HUMAN FACTORS, 1986, 28(6), 703–716

Modeling Fault Diagnosis as the Activation and Use of a Frame System

PHILIP J. SMITH,[1] WALTER C. GIFFIN, THOMAS H. ROCKWELL, *and* MARK THOMAS, *Ohio State University, Columbus, Ohio*

Twenty pilots with instrument flight ratings were asked to perform a fault-diagnosis task for which they had relevant domain knowledge. The pilots were asked to think out loud as they requested and interpreted information. Performances were then modeled as the activation and use of a frame system. Cognitive biases, memory distortions and losses, and failures to correctly diagnose the problem were studied in the context of this frame system model.

INTRODUCTION

This study addresses the question of how domain-specific knowledge is used in fault diagnosis (Gentner and Stevens, 1983; Rasmussen and Rouse, 1981). The fault studied is the failure of the vacuum system in an airplane. The domain-specific knowledge under consideration is the knowledge possessed by instrument-rated pilots.

It should be noted that—although the subjects were pilots and the problem-solving task required knowledge of aviation systems—this is not a study of how pilots diagnose faults while actually in flight. Rather it is an attempt to (1) explore how people use relevant domain-specific knowledge to solve a problem, and (2) represent this problem-solving performance as the activation and use of knowledge stored in a frame system (Aikens, 1983; Minsky, 1975). This line of research offers the potential to refine our theories about the nature and causes of cognitive biases in the performance of diagnostic tasks. Three pertinent questions arise:

(1) Does domain-specific knowledge, organized as frames, help people to overcome certain *cognitive biases?*
(2) Can a more complete description be developed that provides insight into the mechanisms causing confirmation bias and like phenomena?
(3) Are some problem-solving strategies more likely than others to protect a person from the deleterious effects of cognitive biases and limitations?

It is hypothesized, then, that the pilots' relevant domain-specific knowledge can be represented as a frame system. The basic knowledge structure within such a system is a frame. Minsky (1975) defines a frame as:

> . . . a data structure for representing a stereotyped situation like being in a living room or going to a child's birthday party. Attached to each frame are several kinds of information. Some of this information is about how to use the frame. Some is about what one can expect to happen next. Some is about what to do if these expectations are not confirmed (p. 212).

In an aviation setting, one such frame or "stereotyped situation" might be a plane in a descent. A model is presented below in which such knowledge representations play a cen-

[1] Requests for reprints should be sent to Philip J. Smith, Department of Industrial and Systems Engineering, Ohio State University, 1971 Neil Avenue, Columbus, OH 43210.

tral role in directing fault-diagnosis performance.

METHOD

Pilots were read a scenario that provided certain instrument indications and background information pertaining to a flight over the New England area. They were told that a problem existed at a point in time described by the scenario and were asked to try to diagnose the cause of the problem. In order to perform this task, they were allowed to request any information that would normally be available to a pilot under the conditions specified by the scenario. Requested information was provided orally by an experimenter.

Each subject was tested in a separate session. The entire session was tape-recorded.

Subjects

Twenty-six pilots with instrument ratings served as subjects. Pilots were paid $10 for a single session that lasted from one to two hours.

Procedure

The two primary tasks involved fault diagnosis (Task 1) and a recall test (Task 2). In addition, three other supplementary tasks were carried out. Tasks were performed in numerical order by all subjects. Since Tasks 1 and 2 were always performed first, there is a potential for confounding of results on the remaining tasks. Consequently, the information provided by the latter tasks will be presented only insofar as it supports or contradicts conclusions drawn from the two primary tasks. Similar caution must be applied when interpreting results from Task 2.

Task 1. Each subject was asked to describe "what a pilot should do in order to determine the cause of a problem that has developed while flying a Cherokee Arrow that has a 200-horsepower, fuel-injected Lycoming engine. This particular plane is not turbocharged and does not have an autopilot."

The subject was asked to think out loud while trying to diagnose the problem. (Before starting, subjects were given training in thinking out loud on a scenario involving selection of a restaurant.)

The scenario (presented below) was one in which a plane's vacuum pump failed. This fact was indicated by a zero reading on the suction gauge. The vacuum pump drives the artificial horizon and directional gyro. As the artificial horizon lost its drive, it started to sag to the right and the pilot compensated (unconsciously) by turning left. This leveled the artificial horizon and put the plane in a descending left bank. The resulting nose-down attitude caused an increase in airspeed and a descent.

At the point in time represented in the scenario, the plane had faulty readings on the artificial horizon and directional gyro. The plane was descending nose-down and was in a left bank while these instruments indicated straight and level flight.

The scenario was as follows:

> Imagine that this pilot is making a day trip from Augusta, Maine, to Lebanon, New Hampshire. He flies out of Augusta at 9 A.M. cleared Victor 39 to Neets intersection, Victor 496 to Lebanon. He climbs to a cruising altitude of 6000 feet. After 15 minutes of routine flying in instrument conditions in the clouds, the instruments indicate an increase in airspeed, a steadily decreasing altitude, and zero pitch. So, the instruments indicate an increase in airspeed, a steadily decreasing altitude, and zero pitch. How should this pilot go about identifying his problem?

