

Conceptual Blending and Airplane Navigation

Barbara Holder and Seana Coulson

The Boeing Company¹
P.O. Box 3707 MC 67-TC
Seattle, WA 98124-2207
barbara.e.holder@boeing.com

University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92126-0515
coulson@cogsci.ucsd.edu

Abstract

Fauconnier and Turner's conceptual blending theory (1998) has been used to analyze many types of cognitive phenomena, particularly the role blending plays in language comprehension. In this paper we illustrate how blending also has a role in mediating expert performance, such as understanding the meaning of a chart or a navigation display. Pilots use blends to assess flight situations, predict what the airplane will do next, and make decisions. These blended conceptualizations of navigational tools help the pilot to interact appropriately with the aircraft. We propose that the principles of conceptual blending theory aid in the design and testing of novel display designs.

Introduction

Advances in display technology have made it possible for designers of modern flight decks to create new representations for conventional flight tasks. Pilots like these displays because they combine numerous instruments and dials into a single integrated format. One such display is the navigation display (ND). The navigation display presents a dynamic moving map of the airplane's planned flight route. Interpreting the navigation display involves *conceptual blending*, processes for the creative combination of information from different domains. Below we introduce the theory of conceptual blending and show how blending processes play an important role in mediating expert performance in airplane navigation.

Conceptual Blending

Conceptual blending is a theoretical framework for exploring human information integration (Coulson, 2000; Fauconnier & Turner, 2002). In this framework, a small set of processes operate in the construction of meaning in a variety of cognitive phenomena, including language comprehension, concept combination, and analogical reasoning. In conceptual blending, elements from two or more

input models are projected into a *blended space* in order to construct a hybrid cognitive model, frequently referred to as the *blend*. Besides inheriting partial structure from each of the inputs, the blend can develop emergent structure of its own via the activation of background knowledge.

For example, Fauconnier & Turner (1998; 2002) analyze a linguistic example that involves blending information about a modern catamaran and a clipper ship from the 19th century in order to construe the two as racing one another. A sailing magazine reported the event this way: "As we went to press, Rich Wilson and Bill Biewenga were barely maintaining a 4.5 day lead over the ghost of the clipper *Northern Light*, whose record run from San Francisco to Boston they're trying to beat. In 1853 the clipper made the passage in 76 days 8 hours."

In this example, one input involves a cognitive model of the event that occurred in 1853, in which a Clipper ship sailed from San Francisco to Boston in 76 days and 8 hours. The other input involves a cognitive model of the event that occurred in 1993, in which a catamaran sailed from San Francisco to Boston. The blended space involves a cognitive model of an event that exists only in the imagination of the modern-day sailors, in which the catamaran is involved in a race with the "ghost" of the clipper ship (see Figure 1).

The hybrid cognitive model, or the *blend*, represents a race that never actually happened, as the two boats never sailed the same course at the same time. However, constructing this blend allows speakers to talk of the modern-day catamaran as being either "ahead" of the clipper ship, or "behind" it – by comparing where each boat was on the first, second, or third (and so on) day of their respective voyages. Conceptual blending theory is meant to capture the extent to which people reading the above passage about the *Northern Light* understand that each boat sailed the course at a different time, but are able to think and talk about the catamaran's progress on the course as if it were engaged in a race with the other boat.

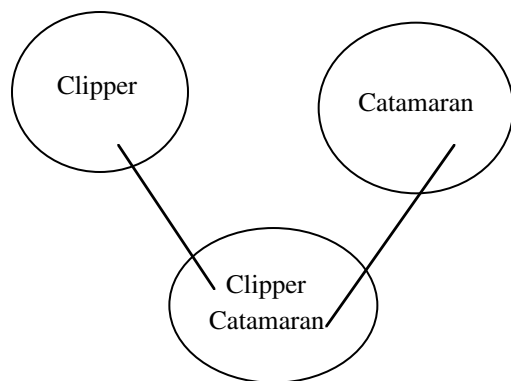
Moreover, once the reader (and the sailors) adopt the race model in the blend, it is possible to react differently to progress along the course from Boston to San Francisco.

