the case of Wumpus information, I think this is reasonable. Wumpus warnings propagate up to two caves away. Thus when a warning is given two caves away from the Wumpus, there are often many unknown caves (i.e. neighbours of unvisited neighbours of the cave where the warning was given) that might contain the Wumpus. Therefore, it is easier to compute where the Wumpus is not and useful to record such information to help provide safe routes and to help in the process of elimination for Wumpus deduction. On the other hand, in the case of bats and pits whose warnings only propagate to one cave away, I think it is easier to compute and more useful to record potential bat and pit locations rather than noting where they are not.

- what caves they have heard of

A few subjects recorded all the caves they had heard of in a list to give themselves a metric on how much of the warren was as yet unknown.

- hypothetical cave structures

This was a device used by two subjects to help them in figuring out possible locations of the Wumpus. They pencilled onto their graph alternative cave structures for unseen caves that would be consistent with data they had thus far. This helped them avoid premature conclusions about Wumpus location. Such a form of hypothesis construction and testing seemed very powerful in assisting with the most difficult aspects of Wumpus deduction.

As was mentioned previously, only five subjects used graphs exclusively to record whatever information they chose to record. Two subjects evolved to

the use of graphs - one went from a table to a listlike graph and finally to a graph, the other went from a listlike graph to a graph. One subject used his graph to record only the connections between caves whilst maintaining warnings in a separate table. Another subject recorded almost everything possible on the graph but in addition kept lists of (1) caves visited, neighbours, and warnings, (2) potential bat locations, (3) potential pit locations, and (4) caves safe from the Wumpus. In addition, a few subjects who used graphs also kept a separate table of caves heard of with notes about potential dangers next to the appropriate cave number in the table. Thus there seems to be quite a variation in what information people decide to represent and, further, whether they represent it on a graph, in a table or list, or both.

#### POTENTIAL COMPUTER CONTROLLED REPRESENTATIONS

There are many reasons why it would be desirable to have the computer coach provide the representation for the player. One argument in favour of this is that the players encounter difficulties related to their representations which such an extended coach could help avoid. For instance, many subjects became preoccupied with the goal of drawing the best possible graph of the connections between caves. This is not only unnecessary but also distracts their attention from learning the game itself. This preoccupation becomes especially noticeable if players assume the warren is a rectangular grid. Other problems with representations included:

- forgetting to record a given piece of information
- recording something incorrectly
- and graphs and tables that were messy and/or too small causing them to be

very hard to read and get information from.


The computer coach could control the representation so that it was tidy, large enough, and the best possible graph of cave connections. This would thereby encourage people to use graphs which are easier to work from than tables. It would also avoid the preoccupation with drawing the best graph since the coach would do this for the player. Furthermore, the graph would be easily legible unlike those of most of the subjects observed. In addition, the coach could check for omissions, inconsistencies, and errors in the player's representation of information thus helping the player. Furthermore, it would ensure that if the player makes a bad move it is due to his reasoning and not just an error in encoding. This would help the coach to diagnose problems more accurately.

An even more compelling set of reasons why the coach should provide the representation facility for the player is that the coach could then be programmed to use information about what the player has represented to build a model of the player. At present the coach must rely solely on inferring the player's knowledge on the basis of what moves he makes. However, a human teacher utilizes very heavily the alterations a player makes to his representation after each move to infer what the player knows and what he needs help with. This ability could be programmed into the coach. For example, if a player utilized a given skill (rule) after receiving information you could predict how he should alter his representation. Then you could utilize whether he alters his representation in the predicted fashion to help infer whether or not he knows that particular skill (rule). Thus letting the computer coach provide the representation facility would allow the coach to build a better model of the player.


Having argued that the computer should provide the representation, the question now arises of what sort of representation facility is appropriate?

How much choice should the player have? For example, should he be able to use either graphs or tables or both? Does it hinder the player if his choice is limited or does it thereby help avoid unnecessary pitfalls?