After hearing the scenario, pilots began requesting information in an effort to diagnose the fault. They continued until they arrived at a conclusion or decided that it was impossible to arrive at a conclusion with the available information.

Task 2. Pilots were asked to recall everything they remembered about this flight. They were told to be very specific about any instrument readings or conditions they remembered.

Task 3. Pilots were given a knowledge test. They were asked to describe the information they would collect in order to determine whether a plane had one of a number of different problems.

Task 4. Pilots were asked to think back to the original problem-solving task. They were asked to describe what had been their impression of the plane's physical orientation while they were trying to diagnose the problem.

Task 5. Pilots were asked whether they formed a visual image of an instrument panel while performing Task 1. If the answer was yes, they were asked what they visualized.

RESULTS AND DISCUSSION

Before presenting summary statistics for the results of the experiment, two full verbal protocols will be presented:

Subject #3

Query 1. "Steadily decreasing altitude. Then I would also assume that that also includes then a showing a descent on the vertical velocity indicator?" [Why?] "Is the vertical speed indicator having a reading consistent with the altimeter. To try to narrow down is it a pitot-static system problem.
"That indicates to me that the vertical speed indicator is consistent with the altimeter."

Query 2. "At this point, then, I would then change my attention away from the, no, I take that back. The airspeed indicator is indicating an increase in airspeed. Is that correct?
"At this point, I will rule out the pitot-static system. Those instruments all seem to be consistent."

Query 3. "With an increase in airspeed, then, the next question is, is the manifold pressure, what is the trend of the manifold pressure gauge?" [Why?] "To try to narrow down is it an engine problem of some sort, am I losing engine power?"

Query 4. "Also with regard the engine just to get information as to whether the engine and prop in this case is working correctly, what is the RPM reading?
"At this point it seems that the pitot-static system is correct. The engine seems to be functioning correctly. The engine seems to be running, producing power."

Query 5. "My next line of thought would be some sort of control problem. I was going to ask a question about the trim, but I'm assuming the trim hasn't been played with. I just, a new thought came to mind and the new thought is that if I am decreasing altitude and zero pitch change, in other words, I haven't evidently put in any control input to affect the elevator. Well, let me phrase it as a question. Is there ice, am I receiving ice on the wings of any sort?
"That takes care of that problem."

Query 6. "Then let's go back to the controls. Is the pitch trim operating correctly, the trim wheel? Has the trim wheel changed position?"

Query 7. "At this point I'm becoming stumped. Let me ask another question which maybe clarifies the initial conditions. That is, I have zero pitch, meaning that indicates that I haven't had a forward deflection in the control wheel. I haven't added down elevator. I'm losing altitude, gaining airspeed, but have not had a, is the nose pitched over is what I'm trying to determine at this point. I'm in the clouds. The only way to determine that is either through the altitude change, which obviously is down, but the next thing to check would be the attitude indicator and I'm assuming that the attitude indicator is indicating level because the initial condition saying there was no pitch change. Ah! I have just rung a bell! Next question: Is the vacuum, what is the reading on the vacuum gauge?"

Conclusion. "My problem is with the vacuum system and I'm losing pressure to my gyroscopic instruments."

Subject #1

Query 1. "The first thing I would think about is with the decreasing altitude and increasing airspeed, that for some reason the plane is starting to go down and I would look to confirm that right away with the attitude indicator. There is zero pitch in there.

> It should show down pitch. So the first instrument I would look at since it runs off of suction, would be over at the suction, to see if it's producing any vacuum. What does the suction gauge show?"

Conclusion. "You have a nose-down attitude and the vacuum pump's gone."

These two protocols illustrate the apparent heterogeneity of the subjects' performances. Subject #1 asked one question while Subject #3 asked eight, yet both arrived at the same conclusion.

Although data were collected for 26 subjects, the following analysis will be based on only 20 of them. Since the objective of this study was to model the way pilots use their knowledge structures (as opposed to whether they have the necessary knowledge), any subjects demonstrating knowledge errors (Task 3) that would prevent them from solving the problem were deleted from the data set. Four subjects were deleted for this reason.

A fifth subject was deleted for failing to follow instructions, and a sixth subject was eliminated because he misinterpreted the meaning of the scenario.

Summary Statistics

Of the 20 final subjects, 11 concluded that there was a vacuum system failure (Group A), 4 stated that the problem was a malfunctioning artificial horizon (Group B), 1 decided the problem was a downdraft (Group C), and 3 concluded that the problem could not be diagnosed with the available information (Group D). A final pilot (Group E) detected the faulty artificial horizon but then concluded that he could not diagnose the problem (for reasons that will be explained later). These pilots ranged in age from 21 to 59 years (with a mean of 33) and in flight experience from 200 to 20 000 hours (with a mean of 1900 hours).

