Reconceptualizing Expertise
Explaining an Expert's Error

Stephen Deutsch
BBN Technologies
sdeutsch@bbn.com

ABSTRACT

A human performance model that produces credible human-like behaviors requires an architecture and implementation strongly grounded in theory. Through an iterative process of model development and theory refinement, a better understanding of the sources of human error can then be achieved. Aircrew and air traffic controller models were used in an analysis of the 1994 windshear accident at Charlotte, NC. The analysis led to a focus on the captain's briefing for the plan to address the immanent windshear threat. The skilled-based decision to plan for the necessary microburst escape maneuver was not well established on the part of the captain. This led to what proved to be a creative, but flawed plan. The theory, as refined during the accident analysis, suggested that the decision-making process was grounded in narrative synthesis—a process hypothesized to underlay much of decision-making and action selection both at the conscious and automatic levels.

Keywords

Human performance modeling, decision-making, human error.

INTRODUCTION

"On July 2, 1994, about 18:43 eastern daylight time, a Douglas DC-9-31, N954VJ, operated by USAir, Inc. as flight 1016, collided with trees and a private residence near the Charlotte/Douglas International Airport, Charlotte, North Carolina, shortly after the flightcrew executed a missed approach from the instrument landing system approach to runway 18R (NTSB, 1995)." The aircraft had entered a weather cell at the threshold of the runway just at the onset of a microburst. The tactic for coping with the potential microburst that the captain had planned and communicated to the first officer, as the pilot flying, was not sufficient to address the situation at hand.

Prior to selecting the USAir 1016 accident for further investigation, we had reviewed a series of accidents, culled from a set identified by Orasanu as involving "plan continuation errors" (Orasanu, Martin, & Davison, 1998)—accidents preceded by a series of cues that suggested another course of action, but for which the aircrew held fast to its original plan. And in fact, the NTSB (1995) report focused on the captain's decision to continue the approach and landing in the face of a series of cues that strongly suggested aborting the landing at an earlier point in the approach.

The USAir 1016 captain was attending and acting on the cues presenting themselves. As the aircraft was on the final approach to runway 18R, in response to the weather cell's threat, the captain briefed the tactic that he wanted executed should they have to abort the landing: "if we have to bail out … it looks like we bail out to the right (NTSB, 1995)." There was a tacit agreement between the captain and the first officer on the suitability of the plan. When the situation deteriorated to the point where the captain called for its execution, the first officers executed the plan as arranged. Several seconds later, it became clear that the plan was not adequate. The captain's new orders and the first officer's execution of them were not successful in attempting to escape the microburst.

The captain's briefing drew little attention in the accident report, yet our review of the accident data suggested that further investigation would prove fruitful. We found the briefing to be central to a better understanding of the causes of the accident. In particular, we wanted to better understand how the captain formulated the plan to gain further insight into how error can intrude on otherwise expert performance.

METHODOLOGY

Simulation provides a safe and relatively inexpensive means to rigorously explore multiple dimensions of complex problems. By including human performance models in a simulation, we can provide the means for the detailed examination of an aircrew's decision-making processes. Exploring alternate paths through a scenario add to the list of questions asked and can require revised answers to previously considered questions. To obtain model behaviors that play out in a human-like manner, satisfactory answers must be found for the detailed questions posed at each step in the modeling process. The requirement to provide answers that lead to realistic model behaviors means that hard questions cannot easily be glossed over or set aside.

The modeling process can help us to better understand the decision-making process and the sources of human error. It provides a framework in which to examine the means to prevent the occurrence of error and explore ways to effectively trap and correct errors when they are not prevented. The modeling aspect of the effort has made contributions on two basic levels. First, it has contributed to the rigorousness of the study process. For the model to faithfully represent individual human performance, it must faithfully capture *why* actions are taken and *how* they are carried out. Much of the time, the why and how questions are answered and progress is made in our understanding and in producing the model's behaviors. Some of the time, the questions do not have ready answers and become placeholders for work to be done. The current theory provides the framework for understanding how errors, such as those described in the NTSB accident reports come about. The difficult questions drive the refinement of the theory leading to more robust models that improve our understanding of human behaviors.