Unfortunately, for the computer coach to be able to interpret the representation and changes made to it, there will have to be restrictions on what and how the player can represent information. An extreme case would be where the coach creates the representation without the player having any part in it. An example of this is a Wumpus program written by Danny Hillis. In this version of Wumpus, all the information is presented graphically by the computer and the player can move up, down, left, and right through the graph thereby gaining information and avoiding dangers until he thinks he knows where the Wumpus is and decides to shoot. The graph indicates connections between caves (the warren is a grid of 64 points with 25 removed randomly for each new game so that the warren is not a perfect grid) and indicates warnings at their point of origin. The player thus has no part in creating the representation. This version of the game appears to encourage rapid play with little deduction. Possibly this is because the player is not given any facility for recording hypothesis as to the location of dangers. If one were to employ a similar version of representation for the computer coach one might want to add a facility allowing the representaion of such hypothesis either on the graph itself e.g.



or else in the form of external lists, e.g. LIST1: CAVES THAT MIGHT CONTAIN BATS - 2, 13. In addition, one would want to allow the player to indicate on the graph what caves he had visited and "definitely known" dangers. In this example,



the player has visited caves ten, three, and two and found a bat in cave two. The graph itself should be even less gridlike than the style Danny created since in the coach version of Wumpus the warren is not gridlike and representing it as such might mislead players into inaccurate deductions about the location and interconnections of missing caves. The above information, i.e., cave connections, warnings, caves visited, and potential danger locations, could all be recorded in a single graph, however, (1) most players did not do this but rather used lists or tables as well and (2) list structures such as (i) POTENTIAL BAT LOCATIONS and (ii) POTENTIAL PIT LOCATIONS are useful in eliminating possibilities. For example, from the following list in a game with two bats \* POTENTIAL BAT LOCATIONS: - 7,14; - 14,3,1; - 2,9,15 \* you can infer that cave 14 must contain a bat and that caves 7,3, and 1 cannot. Thus some list structures help facilitate certain deductions. Furthermore, in an ideal environment it would be helpful to the player and the coach to add the capability of allowing the player to draw in hypothetical cave structures for missing caves to help in Wumpus deduction. Thus, I think, the player should be presented with a graph indicating connections between caves and marking warnings and caves visited with the added facility of the player being able to record certain dangers on the graph and potential dangers either on the graph itself or in separate lists,

e.g. L1: POTENTIAL BAT LOCATIONS , or else both.

L2: POTENTIAL PIT LOCATIONS

L3: WHERE THE WUMPUS CANNOT BE

L4: WHERE THE WUMPUS MIGHT BE

GOALS, PROBLEMS, and TEACHING METHODS

From the observation of ten subjects playing Wumpus, I have tried to devise a classification scheme for the top level goals people pursue whilst playing the game.

The first category of player is the "Topologist". This player decides that the most important component of Wumpus play concerns drawing the best possible map of the connections between caves. After each move he redraws his map of the warren so as to optimally incorporate the new connection data. This was an unhealthy preoccupation shared by several subjects. All were readily persuaded to give it up after having labouriously redrawn their map three or four times.

The second category of player is the "Gambler". This person views Wumpus as a game similar to roulette where chance rather than deductive thinking rules your fate. The extreme gambler just wanders aimlessly through the warren getting picked up by bats until he eventually falls into a pit or gets eaten by the Wumpus. The Masochistic Gambler and the Conservative Gambler, on the other hand, do employ a little deduction using positive evidence about bats, pits, and the Wumpus to guide their path. The Masochist is trying to get killed and seeks out dangers whilst the Conservative is trying not to get killed and hence avoids dangers. However, neither of them engage in the subtleties of Wumpus deduction nor do they actively seek out information about the Wumpus.

In contrast to the Gambler, the next category of player, the Hunter,

engages in much deduction and actively pursues the Wumpus. There is the Inferencing Expert, who resembles closely the Expert of the computer coach, employing standard rules of deduction and route direction with the exception that he does not use advanced probability calculations. Then there is the Ruthless Hunter who is similar to the Inferencing Expert, however, the Ruthless Hunter seeks out the Wumpus no matter what. That is, given a choice between a safe move which may or may not give information about the Wumpus versus a dangerous move which would, nonetheless, give Wumpus information, he will pick the dangerous move in order to optimize information gain. Finally, there is the Intellectual Sophisticate who employs the rules of deduction and route direction of the Inferencing Expert but in addition engages in calculating advanced probabilities. This is the most successful level of play and was eventually obtained by five out of ten subjects.

