Conceptual Models for Understanding an Encounter with a Mountain Wave

Edwin Hutchins
University of California, San Diego
ehutchins@ucsd.edu

Barbara Holder
University of California, San Diego
bholder@cogsci.ucsd.edu

ABSTRACT

Conceptual models are one of many resources available to pilots for making sense of the flight environment. In this paper we describe the conceptual models a pilot uses, in interviews, to explicate his encounter with a mountain wave while flying an Airbus airplane on the line. In the discourse four models emerge: mountain wave, thrust to control speed, pitch to control speed, and climb-descend to control speed. The models he utilizes in his descriptions have a different conceptual structure from models presented in training materials. The analysis suggests that the pilot's conceptual models have an operational organization, and this structure is somehow adapted for managing specific flight situations.

Keywords

Conceptual models, autoflight, flight training

As part of an on-going study of the cognitive consequences of flight deck automation we are attempting to track the development of pilots' conceptual models of autoflight system operation through their flying careers. It is important to understand the structure of conceptual models pilots actually use because these models, unlike engineering models, are very likely to be adapted to the conditions of use on the line. Differences between the models that are provided in training and the models that are used on the line may indicate that the organization of the models provided in training are somehow not fitted to or adapted to use on the line. In this paper we present excerpts from two interviews with one of our participating pilots who encountered a mountain wave. The pilot contacted us shortly after the event and we arranged the initial interview in which he described the event. We subsequently contacted him and conducted a second interview to clarify some specific points made in the first interview.

LOCATING CONCEPTUAL MODELS IN DISCOURSE

Interviews such as the ones presented here do not permit us to know what actually happened. The pilot may be mistaken about the course of events, he may have wrongly remembered them, or he may have been unaware of important aspects of the events. In the most radical sense, we cannot even know *that* the event happened. It would be very surprising for a pilot to construct an account like the one we see here if, in fact, nothing like this actually happened. However, we have no tangible evidence about the event aside from the account offered by the pilot. If we cannot know what happened, or even whether anything happened at all, what can the interviews tell us?

To see what interview data *can* tell us, we need to consider the processes involved in the production of narrative accounts. Communication is grounded in shared models of the world [1,2,3,4]. For example, when thinking about how minds work, Americans believe that experiences lead to feelings, feelings may lead to intentions, and intentions may lead to actions [5]. When attempting to interpret the behaviors of others (or ourselves for that matter) we often do so in terms of inferred (but unobserved) intentions and feelings. A pilot may say, "He asked to descend early because he wanted to be sure to make the altitude restriction at the VOR." In saying that, the speaker is constructing a link between some observed behavior and an unobserved intention. The link is forged from two conceptual models. One model establishes the relevance of early descent to compliance with an altitude restriction. The terms of this model are a set of beliefs about the behaviors of airplanes. Some pilots refer to it colloquially as the "dive and drive" maneuver. This model is universally shared in the pilot community, but is not part of the everyday thinking of non-pilots, who may be puzzled by the idea that it can be difficult to get an airplane to a lower altitude. The second model establishes the relevance of the *behavior*, requesting an early descent, to the *intention*, making the altitude restriction. The behavior is understood to be undertaken because it is a way to satisfy the intention. This is a very widely shared model of the operation of the mind. People account for behavior in terms of intentions, and while they may question whether the intentions have been correctly identified, they normally do not question the strategy of explaining behavior in terms of intentions [6]. Together these two models connect the elements of the statement as meaningful parts of understood worlds of action. Neither model is explicitly stated. Rather, they are implicitly presupposed by the speaker to be available to the listener as resources that can be used to understand the statements.

The target of our analysis is not the true nature of the events that are being described. We have no reason to suspect that our informant wishes to deceive us, but we also have no independent means of assessing the accuracy of the account. Instead, the targets of our analysis are the conceptual models that are used to render the account meaningful.