In Task 4, pilots were asked what their conclusion had been regarding the plane's physical orientation (during the Scenario Task). Fourteen reported that, at the end of the Scenario Task, they thought the plane was in a straight nose-down descent. Five (all in Group A) thought the plane was descending in a left-bank with the nose down. The pilot in Group C thought the plane was in a straight and level descent, with the nose on the horizon.

In the recall task, the question of interest is whether or not the pilots remembered the three instrument indications given in the scenario. Assuming that the probability of retrieving specific information from memory is related to the amount of attention that information was given during the problem-solving task, the subjects' responses to the recall task provide evidence of the salience of these three instrument readings. Of the 20 pilots, all but one in Group A recalled the indications of increasing airspeed and decreasing altitude. All of the pilots in Groups A, B, and C recalled the indication of zero pitch, whereas none of the pilots in Groups D and E recalled this indication.

Initially Activated Frames

Table 1 shows the initial queries for the 20 pilots studied. Based on an analysis of the associated verbal protocols, labels for the following initially activated frames were hypothesized: the plane is in a *descent;* there is a *power loss;* there is *icing;* there is a *pitot-static system malfunction;* there is *static port icing;* there is a *blocked static port;* there is *pitot tube icing;* there is an *airspeed indicator malfunction;* there is a *vacuum system malfunction;* my *memory* may be in *error*. Two subjects seem to have a frame dealing with beliefs about their own limitations and abilities (Norman, 1983). This last frame, then, is concerned with the possibility that the pilot has not recalled the scenario information correctly.

Thus, the evidence is consistent with a model in which a variety of frames exist in

TABLE 1

Initial Queries

	Number of Subjects Asking
What is the reading on the:	
vertical speed indicator?	4
airspeed indicator?	2
manifold pressure gauge?	3
tachometer?	2
outside air temperature gauge?	2
suction gauge?	2
What happens if:	
the alternate static source is opened?	2
Is:	
the pitot heat on?	1
there visible moisture in the air?	1
there an increase in wind noise outside the plane?	1

the pilots' memories. Furthermore, the results suggest that the same "stimulus" (reading of the scenario) can lead to the activation of different frames in different pilots' memories.

Initial Activation of Frames

The preceding analysis identified 10 frames that were used by the 20 pilots to generate their initial queries. (Not all pilots activated the same frame.) This subsection addresses the next question: How were these frames activated? The goal is to understand better the mental processes that occurred between the time the experimenter began reading the scenario and the time at which the pilot made the first query.

By its very nature, protocol analysis provides only fleeting glimpses into the mental processes occurring within any one subject. Subjects do not report all of their thoughts. Furthermore, even if two subjects activate the same set of mental processes, their comments may provide evidence relevant to different portions of these processes. Thus, in order to construct a model that is even somewhat complete, it is desirable to make the following assumption: Unless evidence to the contrary exists, one can assume that if two pilots ask the same question (e.g., What is the reading on the manifold pressure gauge?), that question was produced by the same mental processes (at least in terms of important characteristics). This assumption is based on the goal of developing a parsimonious explanation of performance (a desire to introduce individual differences only when necessary).

The scenario that was read can be thought of as a set of cues or clues indicating what the problem was. The first questions to be addressed are: (1) What are the cues that subjects are attending to? and (2) What frames are being activated by these cues? (Pauker, Gorry, Kassirer, and Schwartz, 1976.) Evidence that a cue has been given attention is the fact that the pilot repeats that cue out loud. There may, of course, be other cues that have received attention but have not been repeated by the pilot.

This type of analysis was applied to the data for all 20 subjects. The data used for the analysis were the spontaneous comments of a pilot before the first query, the first query itself, and the pilot's comments immediately after the first query (spontaneous or in response to the prod: "Why are you interested in that information?"). Thus, the data used consisted of all statements made after the reading of the scenario, but before the asking of a second query by the pilot.

In addition, subjects' statements regarding their initial perceptions of the plane's orientation (Task 4) were used to distinguish between subjects who thought the plane was in a nose-down as opposed to a nose-level descent.

Figure 1 shows the results of this analysis for 18 of the pilots (those pilots who attended to only a subset of the available cues). The data suggest that eight of these pilots initially activated a frame hypothesizing that the plane had a pitot-static system malfunction. Four pilots focused attention on the possibility of a nose-level descent. Five other subjects activated a frame representing the plane as being in a nose-down descent. Finally, one pilot asked, "Is there an increase in noise outside the plane?" [Why?] "That would indicate an increase in airspeed, would back up that instrument indication." He apparently activated a frame for a possible airspeed indicator malfunction.

The two remaining pilots (not shown in Figure 1) were able to attend to and properly utilize all three cues (increase on the airspeed indicator, decrease on the altimeter, and zero pitch on the artificial horizon). They immediately concluded correctly that the plane was in a nose-down descent and that there was probably something wrong with the vacuum system.

Slots. The analysis depicted in Figure 1 provides a number of insights into the mental processes and knowledge structures involved in performance on this problem-solving task. First, it indicates a basic structure in which each frame, or stereotypic situation, has two slots: (1) possible causes of that state of nature (e.g., a power loss can cause a nose-down descent); and (2) expectations (expected instrument readings and other observable conditions such as ice on the wings) given the fact that the stereotypic situation exists.