¹ This research was conducted while the author was a postdoctoral researcher at the University of California, San Diego

For instance, knowing that one is at a certain position along the course is one thing, but being 4.5 days ahead of the *Northern Light* is cause for excitement. Similarly, “falling behind” the *Northern Light* is distressing when the voyage is construed as a race, and perhaps motivates the sailors to work harder to “catch up”.

In conceptual blending, selected aspects of each model are mapped into the blended space. In this example, the clipper ship from 1853 and the catamaran from 1993 are each mapped into the blended space, and each have a distinct counterpart in the blend. In contrast, the route from Boston to San Francisco is mapped from each input onto a single route in the blend – presumably because it is the same route. Emergent structure arises because two boats sailing along the same course activates the “boat race” frame from long-term memory.

Blending processes depend centrally on projection mapping and dynamic simulation to develop emergent structure, and to promote novel conceptualizations, involving the generation of inferences, emotional reactions, and rhetorical force. There are many other examples of blending in language (e.g. Fauconnier & Turner, 1998; Coulson, 2000), but little research has investigated the role of blending in the performance of situated cognitive activities (though see Maglio & Matlock, 1996 for one exception).



1993 Boat Race
Figure 1. Boat Race Blend

Blending and Navigation

In his book, *Cognition In The Wild*, Ed Hutchins details the conceptual shifts navigators make while determining the position and path of their vessel (1995). He argues that a navigator’s ability to conceptualize the movement and position of the ship in terms of a two-dimensional chart requires experience in mapping between the chart and the local environment. In fact, using a chart involves an imaginative leap in which the navigator understands himself and his immediate surrounds from the perspective of the chart. Hutchins describes the conceptual demands of navigation as follows:

When a Western navigator takes a bearing of a landmark, he has a real point of view on a real space.

However as soon as he leans over his chart, he is no longer conceptually on the boat; he is over the sea surface, looking down on the position of his craft in a representation of the real local space. Novice navigators sometimes find this change of viewpoint disorienting, especially if the orientation of the chart does not happen to correspond to the orientation of objects in the world. (Hutchins, 1995:79-80).

Navigation thus involves a blend between the navigator’s conceptualization of the chart, and his conceptualization of the physical world represented in the chart. One input is the chart, a piece of paper whose markings are in systematic correspondence with the physical world. The chart represents the spatial relationships between elements in the seascape from an aerial perspective. The other input is the navigator’s perspective of the seascape from the bow of the boat. In order to exploit the correspondences between the chart and the seascape, the navigator must construct a hybrid model that integrates his conceptualization of the chart with his conceptualization of the position of the ship. The blend involves imagining a huge ship moving along a flat piece of paper, and that the navigator can view his ship’s motion from a bird’s eye perspective.

The ability to construct and elaborate this blend is one dimension of a navigator’s expertise. Hutchins’ note that novices find it disorienting to use the chart highlights the fact that blending results in a construct that is more than an understanding of the correspondences between two domains. The novice understands that there is a systematic correspondence between the physical world and the representation on the chart, and yet cannot always exploit those correspondences inferentially. The ability to use correspondences to navigate is mediated by the hybrid cognitive model in the blended space.

Modern Airliner Navigation Display

The aeronautical environment presents many of the same conceptual issues navigators face in the nautical world only the representations are different. Modern jetliners are highly automated and present integrated flight information to pilots on redundant sets of display screens. Although pilots of modern jetliners still use charts to navigate, they also use an interactive navigation display whose properties differ from those of a paper chart. The display technology supports the integrated depiction of relationships that are not possible with the older technology as it presents a picture that pulls many abstract flight concepts and physical objects into a single two-dimensional image.