OPERATIONAL MODELS

There are many possible ways to organize a conceptual model of autoflight operations. Systems engineering descriptions typically emphasize the components of a system and the relationships among the components. Aerodynamical models emphasize processes involved in changing the airplane's energy. Airplane operating handbooks present component based models and may add operating characteristics of particular functions of a system. Descriptions are sometimes given in terms of features that may be used to discriminate among functional states of the system. In the case of the Airbus airplanes for example, the autoflight system is presented as a complex hardware system in which various kinds of information move among the components and the components perform a specific set of computations. The operating manual for the Airbus A320 includes feature descriptions of the space of flight guidance modes, distinguishing lateral from vertical modes and managed from selected modes in a crosscutting arrangement. The operating manual also makes explicit a crucial dependency between conceptual features of vertical navigation modes and autothrust modes. If the performance target for a vertical navigation mode is a speed, then the associated autothrust mode will have a thrust limit as its target. If the performance target for a vertical navigation mode is a trajectory, then the associated autothrust mode will have a speed as its target. However, at least in talking about autoflight guidance modes, pilots seem to prefer a different sort of conceptual model. They use an operational model that is organized around particular functional properties of the various modes in specific flight situations. For example, one pilot described the difference between managed and selected (open) modes.

I: what sorts of things would you use the managed // modes for.

P: um // difficult uh departure or arrivals that have-uh restrictions? to be-be at certain fixes at certain altitudes-(are often) managed. the o-and then the open mode is just used when you need to expedite more or less you know like when you're in a hurry to get down or // go to open if-you know you're gonna (unint) // best // rates of climb and descent and all that.

Another pilot offered a similar distinction.

P: to where now it's basically y-if if you have an arrival to to fly you go ahead and go in managed descent mode in order for it to calculate all all the cal-uh all the crossing restrictions, and then once you get down to about ten thousand feet or fourteen thousand feet they just give you a-a uh a descent, (1.5) either going into open descent just to go ahead and go down, um // as fast as possible or go into a vertical // speed descent, if the uh // if-if if you're still fairly far out.

In both of these cases, the pilot is matching the operational characteristics of flight guidance modes to particular conditions or problems encountered in flight. The key feature of the managed modes for these pilots is that these modes compute and comply with altitude restrictions that are associated with particular waypoints in the flight plan. The key feature of the open modes for these pilots is that they can be used to produce maximum rates of climb or descent. Open descent mode is especially marked as the mode to use when one is "in a hurry to get down" or one wants to "go down as fast as possible." These functional models are specific to a class of flight situations, often associated with a phase of flight, and involve strategies for managing the task given the specific flight context. We observe similar functional models in the interview data with all of our participants.

ENCOUNTER WITH A MOUNTAIN WAVE

The pilot described flying eastbound in an Airbus A320 at FL370 on a night flight over the Rocky Mountains. He said he heard airplanes ahead of his reporting variations of three hundred feet of altitude and plus or minus 30 knots of speed in mountain wave conditions. Mountain wave is described in the Airman's Information Manual:

> Mountain waves occur when air is being blown over a mountain range or even the ridge of a sharp bluff area. As the air hits the upwind side of the range, it starts to climb, thus creating what is generally a smooth updraft which turns into a turbulent downdraft as the air passes the crest of the ridge. From this point, for many miles downwind, there will be a series of downdrafts and updrafts. Satellite photos of the Rockies have shown mountain waves extending as far as 700 miles downwind of the range [7].

To understand how mountain wave conditions can produce variations in altitude and airspeed, and especially to understand the relations between altitude excursions and airspeed excursions, conceptual models are required. It is possible to construct these models in many ways.

The pilot gave the following description of the plan he formed to handle this situation.

P: so we-we requested a block. I-it was the first officer's leg, and I requested a block thirty five to thirty seven. and I asked him to go ahead and fly thirty six, [and that] he was to protect the airspeed, and just fly the wave. 'kay so let the wave you know bump us up a thousand feet, and then let him-let the wave bring us back down two thousand feet and just hold the speed. okay just protect speed.