Slot-fillers. Each of the two slots (*causes* and *expectations*) for a frame has a set of slot-fillers associated with it. As an illustration, Figure 2 shows a frame representing a nose-down descent. It has eight slot-fillers associated with the *causes* slot (structural icing, power loss, etc.) and five slot-fillers associated with the *expectations* slot. These slot-fillers can be used to direct a pilot's data-collection activities. For instance, in order to determine if the plane is actually in a nose-down descent, the pilot could check to see if the expectations are met by looking at the artificial horizon, at the altimeter, etc. (see Figure 2). To look for the presence of a particular cause of a nose-down descent, the pilot could activate the frames associated with each of the slot-fillers listed in the *causes* slot (see Figure 2).

The psychology literature suggests that people tend to collect data that will help to confirm a hypothesis that they wish to test, rather than trying to seek data that might help reject or falsify that hypothesis (Mynatt, Doherty, and Tweney, 1977). The present study, however, shows many instances in which an information request (generated by a slot-filler in an activated frame) provided data that allowed the pilot immediately to reject the hypothesis that was being considered (see Figure 1). Inquiries about the reading on the outside air temperature gauge, for instance, caused pilots immediately to reject the possibility of icing (it was too warm outside). Similarly, testing the effect of opening the alternate static source led to prompt elimination of the possibility of a blocked static-port.

Mynatt et al. (1977) suggest that it is "very difficult to elicit behaviour by which subjects can prove the falsity of a hypothesis that they are entertaining" (p. 95). The findings of our study suggest that the knowledge structures (i.e., the slot-fillers in the *expectations* slot) containing domain-specific knowledge may be very effective in helping people to overcome this difficulty.

Instructions for use. An analysis of the data indicates that the instructions for use in a frame are based on one of the following two lines of reasoning.

(1) If a frame is a valid representation of the state of nature (e.g., nose-down descent), then the expected readings on certain instruments (listed in the *expectations* slot of that frame) should be present. To assess that frame's validity, the pilot should ask for the readings on those instruments.