Players do not necessarily use the same strategy throughout their Wumpus career. Often their strategy changes as their game evolves. For example, one subject went from being a Topologist to a Conservative Gambler and finally to an Inferencing Expert. Some players even offered to change their strategy because they thought that their strategy was not how I wanted them to play. For instance, one Ruthless Hunter whilst foaming at the mouth said, "I suppose you would prefer me to play it safe," and a Gambler was heard to say, "well I guess you really want me to stop and think". The computer coach assumes that the player will adopt the definition of the game corresponding to the Intellectual Sophisticate. This is not the case - some subjects defined it as a topological problem, some as a game of chance, others as a problem of deductive reasoning and still others saw probability knowledge as being relevant. Since subjects are usually insecure about their strategy, the computer coach could most probably successfully impose its definition onto the player. However, the Coach may be depriving the person of an important aspect



of learning the game, that is, defining the problem and creating suitable goals. This alludes to a class of problem solving expertise that the current version of the computer coach lacks and hence cannot teach. This includes such things as:

- defining the problem
- creating goals and subgoals
- developing an adequate representation
- playing scientist by building and testing theories of the game both by inductive and deductive methods
- consciously thinking aloud about the problem
- graphically working through possible danger locations
  - especially the Wumpus's whereabouts
- being systematic and not omitting important deductions
- making decisions from incomplete knowledge

Unfortunately, the teaching of such skills might be severely handicapped at this point by the coach's inability to comprehend English.

One could argue that by imposing its strategy onto the player, the coach is depriving some people, such as Gamblers and Ruthless Hunters, of a version of the game that is more natural and enjoyable to them. After all the subject is not just a tabular rasa waiting to be filled. He already has theories which he is applying to the game and these need to be modelled and utilized in the learning situation. If the coach were to be modified so that it could adopt different strategies, it is a little difficult to see what it could then teach. This is especially true of the Gambler since he is playing for the fun of challenging chance and hence tutoring deductive reasoning would be inappropriate. Possibly some probability information could be introduced in the context of working out the odds of survival and thereby sneak in the deductive rules under the guise of beating the odds. In the case of the

Ruthless Hunter, the deductive reasoning would still be appropriate, however, most of the probability calculations would not since the player wishes to optimize information gain and is not concerned with survival. For the coach to tutor the Topologist, the coach would have to be given additional expertise either in the form of:

- creating the map for the player hence avoiding the problem,
- knowing about such maps so that it tutors a form of graph theory,
- or else implementing some of the problem solving expertise outlined previously so that it could change the players' goal to a more appropriate one.

This again raises the question of what could and should the computer coach teach. Perhaps this is best approached by looking at the problems or bugs subjects exhibited whilst playing Wumpus.

The following list is an enumeration of bugs that subjects manifested whilst playing Wumpus:-

#### Reasoning Bugs

- (1) not realizing that more than one bat or pit is possibly implied from a single bat or pit warning.
- (2) believing that multiple evidence implies a definite danger.
- (3) not believing that multiple evidence makes a danger more likely.
- (4) not being able to infer that a cave is two away from the Wumpus.
- (5) believing that if a cave is two away from the Wumpus then all of its' neighbours must be one away.
- (6) not making use of one away versus two away in Wumpus location deductions.
- (7) not using negative evidence to eliminate possible danger locations.

### Route Directing Bugs

- (1) not realizing that you can retrace your steps.
- (2) not actively seeking out information.
- (3) not playing it safe.

### Representation Bugs

- (1) not using graphs.
- (2) not recording dangers or warnings at all.
- (3) making errors in recording information:-
  - repeating cave numbers
  - recording warnings or connections incorrectly.
- (4) getting preoccupied with drawing the best possible graph of cave connections.
- (5) not recording certain dangers.
- (6) not recording where you have been.
- (7) using inconsistent notation.

### Problem Solving Bugs

- (1) not using all the appropriate rules even though they are known and used in other situations.
- (2) not seeing probability knowledge as being relevant to the game.
- (3) not working out possible structures for unseen caves near Wumpus that are consistent with the data.

The reasoning and route directing bugs are all tutored by Carr's Wusor II with the exception of being able to explain in depth that multiple evidence makes a danger more likely. This points to a problem the coach has of not being able to further expand explanations of its' rules. That is, it assumes

that you will believe statements like "pits are more dangerous than bats" or "it is sometimes useful to seek out bats" and hence does not possess the capability of further convincing the player by a more extensive explanation of the "laws of Wumpus".

The representation bugs could all be handled by getting the coach to provide and monitor the representation facility as previously suggested in the section on representation in this paper.