Experienced pilots will, no doubt, fill in the details of the relationships between airspeed and altitude in mountain wave conditions. A key concept here is that airspeed is something that needs protecting. High altitude flight takes place in a small region of speed space between a stall and an overspeed condition [8]. Either the stall or the overspeed condition threatens the control of the airplane. Thus, it is necessary to protect airspeed. So far, we can see that airspeed has some relation to altitude in this event, and that the FO has been instructed to allow altitude to vary as required to maintain airspeed relatively constant ("protect airspeed" and "hold the speed").

The relation of airspeed to other elements of the mountain wave model become clearer in the pilot's description of the event, but first let us summarize the pilot's account of the wave encounter and then step through the elements of the account. The airplane entered the mountain wave at FL360 with the autopilot engaged in altitude holding, ALT CRZ, mode. The autothrust system was engaged in SPD/MACH mode, maintaining a reduced cruising speed of .75 Mach. The onset of the wave was more abrupt than the pilot had ever experienced before. The airplane rapidly accelerated to the Maximum Mach number (MMO = .82) and then exceeded that value. The overspeed warning was triggered (aural tone, Master Warning light illumination, and ECAM message) when the airplane passed MMO. As the airplane accelerated through .83 Mach, the overspeed protections engaged. This caused the autopilot to disconnect, and caused the airplane to pitch up. The first oscillation of the wave was flown by the protection systems. The FO (controlling pitch) flew two subsequent updraft and downdraft oscillations by hand with autothrust system engaged in SPD/Mach mode.

The pilot said that he instructed the FO to protect the speed and then described the entry to the mountain wave as follows:

P: okay. so he says okay no problem we-we're cruising along and uh // I mean just out of nowhere (1.5) boom we're in this wave, the speed trend arrow? are you familiar with it?

I: yeah I know what it looks like.

P: okay the speed trend arrow ***pegs off the scale,*** with this increase in speed.

The key elements of this description are the suddenness of the onset of the event and the dramatic increase in speed that it caused.

THE USE THRUST TO CONTROL SPEED MODEL

A little later in the interview the pilot returned to the onset of the event and introduced the rate of reduction of thrust by the autothrust system as a relevant factor in the event.

P: so // the speed goes crazy. in ***my*** opinion (2.0) the thrust reduction required was not su-a-a-abrupt enough.

This statement presupposes that the rate of thrust reduction should somehow match the rate of speed increase. This is based on a basic pilot model of the relation of thrust to speed in altitude holding modes. When the airplane is kept at a constant altitude (by whatever means) changes in thrust are mirrored by changes in speed. Reduction in thrust produces reduction in speed. He then elaborates on the role of the rate of thrust reduction commanded by the autothrust system and its relation to the increase in speed.

P: the airplane // should have pulled all of the power off with the-you know a significant speed increase like we had. and it didn't do that. and I think that // that was partially the cause for the overspeed initially. was slow uh-uh slow response by the autothrust.

Use of the phrase "the airplane should have" marks that it is the autothrust system that is commanding thrust at this point. It also marks an implicit disjunction with what was observed, that the airplane did not seem to "pull all of the power off." The focus on the rate of response by the autothrust is an embellishment on the basic model of the relation of thrust to speed. Now the relative rates of change of thrust and speed are being linked. Under normal flight conditions, this model could be invoked as an explanation of the acceleration or deceleration of an airplane in response to changes in thrust. In the mountain wave situation, the movement of the air stream in the mountain wave causes changes in the speed of the airplane which can be **countered by** the changes in thrust that are commanded by the autothrust system.

In the following segment, the pilot compares the behavior of the autothrust system to what he could have done via manual control of the thrust levers to counter the acceleration that was induced by the mountain wave. Having established the rate of change of thrust as the key variable, he claims that he could have produced a more rapid – and, therefore, more appropriate – reduction in thrust.