Figure 2 depicts a modern navigation display used in the Airbus A320 aircraft. The airplane symbol corresponds to the position of the actual airplane being flown. The display also presents waypoints, the landmarks pilots use to track their progress on their course, and restrictions they must meet. Waypoints are represented as diamonds, and the green line connecting the airplane to the next waypoint corresponds to the path the airplane is currently flying.

Figure 2 represents a case in which there is an altitude restriction of 13,000 feet at the MAJIC waypoint.

The numbered arc at the top of the display is a compass rose that is always oriented with the current heading at the top. Because wind can cause the aircraft's heading (direction the nose is pointing) to differ from its track (direction airplane is flying over the ground), the heading and track are each represented with geometric figures superimposed on the compass (a triangle and a diamond, respectively). Finally, the wind is represented via an arrow pointing in the direction the wind is coming at the airplane, as well as being given in degrees and speed (in Figure 2 this is indicated by the 210/36).

Pilots like this display because it makes it easy for them to conceptualize their flight path at the present time and to integrate that understanding with where the plane will be in the near future. The navigation display in Figure 2 represents a great advance in navigation technology as there is no single display in conventional airplanes that presents heading, track, wind speed, wind direction, projected course, waypoints, and altitude restrictions. Formerly it was necessary to mentally integrate data from a chart and several instruments in order to arrive at the same information.



Figure 2. A Modern Navigation Display.

Interpreting the Navigation Display

Because we are able to interpret the meaning of symbolic and iconic representations so easily, we are prone to forget that, at some level, the symbols on the navigation display are just pixels on a screen. In fact, interpretation of the navigation display involves many different kinds of conceptual blending, and recruits concepts from at least four inputs: navigation, calculations, display, and airplane (Table 1). Elements from each of the input spaces are projected into the blended space. Pilots interpret the display as integrated information from each space that requires experience to develop. When a non-pilot views at this

display he sees a set of meaningless symbols. However, when a pilot views the display he sees the aqua triangle as the direction the nose of the airplane is pointing.

Navigating an aircraft means pilots follow a three-dimensional flight path using landmarks or waypoints to track their progress on the course. They have restrictions that require them to be at specified altitudes or speeds at specific points on the path. Pilots know that if they do not meet altitude restrictions and other federal regulations there may be serious consequences. Like ship navigators, pilots have to develop expertise with the coordinate system employed in the navigation display, and must also perform some of the same imaginative leaps that ship navigators do to effectively use their navigation tools.

Table 1. Inputs to the Navigation Display Blend

| Input | Structure |
|--------------|---|
| Navigation | Charts Instrument Readings Waypoints Weather Predictions |
| Calculations | Estimated time of arrival Vertical rate Ground speed Distance to waypoint Wind correction |
| Airplane | Current position and path Future position and path Heading and track |
| Display | Shapes Lines Colors Letters Numbers Symbols |

As mentioned earlier, the navigation display involves a number of different blends. For instance it integrates the representation of lateral points along the course with vertical points along the course. It also provides an integrated representation of current aircraft information and predicted aircraft information. The green line represents the future path the airplane is expected to fly. The blue arrow is the point where the airplane will level at the altitude specified on the glare shield. The blue arrow represents a point that is simultaneously positioned in vertical and lateral space. The magenta circle indicates a constraint—this is an altitude constraint at 13,000 feet. The green diamond with the INTRCPT label is a specific kind of waypoint, it indicates the point where the airplane will meet the computed lateral flight path. The jagged arrow is the point where the airplane will intercept the computed vertical flight path.

Once pilots learn these mappings, they see the representations on the display as an integrated blend of the aircraft's current position and the airplane's predicted performance in future space and time. When experienced pilots look at the display they interpret it in 3-D (Figure 3.). Furthermore they use these correspondences to make inferences about the craft's predicted course and to make decisions on how to manipulate the airplane to maintain its course and that is the role of conceptual blending in navigating with this tool. In order to study the role of blending in navigation, we interviewed pilots working for a major U.S. airline and asked them how they use the navigation display when they fly.