MODEL THEORY: PROCESS AS DOMINANT

At its best, creating a theoretical foundation for human performance modeling is an iterative process. The literatures of a broad range of disciplines, cognitive, experimental, and social psychology, neuroscience and cognitive neuroscience, and anthropology to name a few, provide a wealth of *empirical* data to support theory formation. These literatures together with that of philosophy similarly provide *theoretical* foundations to support aspects of theory development and refinement. With theoretical foundations in place, model development, experiment design, and scenario trials enable the evaluation of an implementation; and then each step of the process may be revisited. There are numerous complications. To many interesting questions there are often two sides, each supported in the literature by theory and empirical data. And our basic computational tool for modeling, the symbol-based von Neumann machine, is very different than the neuron-based parallel processing human brain. In particular, the literature on decision-making and action selection is very extensive, but it seldom suggests how the findings might fit into a more complete person model.

In furthering model development, we have taken a strong stance on the issue of representation. Symbolic models often accord declarative and procedural knowledge equal status and, more often than not, provide a large cache of declarative knowledge, as if one part of the brain is database and another composed of processes that access and reason over those data. Sometimes explicit in this formulation, and sometimes implicit is the assertion that most of this reasoning is thoughtful conscious processing. An alternate view is that in much of what we do, particularly those things at which we are very good, the process is largely automatic, even "the most intellectual activities in the application environment" (Logan, 1988a; Bargh & Chartrand, 1999).

One challenge has been to provide a reasonable symbolic representation for this mix of thoughtful and automatic processes, and within that framework, provide a place for the declarative knowledge. In the D-OMAR[1] models developed in this research effort, process is represented as a network of procedures. The network is established by the goals that structure an actor's proactive behaviors and channel reactive behaviors. As such, the network represents much of the actor's long-term memory. Declarative memory is defined and modeled as attribute-values of the procedures. Just as the vision system links distinct aspects of visual perception across the occipital nodes (e.g., color and motion in vision areas V4 and MT), selective large-scale brain structure is posited to link process with declarative elements of action selection and execution thereby creating the context in which declarative memory elements are grounded.

Hence, scripts or schemata are long-term memory elements of the model. They are represented in the model as goals, sub-goals, and procedures—the things that people know how to do (Glenberg, 1997). Goals express the purposes or intent of the actions while procedures are the templates for executing the steps to accomplish the goals and sub-goals. The completion of a task is typically accomplished through the coordinated execution of a network of goals and sub-goals that bring together the mix of human perceptual, cognitive, and motor capabilities needed for the task. As a goal and its procedures play out, impinging events lead to decisions that determine the particular execution path through the sub-goals and procedures as procedure attributes acquire particular values. The executed form of the goals and procedures are instances particularized from a script or schema.

In building aircrew models (Deutsch & Pew, 2002; Deutsch, 1998; Deutsch & Adams, 1995), much of the work lies in capturing the procedures that an aircrew executes—domain-specific knowledge in procedural form. The domain-specific procedures are supported by a layer of "basic person" procedures that support the skilled tasks of operating flight deck controls, managing the more complex tasks of handling flight deck conversations and ATC communications, and the interleaving of multiple tasks as situations demand. Existing aircrew procedures in combination with air traffic controller procedures were augmented to build the nominal models for an approach and landing scenario. The task networks, representing skilled aircrew behaviors, played out in appropriate ways depending on the particular situation presented.

[1] Information on the Distributed Operator Model Architecture (D-OMAR) is available at http://omar.bbn.com.

The task at hand was then to probe the seams of the expert behaviors to explain how aircrew errors emerged and combined to lead to the accident described in the NTSB report.

THE MICROBURST THREAT TO AN AIRCRAFT

"A microburst is defined as a 'precipitation induced downdraft, which in turn, produces an outflow when the downdraft reaches the earth's surface' …" (NTSB, 1995). An aircraft entering the microburst experiences a rapid increase in airspeed. On transiting the core of the microburst, the aircraft is then chasing the outflow and hence, experiences a rapid decrease in airspeed. When late in the approach and landing process, the rapid decrease in air speed presents a significant threat. A microburst, a rapid thunderstorm-induced change in wind speed is one form of windshear.