P: uh and I really think that if I had had manual control of the thrust levers, that I think I could've done a little better. I'm not saying that we probably would not have // gone into overspeed but I-I'm pretty sure that // it-it would not have uh gone to a full overspeed condition.

He later said that in retrospect it would have been better had he taken manual control of the thrust to make a reduction to idle thrust. He also noted that doing that would conflict with his company's procedures.

> P: uh from a procedure standpoint (1.5) you know if I had to do it all over again knowing what I know now (2.0) at the first indication of the overspeed? I would've just // you know disconnect the autothrust. Go to idle. = and uh and then control the autothrust manually ***but*** // that conflicts with *((airline #2))* procedures.

He then compares what he did not do in this airplane to what he would have done in other airplanes. He implies that he would not have experienced an extended exposure to the overspeed condition in other airplanes because he would have been able to produce a rate of thrust reduction that matched the increase in speed.

P: But again, you know on any other airplane// instinctively I would've just brought the throttles back with that rate of increase.

Up to this point, we have seen several instantiations of a model in which the rate of thrust reduction is linked to the rate of increase in speed. In the following segment, he hints at a second model for the control of speed, this one involving changes in altitude.

P: 'cause the airplane // when the autopilot is on, the airplane basically is being told one thing, maintain your altitude.

The mention of maintaining altitude anticipates a model of speed control in which change in altitude is linked to change in speed. Since the airplane entered the wave event in an altitude hold mode, this model is not yet active. The relevance of this statement to what follows is that the potential contribution of changing altitude to controlling the speed will not be available while the airplane is in an altitude hold mode. An implication of this is that, all of the speed control that is available while in altitude holding mode must be provided by varying thrust.

P: and and because the wave was so abrupt (1.5) you know you were in it before you knew what had just happened. So // you know when the speed starts increasing, you know that thrust should start coming back, but the the-the rate at which the speed increased apparently was significantly more than the rate that the // the thrust you know would've decreased (in-) for the con- =given conditions.

The pilot argues that because of the abrupt nature of the event, and given the slow rate of reduction in thrust provided by the autothrust system, thrust reduction alone was not sufficient to protect the airspeed.

THE CLIMB OR DESCEND TO CONTROL SPEED MODEL

The pilot began his narrative by recounting his hearing airplanes ahead of him reporting altitude changes in response to the wave. In the pilot's description of the plan for dealing with the mountain wave, a block altitude is requested and the FO is instructed to:

P: 'kay so let the wave you know bump us up a thousand feet, and then let him-let the wave bring us back down two thousand feet and just hold the speed. okay just protect speed.

This excerpt implies that a change in altitude is anticipated as part of the response to the wave event. The pilot's plan going into the mountain wave hints at the relation between change in altitude and protection of speed.

Since the wave is conceived as actively bumping the airplane up, a simplified model in which the airplane just goes up and down with the wave without any specific effects on speed cannot be ruled out on the basis of this statement alone. Subsequent statements indicate that a more complex model is at work. Describing the second and third oscillations, which were hand flown, he says,

P: and the speed was well within the normal range, but we're flying the wave. so you know we went through with the climb and then we came back with a descent.

The pilot maintains that they were able to control the speed on the second and third oscillations. The disjunction 'but' marks the notion that the control of speed did not simply happen, it happened because they were flying the wave. He then elaborates on the aspect of flying the wave that is most relevant to speed control, that is, changing altitude with the up and downdraft portions of the wave. Here we begin to see the introduction of the relevance of altitude change to speed control.

> P: we uh we rode up to thirty seven, and then we rode back down to thirty five five. although I'm pretty sure we could've done it with less altitude // uh excursion?

The reference of 'it' in the last line above is protecting the speed. So again the pilot uses the change altitude to protect speed model. In addition, the mention of the possibility of having controlled speed with "less altitude excursion" suggests that there may be a proportionality aspect to this model. That is, the magnitude of the altitude change is seen as proportional to the magnitude of the speed change.