The data we present below was collected as part of a larger study to investigate the cognitive consequences of aircraft automation. In particular, this study addressed how pilots' expertise develops and changes over time, as well as the conceptual models pilots build as they interact with the automation in their planes. As part of this study, we interviewed pilots entering training for the Airbus A320 program at a major U.S. carrier training center. Pilots were interviewed again within a few months of completing training, as well as at six, twelve and eighteen months of flying in revenue service. The interviews were structured to probe how the pilots used automation and to identify any problems they may have had using it. In addition, pilots were observed in flight from the jumpseat in order to record the way in which they actually used the automation.




Figure 3. Navigation display from the perspective of the pilot.

Prediction

We found that many pilots use display correspondences to predict where the aircraft will meet the flight path. Here is an example:

Pilot: I monitor the uh // you know how we have those little lightning bars, and little-the little uh like arrows and things that // they appear on the screen to tell you if you're on profile, or that's where you're gonna intersect profile,

Interviewer: right

Pilot: (and I think)-well I take those and I use *that* information // to uh determine if I'm gonna meet the uh // meet the uh restriction.

The pilot reports he uses the "lightning bars" to determine if his airplane is on the computed vertical flight path. If he is not on the planned flight path, or what he calls *profile*, the lightning bar shows him where he will intercept

the vertical flight path. He uses the arrows to determine where the aircraft will level, and uses that to determine if the airplane will meet the altitude constraint. He said the arrows and lightning bars *tell* him things, but of course we know that at a certain level they are really only pixels on the screen. After dozens of hours of experience using the navigation display in the air, this pilot brings to the interaction a sense of communicating with the representations on the screen to know what to expect. The facility of going from the basic display mappings to reasoning about the aircraft is a function of experience and he's using the same blend ship navigators use, imagining himself in the plane in the display, while simultaneously doing things that affect the performance of the actual plane he's flying.

One function of the display is to allow pilots to assess the flight situation and determine whether or not human intervention is required. When a pilot monitors the display he might see changes in the behavior of the representations that indicate changes in the predictions. Thus when the level off point is before the altitude constraint everything is okay, because it indicates that the plane will meet the altitude restriction before it reaches the relevant waypoint. However, weather conditions may alter the plane's original course so as to require alterations to the automated flight plan. For example, one consequence of a strong tailwind might be that the arrow moves beyond the crossing restriction indicating that the plane will not meet the altitude constraint until after the specified waypoint. In such a case, the pilot must intervene to make the airplane level at or before the constraint. Part of understanding the display is that it triggers the appropriate action from the pilot and the pilot understanding the serious consequences if the restrictions are not met.

Decision-making

Another pilot we interviewed described how he uses display properties to make decisions.

Pilot: And uh what I do in a case where they (meaning air traffic control) give you a crossing restriction, I use that blue hockey stick and uh just roll the vertical speed to wherever you think it should be and see where the blue hockey stick ends up, if that looks right on the map display on the nav display away you go.

The blue hockey stick indicates where the airplane is predicted to be level, so if that point is beyond the altitude constraint then the pilot has to do something to intervene so that the airplane is at the right level at the right time. The pilot says he dials the vertical speed knob and watches the hockey stick move up or down the flight path in correspondence with his dialing. Vertical speed is the rate the airplane will climb or descend in feet per minute (fpm). A higher vertical speed such as 1000fpm will bring an airplane that is above the vertical path to it sooner than a lower vertical speed of 200fpm.

He coordinates turns on the dial with the position of the arrow in relation to the constraint. From the transcript, it is

clear that the pilot conceptualizes dialing the vertical speed knob as directly connected to the movement of the hockey stick on the map. But in fact these two parameters are only indirectly connected. The flight intercept point depends upon a number of variables, including wind speed, the speed of the aircraft, and the distance to the constraint. The pilot's manipulation of the vertical speed knob is merely the modulation of a single variable.