INTERPRETING THE CAPTAIN'S GO-AROUND PLAN

In aircrew decision-making, as in decision-making in general, the decision process may be simple or complex; the decisions themselves may be automatic or require much thoughtful effort. A decision can be complex because there are many interacting factors that go into the decision—numerous cues to be evaluated, further complicated by cue ambiguity (Orasanu & Davison, 1998). In time-pressured situations where an incorrect decision can lead to an incident or accident, risk may be difficult to accurately quantify. There may be conflicts among immediate goals and the consequences of particular actions may be difficult to accurately assess. At the same time, many, even most aircrew decisions are highly practiced—the product of highly structured and often repeated procedures in familiar situations. Extensive experience in often-repeated situations can provide immediate solutions to complex problems with little or no need for conscious, thoughtful processing. And yet, experts with considerable experience can be put in situations where their experience is more limited—situations in which they are not expert decision makers.

The captain's plan, "if we have to bail out … it looks like we bail out to the right," was the most difficult to explain statement in the USAir 1016 cockpit voice recording (CVR; NTBS, 1995). Although there had been no prior mention of a windshear threat, the weather cell at the threshold of the runway was clearly visible through much of the approach. Why didn't the captain explicitly brief a microburst escape procedure? Eighteen seconds after stating his plan, the captain noted the "chance of shear," yet he did not revise his plan. Moreover, why did he stick with his plan when they actually encountered increased airspeed, a definitive cue pointing to a windshear encounter? That turbulence and airspeed were actively being tracked strongly suggests that he was concerned with the windshear cues that had been important elements in his USAir training experience. It was at this point in the model-based exploration of the accident scenario, that we turned to theory-based analysis as a means to understand observed behaviors.

Automaticity in Expert Decision-making

The theory that grounds the model places a strong emphasis on the role of automaticity in expertise. Well-practiced skills are readily and transparently initiated (Logan, 1988b; Bargh & Chartrand, 1999), even when a complex combination of cues must be recognized to identify the correct course of action. The catch here is that most pilots have had limited experience in encounters with a microburst. The USAir 1016 captain and first officer had simulator experience in identifying and escaping a microburst situation, but we do not know from the NTSB report whether or not they had ever encountered a microburst before the July 1994 approach to Charlotte 18R. Given their limited simulator experience, the USAir 1016 pilots were almost certainly not exercising a well-practiced skill. The plan for the maneuver should the landing need to be aborted was most likely not grounded in highly practiced experience. Yet, the first fallback in formulating a plan was his training experience: Why didn't he simply select and brief for a microburst escape maneuver?

The automaticity-action link in the model clearly does not tell the whole story. We are concerned with the difficult juncture at which the highly skilled operator is making decisions at the edges of his or her well-practiced expertise. The problem at hand is one for which he or she has had limited experience. This particular instance has a further complicating twist. Even though the captain may have had limited experience in actual windshear encounters, there seemed to be an obvious plan to call out and then execute when required—the microburst escape maneuver. With the data at hand, we believe that working from our modeling perspective it is possible to suggest an explanation for the captain not selecting the apparently obvious plan. The modeling perspective also provides a suggestion on how the alternate plan might have been constructed.

Characterizing Thoughtful Processes in Decision-making

We have suggested that the authoring of the plan was not the product of an automatized process—the captain is presumed to have had limited experience with windshear encounters. Hence, it was the product of a thoughtful process. There are two prominent theories (Sloman, 1996) each claiming to be the foundation for such thoughtful processes. The first is that we as people ground our thought in *rules*: "if there is a windshear threat, then brief a microburst escape maneuver." The second is that we maintain and reason over *mental models* to create our plans such as the one briefed by the captain of USAir 1016.

Rules such as the one just quoted do little to suggest how the captain's alternate plan came about. They explain nominal behavior but do not help us to understand off-nominal behaviors. The *theory* on which the model is based traces the foundations for thoughtful work to another source. The ideas that we will briefly outline are more closely aligned with mental models, but they replace the, for us, very loose notion of a mental model with a more explicit entity for which we have an alternate label. Our explanation avoids the explicit split between the representation that a mental model implies and the separate process of reasoning over that representation.