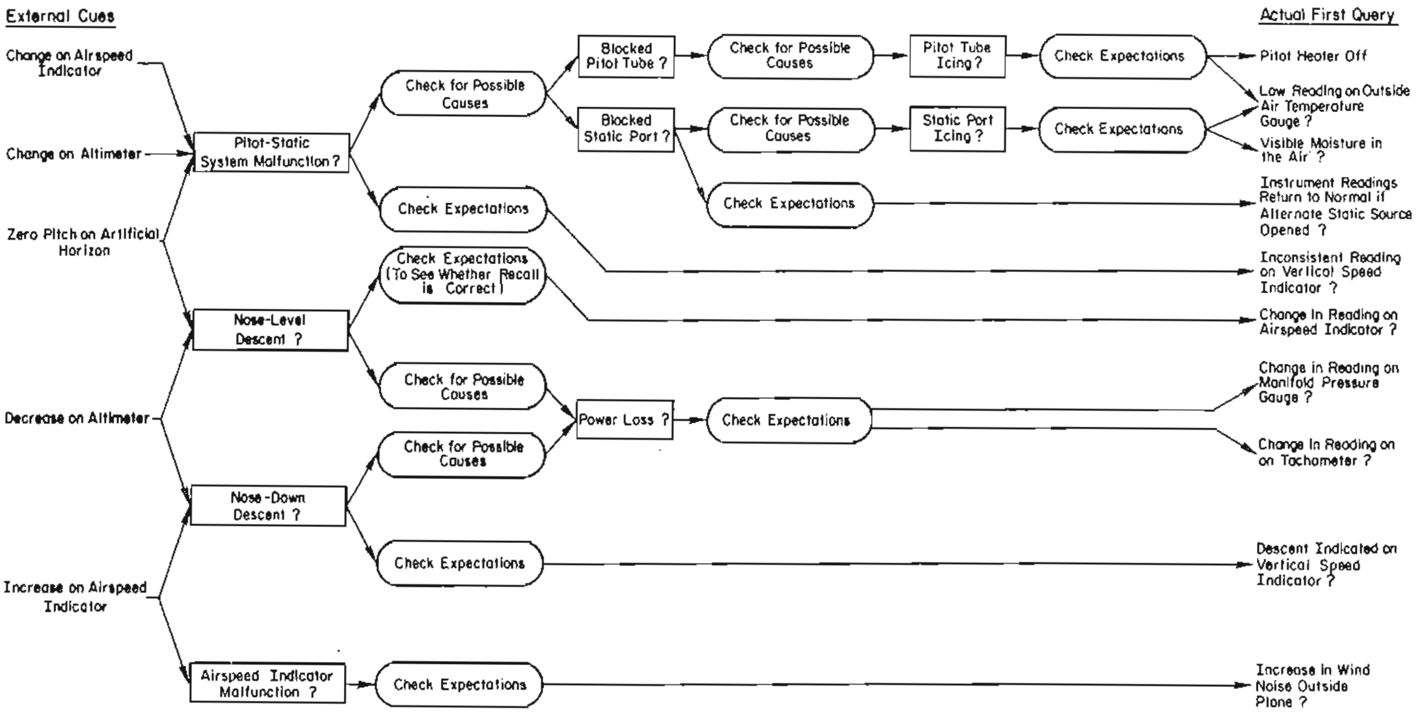

Figure 1. *Hypothesized processing sequence leading to first queries.*

Frame Label : NOSE-DOWN DESCENT

Causes : Structural Icing
Power Loss
Downdraft
Gear Down
Trim Wheel Mispositioned
Yoke Mispositioned (Decreased Backpressure)
Flaps Down
Banked Plane

Expectations : Nose-Down on Artificial Horizon
Descent Shown on Altimeter
Descent Shown on Vertical Speed Indicator
Increase Shown on Airspeed Indicator
If Backpressure Applied to Yoke,
Instrument Indication of Descent
will Cease or be Reduced in Magnitude

Instructions for Use : Check to See if One of the Causes is Present

Figure 2. *Possible frame representing a nose-down descent.*

(2) If a frame is a valid representation of the state of nature, then something must have caused that state to occur. The pilot should assume that the frame is a valid representation and look for a possible cause (from the list in the *causes* slot).

Use of the first line of reasoning is an example of a strategy of top-down refinement (Hasling, Clancey, and Rennels, 1984), a strategy implicitly used in many medical consultation systems that have been developed using knowledge-based systems techniques (Chandrasekaran, 1983). Note that one implication of the present analysis is that not all pilots used a top-down refinement strategy.

Selective attention. Figure 1 also illustrates how selective attention to different subsets of the available cues can lead to the activation of different frames. Pilots activating the nose-down descent frame ignored the information about zero pitch. Those hypothesizing a pitot-static system malfunction failed to note the directions of the changes on the airspeed indicator and altimeter.

What, then, were the determinants of attention for these subjects? What, for instance, made the airspeed and altitude information more salient to the pilots who activated the nose-down descent frame? Two possible causes can be hypothesized.

First, studies of human perception and attention suggest that "the perceptual system actively attempts to reconstruct the external environment in an effort to cope with the massive volume of information it continually encounters. The 'match-mismatch' notion clearly identifies the unexpected as a, if not the, crucial determinant of attention" (Dember and Warm, 1979, p. 131). Extending the same concept to the "perception" or comprehension of text (the scenario), it is predicted that, in this problem, the indications regarding airspeed and altitude should be more salient than that of pitch. Prior to hearing about these instrument indications, the subject was told that the plane has been cruising for 15 minutes at a constant altitude. Thus, a model or reconstruction of the situation would indicate a constant airspeed, no change in altitude, and zero pitch. This means that two of the cues—an increase in airspeed and a decrease in altitude—are unexpected and hence predicted to be highly salient. The third cue, zero pitch, is consistent with the constructed model, and therefore not as likely to attract attention.

A second possible cause for the kind of inordinate attention given the airspeed and altitude cues is indicated by a remark of Bower, Black, and Turner (1979), who state that: "according to schema theory the understander must commit himself to some initial schema in order to understand sentences; yet the most diagnostic information may not appear in the text until later. That is, one can be led down 'garden path' stories" (p. 216). Given the predicted salience of the increasing airspeed and decreasing altitude (the unexpected events) and the fact that these two cues are presented first, the pilot may have already activated a frame for descent before hearing about the zero pitch.

If the activated frame (DESCENT) instructed the pilot to consider the reading on the vertical speed indicator (Subjects 2 and 10) or to consider a possible power loss (Subjects 4, 6, and 18), the information about zero pitch might easily have been ignored as irrelevant. If, on the other hand, the activated frame instructed the pilot to consider the reading on the artificial horizon (Subjects 1 and 5), the salience of the third instrument indication, zero pitch, would be increased and the cue would likely be noticed. In order to avoid information overload, the pilots may have used these types of mechanisms to focus attention selectively on some subset of the cues available in the scenario (Sheridan, 1981).

Thus, it is possible to account for the performances of those pilots who focused on the changes in altitude and airspeed in terms of known models of human cognition. The fact that other pilots focused on different subsets of the available cues (see Figure 1) is more difficult to explain with the available data. Factors such as expectancies and priming (Meyer, Schvaneveldt, and Ruddy, 1975; Morton, 1970; Rummelhart and Siple, 1974) may have played a role, however.

Memory distortion. One pilot provided a very interesting example of how people not only ignore some available data, but may even distort their recall of other available data (Arkes and Harkness, 1980; Bartlett, 1932) to make it consistent with the activated frame. This subject reported that he initially thought the plane was in a nose-level descent (Task 4). Immediately after hearing the scenario, he said "Let me get this straight now: *Increasing airspeed*, decreasing altitude and you mean pitch as far as being above or below the horizon based on the artificial horizon?" [Answer from experimenter: "Yes."] "What's happening to my power? Very definitely we have a situation where we seem to be losing power. The fact that we're decreasing in altitude and our *airspeed is constant* indicates that we are basically in a situation where we are losing altitude. It would stay fairly constant if we're coming down."

