www.computer.org/intelligent

# Making Sense of Sensemaking 2: A Macrocognitive Model

*Gary Klein, Brian Moon, and Robert R. Hoffman*

Vol. 21, No. 5
September/October 2006

# Making Sense of Sensemaking 2: A Macrocognitive Model

**Gary Klein and Brian Moon,** *Klein Associates Division of ARA*
**Robert R. Hoffman,** *Florida Institute for Human & Machine Cognition*

In our first essay on sensemaking,[1] we discussed various possible meanings of the concept and debunked some of the myths that seem current in discussions of cognitive work. The motivation for these two essays is to question whether it makes sense to envision certain kinds of intelligent sensemaking systems. None of the "verdicts" we announced in the first essay mean that intelligent technologies might not assist people in sensemaking. Indeed, intelligent technologies might help; they just won't be the sorts of technologies that people seem to seek.

Gary Klein and his colleagues have laid out a theory of sensemaking that might be useful for intelligent systems applications.[2] It's a general, empirically grounded account of sensemaking that goes significantly beyond the myths and puts forward some nonobvious, testable hypotheses about the process.

When people try to make sense of events, they begin with some perspective, viewpoint, or framework—however minimal. For now, let's use a metaphor and call this a *frame*. We can express frames in various meaningful forms, including stories, maps, organizational diagrams, or scripts, and can use them in subsequent and parallel processes. Even though frames define what count as data, they themselves actually shape the data (for example, a house fire will be perceived differently by the homeowner, the firefighters, and the arson investigators). Furthermore, frames change as we acquire data. In other words, this is a two-way street: Frames shape and define the relevant data, and data mandate that frames change in nontrivial ways.

Figure 1 shows that the basic sensemaking act is data-frame symbiosis. The figure captures a number of sensemaking activities. Sensemaking can involve elaborating the frame by adding details, and questioning the frame and doubting the explanations it provides.[3] A frame functions as a hypothesis about the connections among data. One reaction to doubt is to explain away troublesome data and preserve the frame.[4,5] These two aspects, elaborating the frame and preserving the frame, are part of the elaboration cycle of sensemaking (the left side of figure 1), akin to Jean Piaget's notion of assimilation.

Yet another sensemaking cycle is to reframe (see the figure's right side). Here, questioning the frame leads us to reconsider—to reject the initial frame and seek to replace it with a better one. We might compare alternative frames to determine which seems most accurate. Or we might simply be mystified by the events. The sensemaking activity here, akin to Piaget's notion of accommodation, is to find some sort of frame that plausibly links the events that are being explained.

Each of these aspects of sensemaking has its own dynamics, strategies, and requirements. Recognizing a frame and recognizing data are different from elaborating a frame that has already been adopted, and this is different from explaining away inconsistencies. Different still are the reactions to questioning a frame—choosing between alternative frames and constructing a frame where none exists.

The Data/Frame Theory posits a closed-loop transition sequence between

- mental model formation (which is backward looking and explanatory), and
- mental simulation (which is forward looking and anticipatory).

Think of the simplest transition sequence as a chain of closed loops. Each loop is triggered by a perceived subevent, leading to an effort to refine the existing mental model (backward looking) and an effort to run a new mental simulation (forward looking). You can construct a tran-


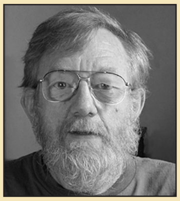
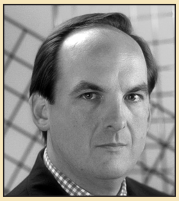

**Editors: Robert R. Hoffman, Patrick J. Hayes, and Kenneth M. Ford**
Institute for Human and Machine Cognition, University of West Florida
rhoffman@ai.uwf.edu


1541-1672/06/$20.00 © 2006 IEEE
*Published by the IEEE Computer Society*

sition sequence retrospectively to generate an explanation of how events and subevents unfolded, or prospectively to imagine how a major causal factor or a situational mix of factors might play out. For illustration, envision a transition sequence using the metaphor of billiards, where a player would anticipate how hitting one ball would lead to motion in a second, and a third, to the shot's completion.

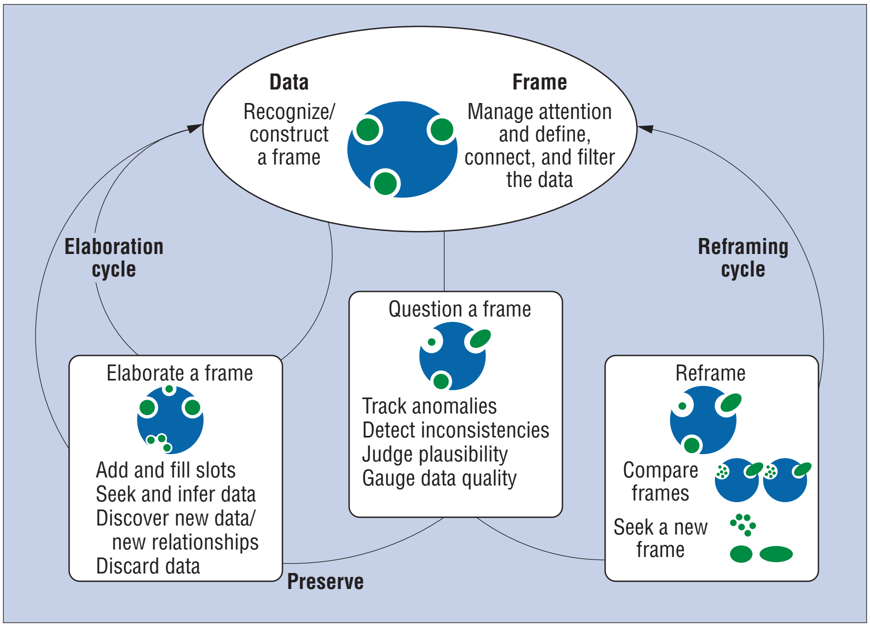


**Figure 1. The Data/Frame Theory of sensemaking.**

## Empirical findings

We examine five areas of empirical findings: causal reasoning, commitment to hypotheses, feedback and learning, sensemaking as a skill, and confirmation bias. In each area the Data/Frame model, and the research it's based on, doesn't align with common beliefs. For that reason, the Data/Frame model cannot be considered a depiction of commonsense views.