The perceived need to climb in order to counter the speed increase is explicitly stated in the following excerpt.

P: We figured we're gonna fly the block so // we got to thirty six we get the first oscillation speed increase so we wanna climb I put thirty seven up there because we wanted to climb

In the follow-up interview, the pilot offered a clarification of the plan. Because the altitude hold mode would not permit the airplane to use change in altitude to control airspeed, yet the crew had briefed the idea of “flying the wave”, we asked why they had entered the event with the autopilot engaged in an altitude hold mode?

I: you went in on the autopilot but you were gonna actually instruct it to climb early in the event or something right?

P: well as soon as we got whichever oscillation we got …so all we were gonna do is set the new altitude whether we had to do up or down on the FCU and then hit open climb, open descent, vertical speed, you know whatever it took.

Finally, in imagining what he would do were he to encounter another mountain wave, the pilot says,

P: number one the thrust is gonna do whatever I'm gonna manually make it do. Number two the autopilot will be intentionally disconnected so that I could fly the block.

In this segment we have a neat combination of the thrust model and the altitude model.

THE USE PITCH TO CONTROL SPEED MODEL

In describing the way that the airplane’s protections prevented the speed from going much beyond the maximum Mach number the pilot said the following.

P: that’s what impressed me about this because the-the first // nose up command // to prevent the overspeed okay? the second-well and then the-first speed loss command, you know at the top of the of the wave of-on the first oscillation, when the speed started increasing, the first nose down command was also issued by the aircraft’s protective-system.

This segment is built on a conceptual model in which nose-up pitch attitude is associated with decrease in speed and nose-down pitch attitude is associated with increase in speed. In the pilot’s view, when the autothrust system was unable to control the increase in speed induced by the encounter with the mountain wave, the airspeed increased beyond the limits that trigger the aircraft’s built-in envelope protections. At that point, the airplane pitched up in order to control the speed. The conceptual model relating pitch attitude to speed comes into play to interpret the behavior of the overspeed protection. This model is shared among all pilots.

In his description of the way the FO flew the second oscillation of the wave this model is again used in a refined form.

P: so you know // um i-as soon as we saw the trend arrow going-you know from a little bit above VLS on up, we knew what we were faced with, so he started nice and early. you know. um (1.5) and then // you know again at the top, coming back the other way, the same thing. as soon as the uh trend arrow reversed itself, he just shoved the nose back over, (and-) we were in business.

The refinement of the model is that the FO is responding to the direction of the speed trend even before the target value has been reached. Thus the FO could begin pitching up while still below the target speed as long as he could see that the speed trend indicated acceleration. We assume that at this point in the event, the A/THR is also controlling speed because “If no AP/FD mode is engaged – the A/THR system controls SPD/MACH mode.” According to the pilot’s account, no AP/FD mode could be engaged at this point because the autopilot disengaged when the overspeed protections engaged.

The mountain wave model produces a set of expectations about the shape of the path of the airplane through the rising and falling stream of air. The mountain wave model also generates an expectation that airspeed will need to be protected, but it does not specify how that protection is to be accomplished. Three additional models are introduced to explain the relation of airspeed to the other elements of the mountain wave event. One model links thrust to speed and links rate of change of thrust to rate of change of speed. A second model links changes in altitude to the control of speed. The third model links pitch attitude to change of speed. The second and third models are applied independently at many points in the interviews. They are, however, also linked conceptually to each other. The following two excerpts establish the relation of the pitch model to the altitude model. The first simply describes the relationship of pitch to holding altitude in the updraft portion of the wave.

P: alright. so now as we go to the first oscillation the-the-the wind comes and starts going up over the Rockies right? so that is an updraft. for you to maintain your altitude, you're gonna point the nose down.