Rethinking Memory

Memory or more specifically, a new perspective on what memory is fundamentally about is the starting point for the theory. The idea is a simple one: Memory's primary function is to retain key elements of today's experience so that we will be able use that experience to perform more effectively in the future (Glenberg, 1997). Rather than the prominence given to declarative knowledge, what we know, the focus is on *what we know how to do*—episodic memory, not in the form of the remembrance of what we did, but rather the remembrance of what we know how to do—that is, procedural memory. And indeed, most of the memory in our models is in the form of goals and procedures—what the model knows how to do.

In the modeled Charlotte scenario, one of the low-level background activities that both the captain and the first officer execute is a scan that includes watching for weather cells. The sighting of the weather cell (the cell at the threshold of the runway is the only one currently modeled) first kicks off a discussion between the crewmembers about the cell and leads to the tracking of the several cues needed to assess its impact as they proceed on the approach. The focus is on the steps the crew *executes* (drawn from long-term memory) in response to the cues and the supporting knowledge needed to select and execute those actions.

However, even the mundane things that we do each day have slight twists and variations on how we did them yesterday. As we make our way through each day, we are always making the adjustments and refinements necessary to address the particular immediate situation. As we conduct these activities, we are making slight additions, adjustments, and refinements in our *remembrance of what we know how to do*. The conscious representation of the goals and procedures that we execute as we go about our daily business are *narratives* that we employ to guide our actions. This introspective view of what we know how to do in the form of narratives is constantly subject to real world experience. That experience serves as the critique helping to nurture the refinements that improve our skill base.

Memory's essential function is the refinement of the representation of what we know how to do that helps us better guide future actions—to act more effectively in the world today than we were able to yesterday based on our experience (Glenberg, 1997). We claim that this is memory's primary function and most important accomplishment. Two of memory's more familiar functional capabilities devolved from this starting point. The first of these, captured by the conventional view of memory, is simply remembrance of what happened in the past. The second further capability is the ability to use a remembrance as a building block to explore what I might do in the future and gauge how it might come out. The ability to introspect on what we know how to do, in the form of narratives, may be used to reflect on and refine the actions that we then take in a given situation. Remembrance may lead to immediate action selection (automaticity) or it may be used as the basis for thoughtful deliberation about future actions.

In building our models based on this perspective on memory, we have redefined the balance between declarative and procedural memory. Virtually all human performance modeling frameworks grant equal prominence to distinctly separate declarative and procedural memory resources (e.g., Anderson & Lebiere, 1998; Corker & Smith, 1993; Laird, Newell, & Rosenbloom, 1987; Meyer & Kieras, 1997). Declarative memory is selectively retrieved and used as the basis to select and retrieve procedures that then execute more or less successfully. In our models, procedural memory, what we know how to do, has precedence. Declarative memory exists as attributes of the procedures—memory for a particular entity is tightly tied to the context in which it is used. There is no declarative memory box per se. Declarative memory items do not exist context free. With this as background, we can now look again at the process leading to the captain's briefing for addressing the microburst threat.

Understanding the Briefing: Narratives and their Role in Action Selection

Goals and procedures capture what the aircrew and controllers know how to do and successfully modeled most of the events in the USAir 1016 approach to Charlotte 18R, but they did not capture the plan that the captain briefed on the approach. In the model, the captain having viewed the weather cell on the approach path briefed a microburst escape maneuver. Given the USAir training, this is just what he would have been expected to do. Here we were at a point where the theory underlying the model itself suggested a path to an understanding of *why* the briefing actually given was constructed with its unexpected content. The theory also offered a potential explanation for *how* the plan might have been constructed.

The problem that we as modelers were confronted with here usually comes in a slightly different form: A working model exists with capabilities to successfully address a range of situations, but within one of those situations, the cues are in some way slightly different in a manner such that the behavior the model produces fails to be human-like in some sense. As humans in the same situation, we might readily detect the variance in the cues and stitch together a slightly adjusted response to meet the requirements of the particular situation. To date, for most of these situations, we simply do not know how to built models that achieve this level of human-like capability.