These data suggest the following sequence:

(1) The pilot listened to and noted all three of the symptoms. (Note that he explicitly mentioned the increasing airspeed.)
(2) He then tried to identify a problem(s) consistent with this data. During this initial frame activation process, he focused his attention on the decreasing altitude and zero pitch, and consequently activated a frame for a nose-level descent (indicated directly in the data from Task 4 and implied by the memory distortion to be discussed next).
(3) He asked himself what could cause a nose-level descent and concluded it could be a power loss.
(4) He reviewed the available data to make sure they were consistent with a nose-level descent due to a power loss. His recall that "we're decreasing in altitude and our airspeed is constant" was in fact consistent with such a situation.

Note that within a time span of less than 30 seconds the pilot has distorted his recall. Originally he stated that there was an indication of increasing airspeed. After activating the frame for a nose-level descent, he stated that the airspeed was constant. This indicates a rather self-defeating process. The nose-level descent frame, which the pilot is trying to test by reviewing the symptoms he has heard, is being used to help recall or reconstruct the set of symptoms. The role of the activated frame is so powerful in this recall/reconstruction process that the pilot "remembers" symptoms consistent with that frame rather than the symptoms actually presented.

Organization of the knowledge structures. The previous analysis identified ten different frame labels based on the pilots' first queries and associated statements. Applying the same form of analysis to the remainder of the verbal protocols, there is evidence for eight additional frames: *artificial horizon malfunction, structural icing (wings), trim wheel mispositioned, downdraft, gear down, flaps down, banked plane, and yoke mispositioned.*

All of the queries made by the 20 pilots can be accounted for in terms of attempts to access or to test the validity of the 18 frames that have been defined. The performances of the pilots can be modeled by linking these individual frames into a frame system, with the links occurring through the *causes* slots. Thus, an instruction to check for possible causes of a nose-down descent results in checking the *causes* slot of the appropriate frame (see Figure 2). This leads to the activation of the structural icing frame, which may then instruct the pilot to check for visible ice and for expected instrument readings. Thus, frame activation and query generation is controlled by focusing attention on one of the slots in the currently activated frame and using the information present either to activate a new frame (a possible cause) or to generate a query (ask about a particular instrument reading). The organization of the frame system thereby serves to generate hypotheses regarding the cause of the existing problem and allows the pilot to diagnose the problem at different levels of specificity.

Directing the problem-solving process. The preceding sections identify knowledge structures (frames, triggers or enabling events for frames, slots for frames, and links among frames) consistent with the available data (the verbal protocols and information requests). They also suggest that when a given frame is activated it instructs the pilot to either make sure that frame is a valid representation of the state of nature (check expectations) or look for a possible cause of that state of nature.

The fault diagnosis performances observed can be described as a process of recursively identifying possible problems and their causes by activating corresponding frames until the person decides that he or she has found the initiating cause. The data indicate that, in order to drive this process, subjects attempted to answer six types of questions:

(1) Is the currently activated frame a goal-state? (Have I diagnosed the problem?)
(2) What is the cause of the state of nature represented by this frame?
(3) Is this frame a valid representation of the state of nature? Are the expected instrument readings and conditions present?
(4) If the currently activated frame has been rejected as a possible state of nature, can I find another frame to activate?
(5) Is there a recall error?
(6) Is there an instrument malfunction?

Patterns of Performance

Group A consists of the subjects who diagnosed the problem as a vacuum system failure. Group B pilots concluded that there was an artificial horizon malfunction. Group C concluded the problem was a downdraft. Group D pilots concluded that the cause of the problem could not be determined with the available information.

Group E (Subject 20) discovered the presence of the artificial horizon malfunction. At that point he thought the plane was in a straight, nose-down descent that was not being indicated on the artificial horizon. He asked what would happen if he applied backpressure on the yoke to arrest the descent. When the expected response did not occur (because the plane was actually in a left bank), he decided there must be some other problem. He failed to discover the left bank and concluded that he could not determine what the problem was.

Groups A, B, C, and D divide the subjects into four classes according to their final conclusions. (Group E is really a special case of Group B.) The most apparent differences among the groups are the contexts in which the six alternative questions (check to see if done, check for causes, test expectations, look for new frame, check for memory error, or consider instrument malfunction) are addressed. Group B, for instance, differs from Group A by its failure to check for possible causes of the artificial horizon malfuncton.