A nose-down pitch attitude is required to maintain altitude in the updraft of the wave. The pitch to speed model predicts that nose down pitch attitude will produce an increase in speed. And remember the airplane entered the wave with the autopilot engaged in an altitude holding mode. The next excerpt describes what happened to other airplanes on that route that night that did not request a block altitude.

P: everybody else // tried to stay at their altitude so their speed excursions were // bigger.

The pilot maintains that those who attempted to stay at altitude had big speed excursions - because they had to put the nose down to stay at altitude. The only way to avoid the speed excursions would be to avoid putting the nose down, but in the updraft of the mountain wave, that means one must climb. This is the justification for requesting the block altitude.

DISCUSSION

The mountain wave is an interesting test case because it is an event that is complicated, and involves the aerodynamics of the airplane. Most engineers would say it is an interesting event from an energy management standpoint, and so there is an opportunity in this event for the pilot to present the kinds of models, energy management for example, that we might expect from one of our engineering colleagues. What we see in fact is that even though it is an event that would be easy to describe that way, that's not the way our pilot describes it. In fact he describes it terms of these much simpler operational models. Each one accounts for the relationships among critical parameters of flight in idealized and simplified contexts. Each model has an implicit *ceteris paribus* (other things being equal) clause. Thus, the model of the relationship of thrust to speed invokes a simplified world in which the airplane is in level flight with no unusual outside influences acting on it. The model of the relationship of pitch attitude to speed assumes thrust and all other parameters are held constant while pitch is varied. The model of the relationship of altitude change to speed predicts that if thrust is held constant in a stable atmosphere, climbing will result in a decrease in speed.

These basic models are not specific to automated airplanes or to mountain waves. They describe general behavioral characteristics of all airplanes in stable atmospheric conditions. It is clear, however, that they are used by the pilot in these interviews to construct an understanding of the behavior of the autoflight and envelope protection systems. These models are also used to account for the behavior of the airplane in the mountain wave event. The conditions of the mountain wave violate the stable atmosphere assumption of the models, so the models must be adapted for use in this setting. First, the model of the relation of pitch attitude to speed is used to explain the acceleration that the airplane experienced in the updraft portion of the wave. Then, the decelerating effects of reduction in thrust, pitch up attitude, and climbing are all seen as elements that can counter the acceleration that was induced by the mountain wave.

The importance of examining the kinds of models that pilots use is that this is how they are making sense out of the world. We assume that the types of models that pilots develop for use on the line are somehow adapted to the conditions encountered on the line. If we understood the structural or organizational characteristics of models that do seem to be adapted to use on the line we might be able to restructure models in training. The goal would not be to train pilots on the models that are observed in use on the line. Rather, we hope to understand the properties of the models that are observed in use on the line that make those models appropriate for use in flight. With an understanding of these properties, we might be able to provide pilots leaving the training center with a better set of tools to deal with the problems they face.

ACKNOWLEDGMENTS

The research reported here is supported by a grant from the NASA Ames Research Center.

REFERENCES

1. Hutchins, E. 1980. *Culture and Inference*. Harvard University Press.
2. Holland, D and N. Quinn, eds. 1987. *Cultural Models in Language and Thought*. Cambridge University Press.
3. Shore, B. 1996. *Culture in Mind*. Oxford University Press.
4. Strauss, C. and N. Quinn 1997. *A cognitive theory of cultural meaning*. Cambridge University Press.
5. D'Andrade, R. 1987. A folk model of the mind. In *Cultural Models in Language and Thought*. Eds., Dorothy Holland and Naomi Quinn. Cambridge University Press, pp. 112-148.
6. D'Andrade, R. 1989. Cultural Cognition. In *Foundations of Cognitive Science*. Ed., Michael I. Posner, Cambridge: MIT press, pp. 795-830.
7. Jeppesen Sanderson AIM, *Safety of flight*. US 7-5-3
8. Webb, J. 1971. *Fly the Wing*. Iowa State University Press, Ames, Iowa.