The problem here was a variation on adapting to new cues in a familiar situation. Here the cues were consistent, at least with past training experiences, but the captain appeared to have felt the need to amend the trained response. The resulting process was the same in each case. A varied response to the situation was determined to be required. The potential sources for that response were just the set of *narratives* that described responses to closely related situations. In such circumstances, a narrative may be selected and slightly amended or two or more narratives may contribute segments that may be stitched together in a new way to provide the response. This is a process leading to what Orasanu and Fischer (1997; Orasanu & Martin, 1998) have termed a *creative* solution to a problem at hand.

For the particular USAir 1016 captain's briefing, potential contributing narratives included: (1) a standard microburst escape maneuver; (2) a standard go-around procedure; and (3) an en route like avoidance of the weather cell by maneuvering to avoid it. The situation the captain appeared to be planning for was the possibility that they would act to avoid the weather cell; he appeared to plan to proceed with a standard go-around, but literally go around the cell to the right (just as in the case of the en route cells he had circumvented earlier in the approach). Pieces from each of the three related narratives may be readily stitched together to produce: "if we have to bail out … it looks like we bail out to the right." They will execute a standard go-around; they will avoid the cell by going around to the right as they had successfully done previously during their approach; and they will monitor the windshear cues as taught for the microburst escape procedure (this last element of the plan was later acted on, but not explicitly briefed). The theory that underlies the model (but that is not yet instantiated in the model) suggested the process by which the plan was generated.

But what about *why* the plan was so generated. Why was a creative solution invented? Might not there have simply been a rule-like automatic behavior: When there is a windshear threat, then brief a microburst escape maneuver. First, it is likely that the captain did not have sufficient experience in microburst situations to lead to a strong automatic link from the cues at hand to the generation and briefing of the expected microburst escape plan. Thus the cues in combination with the related available narrative elements for a new plan may well have been sufficient cause to lead to a *thoughtful* intervention that derailed the generation of the lightly connected automatic response. Not only did the processing of the cues help determine the plan and the resulting actions, recent related decisions and actions influenced the processing of current cues and the selection of the plan for future actions (Holyoak & Simon, 1999). Successful decisions and actions taken with respect to circumventing the earlier weather cells on the approach path suggest themselves as important influences on the thoughtful planning process for avoiding the weather cell at the threshold to Charlotte runway 18R. Recent narrative outcomes (the ready avoidance of the weather cells earlier in the approach) played a role in establishing the briefing decision as a thoughtful rather than automatic process, and they came into play in the decision-making itself. The briefed plan was accepted by the first officer (an interesting issue in its own right), called for three minutes later as they encountered the microburst, and executed by the first-officer as pilot-flying. It was almost four seconds after initiating the planned maneuver before the plan was explicitly revised by the captain as he called for "max power."

DISCUSSION: ADAPTIVE BEHAVIOR THROUGH NARRATIVE SYNTHESIS

Orasanu and Martin (1998) described creative action selection as the response to "situations in which no suitable options are readily available and the decision maker must invent at least one to meet the demands of the situation." In the Charlotte windshear threat situation, there was a straightforward plan for a microburst escape to be put in place, yet a creative action selection process was employed. Rather than a rule-based decision, we have suggested that the straightforward decision would have been a skill-based decision. The skilled-based decision was not made because the captain lacked sufficient experience in windshear situations to firmly establish the skill—previous real world experience was not available and simulator training was inadequate to establish the skill.

Probing the captain's creative action selection shed new light on the process itself that has, in turn, furthered the theory underlying our modeling efforts. The creative action selection process is viewed as central to adaptive behaviors. It is something that, we as people, do many times a day as we extend what we know how to do by adjusting experience-based action selection to meet the demands of the immediate situation. Our new experience is then consolidated as an extension to what we know how to do. Our hypothesis is that our experience base—what we know how to do—can then be accessed as narratives derived from the base level representation. Access to that experience cache as narratives, both at the conscious and non-conscious levels, then makes possible adaptive behaviors through narrative synthesis—a process that draws upon one or more narratives and adjusts and merges them to meet the needs of the immediate situation.