The problem solving bugs are not tutored by the current version of the coach. (N.B. the list of problem solving bugs is probably very incomplete.) They could perhaps be dealt with by implementing a version of the problem solving expert suggested previously. However, I think this alone would be insufficient since these bugs also closely correspond to my list of "things people found most difficult whilst playing Wumpus" (i.e. (1) the final stages of deducing the Wumpus's location and (2) utilizing the appropriate probability knowledge in the appropriate situation). Therefore, I think they require the addition of new deductive and probabilistic expertise and not just general problem solving knowledge. For example, the final stages of deducing where the Wumpus is hiding gave all subjects problems. Determining accurately when they knew for certain where the Wumpus was or knowing what additional information they needed by hypothesizing possible structures for unseen caves was a set of skills that only a few subjects developed. The current coach could be improved by being able to discuss Wumpus deduction in more detail and also by adding the facility for putting hypothetical cave structures on the graph. In addition, many subjects did not see that probability calculations were relevant to the game. Further, those subjects who did calculate probabilities often had trouble determining exactly what probability it was appropriate to calculate. The coach could be extended to include more expertise on which probability calculations are appropriate under what

circumstances and perhaps incorporate this with a general problem solving expert so that the Wumpus coach could be a vehicle for teaching probability problem solving.

My determination of the skills to teach Wumpus players and the order in which to teach those skills correspond rather closely to those of the computer coach except that I also tutored the representation of information as well as the reasoning. For example, my outline of what to teach is as follows:-

- Representation - represent connections
  - record warnings
- Route Direction - pursue safe caves
- Positive Evidence - bats and pits
- Negative Evidence - bats and pits
- Multiple Evidence
  - Route Direction - pursue Wumpus information
  - Positive Evidence - Wumpus
  - Negative Evidence - Wumpus, 2 away vs. 1 away
- Representation - record definite dangers
  - record warnings at potential danger location rather than at source of warning
- Route Direction - it is sometimes useful to seek out bats
- Cave Sets
- Probability Calculations

However, my manner of teaching these skills was somewhat different. The computer coach teaches through the use of a hierarchically ordered set of skills. At a given time, the program is trying to teach a given subset of the skills appropriate to the player's knowledge level. It does so by

interrupting the player whenever he makes a bad move with respect to those skills or skills he is already supposed to know. The interruption consists of a deductive proof (utilizing those skills and ones the player already knows) that the move is not optimal and that there is a better move. So, the player is presented with a chain of deductive reasoning consisting of the application of rules he is learning to the direction of his next move. There are several alternative methods for teaching these skills. The program tends to appear condescending with its "listen, stupid, there is a better move here" technique, also, if the chain of inference gets too deep many people will not pay attention to it. Rather than give a proof of a better move, I often just asked the question, "Can you think of a better move?" or else I just stated a relevant rule such as, "Remember, it is sometimes useful to take a bat" etc.. Alternatively, one could explain in the context of an example which was similar to the current situation or else attempt to teach the same skills in a similar game such as Wumpus Junior, Clue, Master Mind etc.. The latter would also encourage encoding the knowledge in a form that has broader uses. In fact, one criticism of the current coach is that it is teaching skills that only apply to Wumpus. Thus enlarging the context might be useful.

The computer coach could also be modified to teach by creating a synthetic student who didn't know certain rules and getting the student to teach it or else by getting the student to create a synthetic player. Also, the coach could (and already does to some extent) structure its games and modify the game being played so as to produce an optimal learning situation for a given rule. (One problem with the game I was using with 20 caves and 3 bats and 3 pits was that often the Wumpus was surrounded by dangers and you couldn't get next to him to shoot. Also, the odds were too high that the player got killed by chance rather than by faulty reasoning.)

Another difference between the teaching of the computer coach and myself

was in terms of when we intervened with an explanation. The coach only intervened when the player made a bad move with respect to the rules it was trying to teach, whereas, I intervened not only after bad moves but also when the player could have deduced something but didn't. For example, I would say, "Can you figure out where the pit has to be?". Intervening at points other than just after bad moves might be useful, however, I suffered from the bug of intervening too frequently. Several players told me to "shutup". Even the program, which intervened less often than I did, intervenes too often. The coach could be changed so that it only intervenes when the player hasn't demonstrated increasing knowledge of any rule for a specified period of time. This might prove more satisfactory since it would allow subjects to "figure it out for themselves" but step in if they need help before they get too frustrated.

To summarize differences between the coach and myself, I possessed more areas of expertise and thus could help mend bugs and deal with different strategies better. I also could draw upon more clues - facial expressions, gestures, change in representation, English utterances etc. - from which to model what the player knew. However, the coach (1) was more patient - intervening less frequently, (2) had a far superior memory of what the player knew, and (3) stuck to the hierarchy of rules, whereas, I tried to teach advanced rules too soon.

#### CONCLUSIONS

The preceding report has suggested the following modifications and extensions to the Wumpus Advisor:

### Representational Facility

The report has proposed the addition of a representational facility where the player is presented with a graph indicating connections between caves and marking warnings and caves visited with the added facility of the player being able to record certain dangers on the graph and potential dangers either on the graph itself, in separate lists or else both. This would encourage the use of graphs which are easier to work from than tables and, further, would help avoid errors in recording and provide a clear representation of the information. The coach could also have the facility of utilizing the way in which the player updates this representation to infer what the player knows.