Cognitive narrowing. The failure of Group B to determine whether a vacuum pump failure was causing the artificial horizon malfunction can be explained in terms of the ordering of the six questions by application priority. The use of the following simple rule would almost certainly have caused all the Group B pilots to discover the vacuum pump failure: *Always check for possible causes of the state of nature represented by the currently activated frame before asking whether it can be used as a final diagnosis.* During Task 3, all of the pilots in Group B demonstrated that they had knowledge of the relationship between the vacuum system and the functioning of the artificial horizon. This explanation of Group B's failure to seek a broader systemic cause is consistent with pilots' explanations at the end of the experiment as to why they stopped without asking about the suction gauge: "I just narrowed my vision down to one area, tunneled my vision down, and stopped." This finding is also consistent with a known bias of human operators to produce "a sort of 'cognitive tunnel vision' (Sheridan, 1981) in which operators fail to encode or process information that is contradictory to or inconsistent with the initially formulated hypothesis" (Wickens, 1984, p. 97).

In the case of the present study, however, the correct hypothesis is not contradictory to the one generated by the subjects. The artificial horizon is malfunctioning. The failure of the subjects lies in their assuming that this localized malfunction is the sole problem. They have not considered the possibility that this hypothesis, if true, could itself be evidence of a broader, systemic problem. In terms of the frame system model, the failure results from focusing attention exclusively on the *expectations* slot.

Activation of default values. All three pilots in Group D acted as if they had activated the nose-down descent frame. They then proceeded to check for possible causes of the descent. When they failed to find a cause, they stopped and concluded that the cause of the problem could not be determined with the available information.

Given that the plane was in a nose-down descent, what accounts for their failure to find the cause? The answer may lie in the activation of a default value. All three of the pilots reported that they thought the plane was in a straight nose-down descent (Task 4). Subject 18 even reported visualizing the turn and bank indicator, relating that the "turn and bank indicator showed straight and level." (In actuality he had been given no information about the turn and bank indicator, which showed a left bank.)

The data suggest that the pilots in Group D activated a default value—that the descent was straight ahead—for the direction of the nose-down descent. They did so in the absence of any data to support this assumption. (On the other hand, the plane had been cruising straight ahead, and they had not received any information clearly indicating a turn.)

The activation of this default value rules out the actual cause of the descent, which was a banked plane. (Activation of such a default value could accomplish this by activating a new frame representing a *straight,* nose-down descent.) None of the subjects in Group D considered this as a possible cause of the descent.

If it is assumed that these pilots activated a frame representing a straight, nose-down descent, then the observed behavior could be construed as an example of a confirmation bias (a tendency to collect data that is consistent with the hypothesis under consideration and to avoid collection of data that might lead to the rejection of this hypothesis). Such a bias has been found in a number of studies of human decision making (Einhorn and Hogarth, 1978; Mynatt et al., 1977; Schustack and Sternberg, 1981). Modeling performance in terms of the activation of a frame system provides insight into the mechanism causing such a bias in this experiment. Pilots failed to consider a left bank because (1) an incorrect frame (straight, nose-down descent) was activated as a result of the activation of an invalid default value, and (2) a left bank was not a possible cause of this activated frame. (It was not a possible slot-filler in the *causes* slot.) These pilots never thought to question the validity of the activated default value for a straight descent.

Top-down refinement. Some knowledge-based systems use a top-down refinement strategy (Hasling et al., 1984) in order to guide the diagnosis process. Such a problem-solving process "can be characterized as an 'establish-refine' type. Each concept first tries to establish or reject itself. If it succeeds in establishing itself, the refinement process consists of seeing which of its successors can establish itself" (Chandrasekaran, 1983). In the case of the problem-solving task discussed in this paper, such a top-down strategy would imply, for example, establishing that the plane is in a straight, nose-down descent (checking the expectations for that frame) *before* looking for "successors" or causes of such a state.

The pilots in Group D (those concluding that the problem could not be diagnosed) differed from the other pilots in that they never attempted to establish that the frame they had activated (straight, nose-down descent) was valid. They never checked to see if the expectations listed by that frame were present. Instead, they immediately focused attention on the *causes* slot. Then, when none of the possible causes of a straight, nose-down descent was found to be present, they concluded that the problem could not be diagnosed. (In addition, none of the pilots in Group D recalled the indication of zero pitch.)

The pilots in Groups A, B, and E, on the other hand, all checked expectations of the straight, nose-down descent frame, resulting in the identification of a problem with the reading on the artificial horizon. Such findings suggest that a top-down refinement strategy may protect against activation of an incorrect frame due to: (1) an inability to attend to all of the available cues (ignoring the zero-pitch reading); (2) activation of an invalid default value (the plane is descending straight); (3) memory losses (the zero-pitch information was not only ignored, it was forgotten by the pilots when focusing attention on the straight nose-down descent frame).

Thus, this analysis of performance in terms of knowledge representations highlights those mechanisms by which people might "adopt a natural bias to retain an old hypothesis rather than go to the trouble of formulating a new one" (Wickens, 1984, p. 99). The pilots in Group D are retaining the same hypothesis/frame by focusing attention on possible causes rather than on establishing the frame's validity through the checking of expectations. Note that *it is the structure of the frame that segregates these two classes of information.*

GENERAL DISCUSSION

A frame system has been proposed that accounts for the performances of 20 pilots on a fault-diagnosis task. This task was specifically designed to permit pilots to use their domain-specific knowledge in order to diagnose a problem.