However, because the computation of the flight intercept point is distributed between the pilot and the auto-flight system, the pilot's innovative use of the navigation display is highly effective. As the pilot adjusts the vertical speed, the auto-flight system re-calculates the plane's projected flight path and displays the results on the navigation display. The dynamic properties of the display thus promote the pilot's subjective experience that the movements of the dial have direct consequences for the movement of the hockey stick. In this case, the pilot constructs a blended model in which the movement of the dial causes the movement of the blue hockey stick. Moreover, in the context of an effective auto-flight system, this conceptual blend enables the pilot to experience the movements of the dial, the movements of the hockey stick, and compliance with the altitude constraint as a single integrated activity.

In this example the display also becomes a computational resource for determining the rate of descent needed to meet the constraint. The pilot's manipulation of the arrow via the vertical speed knob is an active elaboration of the blend. This movement is an example of what Kirsh and Maglio (1994) call an epistemic action—playing with the vertical speed and watching the arrow gives him an idea of the variance in vertical speed that will meet the constraint and performs that computation for the pilot.

Although not presented in training, the hockey stick heuristic was described by a majority of pilots in our study. Apparently, pilots have discovered this handy heuristic for determining the right vertical rate to meet the altitude restriction by interacting with the display in a real flight context. In training, pilots are presented with the mappings between the display and the relevant navigational concepts. The pilots' independent discovery of the hockey stick heuristic makes it clear that in the hybrid cognitive model in the blended space affords an integrated understanding whose computational properties differ from an appreciation of the mappings in the display. One could study the mappings indefinitely and never arrive at this heuristic.

Action

The auto-flight system is a remarkable piece of technology, but it is not a replacement for pilot skill and judgment. The navigation display occasionally misrepresents aspects of the aircraft's state, especially the parts that pertain to its predicted behavior. Before an airplane takes off pilots program their route into the flight management computer and include waypoints, fuel on board, altitude constraints, the predicted winds and other things. Consequently, the flight management computer will compute the constraint and descent positions based on the predicted winds that

were entered hours before at the originating airport. This can lead to "bad" decisions by the auto-flight system that must be corrected by the pilot. For example, to compensate for a tailwind, the airplane would need to begin its descent earlier than predicted.

One of the pilots we interviewed talked with us about this very issue.

Pilot: it does this perfect calculation to cross this thing, but then as you go down, it can get far behind. the wind picks up or something...it can get behind so instead of rolling out perfectly on the altitude they want you to, it now shows you're gonna be two or three miles past it or something because the wind changed.

Interviewer: oh okay.

Pilot: having uh // less headwind or more headwind or more tailwind it's the other way, and it messes its own internal calculations up, so you just have to manually reach up there, and force the airplane to do what it's supposed to be doing.

When a prediction on the display is wrong, the pilot has to recognize it and then interact with the navigation display and aircraft to reconcile the discrepancy. As the pilot says, he has to, "force the airplane to do what it is supposed to do." The ability to conceptualize the navigation display as equivalent to things in the aviation space and draw inferences and make predictions from it requires interactions between pilot, aircraft, and display. Thus the utility of the display for navigation is both the *cause* and *effect* of pilot's interaction with the airplane.

Design Implications

The data examined indicate that people's conceptual systems are extremely flexible and adaptive. One consequence of this conceptual fluidity is that interfaces will often have unanticipated computational affordances. For example, in the case discussed above, the pilots discovered that they could facilitate the task of meeting altitude constraints by observing the effect of changes to the vertical speed on the blue hockey stick on the navigation display. In this case, the new affordances of the display arise because of humans' propensity to *compress*, or simplify certain sorts of concepts.