### Explanatory Expertise

This report has proposed the following additions to the coach's explanatory capability:

- add a module of problem solving expertise so that the coach could tutor such things as defining the problem and being systematic.
- make it able to diagnose different strategies and adjust what it teaches accordingly.
- make it able to expand what it is trying to teach to a deeper level, for example, be able to explain why multiple evidence makes a danger more likely.
- increase its tutoring capacity for the final stages of Wumpus deduction.
- add more probability related expertise.

### Teaching Methods

It was suggested that the coach should use additional teaching methods



such as:

- use examples
- ask vague questions, e.g., "Can you think of a better move?"
- teach the same skill in a similar game such as Master Mind or Clue.
- remind the student of a skill he already knows in an appropriate situation.
- get the player to teach a synthetic student.
- get the student to create a synthetic player.

It was also suggested that the computer coach change when and how often it intervenes. In the case of "when", it should intervene not only when a student makes a bad move but also when he could have deduced something and didn't, e.g., "Can you figure out where the pit has to be?". In the case of "how often", the coach should only intervene when the player has not improved on any skill after a specified period of time.

REFERENCES

Carr, Brian, The Wumpus Advisor (Master's Thesis), Massachusetts Institute of Technology, Department of Electrical Engineering and Computer Science, May 1977.

Carr, Brian and Goldstein, Ira P., Overlays: A Teory of Modelling for Computer Aided Instruction, Massachusetts Institute of Technology, Artificial Intelligence Laboratory, Memo 406, February 1977.

Goldstein, Ira P., A Preliminary Proposal for Research on The Computer as Coach: An Athletic Paradigm for Intellectual Education, Massachusetts Institute of Technology, Artiificial Intelligence Laboratory, Memo 389, December 1976.

Goldstein, Ira P., and Miller, Mark L., Intelligent Tutoring Programs: A Proposal For Research, Massachusetts Institute of Technology, Artificial Intelligence Laboratory, Working Paper 122, March 1976.

Stansfield, James L., Carr, Brian P., and Goldstein, Ira P., Wumpus Advisor I. A first implementation of a program that tutors logical and probablistic reasoning skills., Massachusetts Institute of Technology, Artificial Intelligence Laboratory, Memo 381, October 1976.

Yob, G., Hunt the Wumpus, Creative Computing, Sept.-Oct. 1975.

APPENDIX

SUMMARY OF SUBJECT DATA

SUBJECT	SEX	AGE	OCCUPATION	REPRESENTATIONS	GOALS PURSUED	BUGS NOTED*
A	M	11	student	graph --> graph & list	Topologist --> Conservative Gambler --> Inferencing Expert	Rep. Bugs 2,4,5,&6 Reas. Bugs 1&2 R.D. Bug 2 P.S. Bugs 1,2,&3
B	F	mid 20's	graduate student	graph	Inferencing Expert	Rep. Bug 3 Reas. Bug 3
C	M	late 20's	research associate	graph	Gambler --> Ruthless Hunter-->Intellectual Sophisticate	Rep. Bug 7 Reas. Bug 3
D	M	late 20's	research associate	table & graph	Inferencing Expert	Rep. Bug 4 Reas. Bugs 6 & 7
E	M	late 20's	administra- tor	listlike graph & list --> graph & list	Conservative Gambler --> Inferencing Expert	Rep. Bugs 2 & 3 Reas. Bug 7
F	F	mid 20's	secretary	table--> list- like graph--> graph	Gambler	Rep. Bug 1 Reas. Bugs 4 & 6
G	M	mid 20's	graduate student	graph	Intellectual Sophisticate	progressed rapidly -no bugs recorded
H	M	mid 20's	graduate student	graph	Intellectual Sophisticate	progressed rapidly -no bugs recorded
I	M	mid 20's	graduate student	graph	Intellectual Sophisticate	progressed rapidly -no bugs recorded
J	F	mid 20's	graduate student	table--> graph --> graph & lists	Topologist -->Cons. Gambler -->Inferencing Expt.-->Int. Sophist.	Rep. Bugs 1,4, & 6 Reas. Bugs 1 & 5 P.S. Bugs 1 & 3

\* lack of bugs being noted does not necessarily imply that the subject does not possess them, only that the experimenter did not record them.

NOTE: None of the subjects had played Wumpus prior to these sessions.