It was found that 18 frames, all having a common structure, are sufficient to explain the data. These frames represent prototypical states of nature (nose-down descent, blocked static port, etc.). With the exception of the memory error frame, each frame has an associated set of triggers (external cues or another frame), a label, and two slots. These frames represent possible physical states of a plane or of some part of a plane. One slot represents possible causes of the state of nature represented by that frame (e.g., power loss can cause descent) and provides links among frames. The other slot represents expected instrument readings and observable conditions if that state of nature exists (e.g., the vertical speed indicator should indicate a descent if the plane is in a descent).

The present study suggests that there may be significant changes in human problem-solving performance when relevant domain-specific knowledge is available in the form of highly organized data structures (in memory). Like the subjects in other studies (Mynatt et al., 1977; Wickens, 1984), these pilots quickly formulated a working hypothesis and focused on it alone rather than comparing competing hypotheses. Unlike what was reported in the other studies, however, there were frequent requests for information that could and did lead to the rejection of the working hypothesis (the activated frame). For example, all seven pilots initially activating a frame representing some pitot-static system malfunction (see Figure 1) quickly requested information that led to the rejection of that working hypothesis.

Furthermore, this study suggests that the failure to employ a top-down refinement strategy (which resulted in the acceptance of activated default values without question) may be a critical factor in the failure of subjects to diagnose the cause of the problem. This failure resulted from the focusing of attention on the *causes* slot rather than on the *expectations* slot. Although this is consistent with other findings demonstrating that subjects have a bias to retain an old hypothesis, the analysis of performance in terms of knowledge and control structures offers a much different perspective than do explanations in terms of cognitive effort or cost of thinking (Einhorn and Hogarth, 1981). The results of this study imply that it may be very fruitful to investigate further the role that knowledge structures play in influencing cognitive biases and strategies.

ACKNOWLEDGMENTS

Funds for the support of this research have been provided by the NASA-AMES Research Center, Grant No. NAG2-112.

REFERENCES

Aikens, J. (1983). Prototypical knowledge for expert systems. *Artificial Intelligence, 20,* 163-210.

Arkes, H., and Harkness, R. R. (1980). The effect of making a diagnosis on subsequent recognition of symptoms. *Journal of Experimental Psychology: Human Learning and Memory, 6,* 568-575.

Bartlett, F. C. (1932). *Remembering.* Cambridge, U.K.: Cambridge University Press.

Bower, G., Black, J., and Turner, T. (1979). Scripts in memory for text. *Cognitive Psychology, 11,* 177-220.

Chandrasekaran, B. (1983 Winter/Spring). Towards a taxonomy of problem solving types. *AI Magazine,* pp. 9-17.

Dember, W., and Warm, J. (1979). *Psychology of perception* (2nd ed.). New York: Holt, Rinehart and Winston.

Einhorn, H., and Hogarth, R. (1978). Confidence in judgment: Persistence of the illusion of validity. *Psychological Review, 85,* 395-416.

Gentner, D., and Stevens, A. (Eds.). (1983). *Mental models.* Hillsdale, NJ: Erlbaum.

Hasling, D. W., Clancey, W. J., and Rennels, G. (1984). Strategic explanations for a diagnostic consultation system. *International Journal for Man-Machine Studies, 20,* 3-19.

Meyer, D. Schvaneveldt, R., and Ruddy, M. (1975). Loci of contextual effects on visual word recognition. In P. Rabbit, and S. Dornic (Eds.), *Attention and performance V.* London: Academic Press.

Minsky, M. (1975). A framework for representing knowledge. In P. Winston, (Ed.), *The psychology of computer vision* (pp. 221-277). New York: McGraw-Hill.

Morton, J. (1970). A functional model for memory. In D. A. Norman (Ed.), *Models of human memory* (pp. 203-254). New York: Academic Press.

Mynatt, C., Doherty, M., and Tweney, R. (1977). Confirmation bias in a simulated research environment: An experimental study of scientific inference. *Quarterly Journal of Experimental Psychology, 29,* 85-95.

Norman, D. (1983). *Learning and memory.* San Francisco: W. H. Freeman.

Pauker, S., Gorry, G., Kassirer, J., and Schwartz, W.

(1976). Towards the simulation of clinical cognition. *The American Journal of Medicine, 60*, 981-996.

Rasmussen, J., and Rouse, W. (Eds.). (1981). *Human detection and diagnosis of system failures.* New York: Plenum.

Rummelhart, D., and Siple, P. (1974). Process of recognizing tachistoscopically presented words. *Psychological Review, 81*, 99-118.

Schustack, M., and Sternberg, R. (1981). Evaluation of evidence in causal inference. *Journal of Experimental Psychology: General, 110*, 101-120.

Sheridan, T. (1981). Understanding human error and aiding human diagnostic behavior in nuclear power plants. In J. Rasmussen and W. Rouse (Eds.), *Human detection and diagnosis of system failures.* New York: Plenum.

Wickens, C. (1984). *Engineering psychology and human performance.* Columbus, OH: Merrill.