ACKNOWLEDGMENTS

The author wishes to acknowledge Judith Orasanu for introducing him to the relevant literature and the guidance that she provided once the research effort was underway. Thanks are due as well to Richard Pew for his contributions to this research effort.

REFERENCES

1. Anderson, J. R., & Lebiere, C. J. (1998). *The atomic components of thought.* Mahwah, NJ: Lawrence Erlbaum Associates.
2. Bargh, J. A, & Chartrand, T. L. (1999). The unbearable automaticity of being. *American Psychologist, 54,* 462-479.
3. Corker, K. M., & Smith, B. R. (1993). An architecture and model for cognitive engineering simulation analysis: Application to advanced aviation automation. In *Proceedings of the AIAA Computing in Aerospace 9 Conference.* San Diego, CA.
4. Deutsch, S. E., & Pew, R. W. (2002). Modeling human error in a real-world teamwork environment. In *Proceedings of the Twentieth-fourth Annual Meeting of the Cognitive Science Society (pp. 274-279).* Fairfax, VA.
5. Deutsch, S. E. (1998). Interdisciplinary foundations for multiple-task human performance modeling in OMAR. In *Proceedings of the Twentieth Annual Meeting of the Cognitive Science Society.* Madison, WI.
6. Deutsch, S. E., & Adams, M. J. (1995). The operator model architecture and its psychological framework. In *Proceedings of the 6th IFAC Symposium on Man-Machine Systems.* Cambridge, MA.
7. Edelman, G. M. (1987). *Neural Darwinism: The theory of neuronal group selection.* New York: Basic Books.
8. Edelman, G. M. (1989). *The remembered present: A biological theory of consciousness.* New York: Basic Books.
9. Glenberg, A. M. (1997). What memory is for. *Behavioral and Brain Sciences, 20,* 1-55.
10. Holyoak, K. J., & Simon, D. (1999). Bidirectional reasoning in decision making by constraint satisfaction. *Journal of Experimental Psychology: General, 128,* 3-31.
11. Laird, J. E., Newell, A., & Rosenbloom, P. S. (1987). SOAR: An architecture for general intelligence. *Artificial Intelligence, 33,* 1-64.
12. Logan, G. D. (1988a). Automaticity, resources, and memory: Theoretical controversies and practical implications. *Human Factors, 30,* 583-598.
13. Logan, G. D. (1988b). Toward an instance theory of automatization. *Psychological Review*, *95*, 492-527.
14. Meyer, D. E., & Kieras. D. E. (1997). A computational theory of executive cognitive processes and multiple-task performance: Part 1. Basic mechanisms. *Psychological Review, 104,* 3-65.
15. National Transportation Safety Board. (1995). *Aircraft accident report: Flight into Terrain during Missed Approach, USAir Flight 1016, DC-9-31, N954VJ, Charlotte International Airport, Charlotte, North Carolina, July 2, 1994.* NTSB/AAR-95/03. Washington, D. C.: U. S. Government Printing Office.
16. Orasanu, J., & Davison, J. (1998). Errors in aviation decision making: A factor in accidents and incidents. In *Proceedings of the Working Conference on Human Error, Safety, and System Development,* (pp. 100-106). Seattle, WA.
17. Orasanu, J, & Fischer, U. (1997). Finding decisions in natural environments: The view from the cockpit. In C. E. Zsambok & G. Klein (Eds.), *Naturalistic decision making* (pp. 343-357). Hillsdale, NJ: Lawrence Erlbaum Associates.
18. Orasanu, J., & Martin, L. (1998). Errors in aviation decision making: A factor in accidents and incidents. In *Proceedings of the Workshop on Human Error, Safety, and Systems Development,* Seattle, WA.
19. Orasanu, J., Martin, L., & Davison, J. (1998). Errors in aviation decision making: Bad decisions or bad luck? In *Proceedings of the Fourth Conference on Naturalistic Decision Making.* Warrenton, VA.
20. Sloman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin, 119,* 3-22.