Studied in detail by Fauconnier & Turner (2002), compression of complex relationships into simpler ones is a basic property of human cognition. Indeed, we suggest that compression is a basic function of an interface. In the navigation display, for example, the entire flight route of thousands of miles, and which unfolds over hours, can be represented in a single display that can be apprehended all at once, as it were. These sorts of representations can greatly facilitate planning tasks, but are less useful for making decisions based on the immediate context – in which a less compressed, and more specific representation is desirable.

Another sort of compression basic to interface design is the shortening of cause-effect chains. People have a tendency to simplify the representation of long causal chains from many steps into a few, or often only a single step (Fauconnier & Turner, 2002). Besides reducing the number of causal steps, cause-effect compression can be used to construe a diffuse or fuzzy case of causation as being more direct than it actually is. Typically, the user reasons by analogy and recruits a cognitive model with a simple, direct causal relationship to help understand the more complicated causal sequence in the target domain.

In many cases cause-effect compression is a desirable function of an interface since the user need not know the complicated underpinnings of his interaction with a device. However, the designer must be aware that the tendency for cause-effect compression will continue in the user's interaction with the interface itself. In the case of our pilots, this compression resulted in the discovery of useful heuristics. However, because this sort of compression can lead to unanticipated drawbacks, it is important to explore the emergence of compressions in usability tests.

In fact, these observations highlight the need for usability tests conducted under realistic conditions. Because of pilots' conceptual fluidity, it is virtually impossible to predict all of the consequences of a design. Testing under realistic conditions both reveals the unanticipated virtues of the design, and allows unanticipated drawbacks to surface. Moreover, realistic testing conditions are warranted by the observation that useful blends depend on what they are used for.

Conclusions

We have illustrated how expert performance is mediated by conceptual integration. In the examples we showed how pilots use their integrated constructs to assess, infer, and control the airplane. There seems to be a dynamic coupling between the display and a pilot's actions. The pilot is not passive he is actively responsive to the display and the display is responsive to his actions. Thus airplane performance is not entirely produced by the automation or by the pilot. It is through their interaction that the aircraft follows the correct course and meets the appropriate restrictions. One of the positive consequences of pilots' interactions with the displays is the discovery of heuristics, such as the hockey stick heuristic, that make flying easier. Conceptual

integration theory offers a valuable descriptive process for understanding expert situated performance that arises when information from different domains must be integrated.

We propose that this approach is also useful for developing and assessing display designs that will be used to support expert performance. This approach moves us beyond usability testing by providing a means of identifying emergent properties in the display that performers can exploit in the conduct of their work. The principles of blending help us identify the mappings between pilot knowledge and elements of the world *and* the novel relationships the display makes possible. When we examine the ways a pilot can interact and manipulate these relations, we can identify how the display might be used in practice and the kinds of relations that pilots might exploit to simplify their cognitive tasks. In the analysis of a display design it is important to understand not only the time it takes to perform various tasks and the errors that might arise from its use, but also in how the pilot detects and exploits emergent structure in the display in the context of real flight tasks.

Acknowledgements

We are grateful to NASA AMES Research Center for funding this research through NAG 2-1313. We are appreciative of the helpful comments we received from our reviewers.

References

- Fauconnier, G. and M. Turner. 2002. *The Way We Think: Conceptual Blending and the Mind's Hidden Complexities*. New York: Basic Books.
- Coulson, S. 2000. *Semantic Leaps: Frame-shifting and Conceptual Blending in Meaning Construction*. Cambridge: Cambridge University Press.
- Fauconnier, G. and M. Turner. 1998. Conceptual integration networks. *Cognitive Science* 22:133-187.
- Hutchins, E. 1995. *Cognition in the Wild*. Cambridge, MA: MIT Press.
- Kirsh, D. & Maglio, P. 1994. On distinguishing epistemic from pragmatic action. *Cognitive Science* 18: 513-539.
- Maglio, P. and T. Matlock. 1999. The conceptual structure of information space. In A. Munro, D. Benyon, and K. Hook (eds.), *Personal and Social Navigation of Information Space*. Berlin: Springer-Verlag.