### Causal reasoning

Studies of domain practitioners' stories about how they understood real-life decision-making situations suggest that transition sequences—beliefs about what converts one situation into another—are typically based on about three to four causal factors. For example, in explaining why one sports team beat another, newspaper accounts typically focus on a single event such as a critical turnover ("and that cost them the game"), or perhaps that plus one or two other events, such as a star player doing poorly or well. Given the game's length, we can see these as oversimplifications, but most people would skim over any account that tried to capture a game's full complexity. That's why we introduced the billiards metaphor earlier, to illustrate a preference for chains of simple cause-effect relationships. A single causal factor at each junction might be the preferred form of explanation,[2,6] although such explanations open the decision maker up to the reductive tendency.[7]

### Consideration of hypotheses

Decision makers are sometimes advised that they can reduce the likelihood of a fixation error by avoiding early consideration of a hypothesis.[8] But the Data/Frame Theory regards early consideration to a hypothesis as advantageous and inevitable. Early consideration—the rapid recognition of a frame—permits more efficient information gathering and more specific expectancies that can be violated by anomalies, permitting adjustment and reframing. Jenny Rudolph[9] found that decision makers must be sufficiently committed to a frame in order to be able to test it effectively and learn from its inadequacies—something that's missing from open-minded and open-ended diagnostic vagabonding. Winston Sieck and his colleagues have found that domain experts are more likely to question data than novices, perhaps because they're more familiar with instances of faulty data.[10] It might also mean that experts are more confident in their frames and therefore more skeptical about contrary evidence, in contrast to novices who are less confident in the frames they identify.

These observations would suggest that efforts to train decision makers to keep an open mind[11] can be counterproductive, and efforts to make machines that do the vagabonding for the human might be similarly unhelpful. We hypothesize that methods designed to prevent premature consideration to a frame will degrade performance under conditions where active attention management is needed (using frames) and where people have difficulty finding useful frames. Spoon-feeding interpretations to the human (via such methods as data fusion) can be counterproductive.

### Feedback and learning

Another implication of the Data/Frame Theory concerns using feedback to promote learning. Frames are by nature reductive. And yet, frames can help overcome the reductive tendency. The commitment to a frame must be coupled with a motive to test the frame to discover when it's inaccurate. This process hinges on feedback of a certain kind. Outcome feedback ("you got it wrong") isn't nearly as useful as process feedback ("you did it wrong"),[12] because knowing *that* performance was inadequate isn't as valuable as understanding *what* to modify in the reasoning process. This includes the frame itself, because that will determine the way feedback is understood. In other words, people need sensemaking to understand the feedback that might improve sensemaking—the cycle as shown in figure 1. The implication is that people might benefit more from intelligent systems that guide the improvement of frames than from systems that generate alternative understandings and hypotheses and foist them on the human.

### Sensemaking as a skill

We haven't seen evidence for a general sensemaking skill. Some incidents we've collected do suggest differences in motivation—an "adaptive mind-set" of actively looking to make sense of events, as illustrated in essay 1's example of the patient with a pacemaker. It might be possible to develop intelligent systems that acknowledge the Pleasure Principle of human-centered computing[13] and promote a positive motivation to question frames and to reframe, or at least not to frustrate the human and thereby detract from intrinsic motivation. Training might be better aimed at increasing the range and richness of frames, particularly causal mental models, and skill at noticing anomalies. Training scenarios and decision support might be developed for all the sensemaking activities in figure 1 (elaborating a frame, questioning a frame, evaluating a frame, comparing alternative frames, reframing a situation, and seeking anchors to generate a useful frame). Training would aim to provide a larger, richer repertoire of frames rather than to improve each aspect of sensemaking as if it were a separate skill.

### Is there a confirmation bias?

The decision research literature suggests that people are inclined to look for and notice information that confirms a view rather than information that disconfirms it.[14,15] And yet more recent research looking at experts has shown just the opposite. For example, expert weather forecasters have sometimes been observed to deliberately look for information that might disconfirm hypotheses about future severe weather.[16]

The Data/Frame Theory provides a richer understanding of what's actually going on here. People don't engage in simple mental operations of confirming or disconfirming a hypothesis. Our cognitive task analyses of real-world decision making show that skilled decision makers shift into an active mode of elaborating a competing frame once they detect the possibility (or become worried) that the current frame is potentially inaccurate. What might look like a confirmation bias might be simply using a frame to guide information seeking. You need not think of it as a bias and assume that the purpose of an intelligent decision support system must be to help the human overcome some inherent reasoning bias.

## Implications for AI: Reframing frames

Now, the other shoe must drop. Not only might the phenomenon of sensemaking illuminate the computational notion of frames—conversely, that computational notion might challenge our notion of sensemaking.

### Reframing frames

As Marvin Minsky described frames, these organizing structures express the values of features that together define meaningful entities or categories—groups of slots into which the values of defined variables are entered.[17] The primary function of frames (in Minsky's original discussion) is recognition, to guide attention to fill in missing parts of the frame, to test a frame by searching for diagnostic information. To Minsky, frames are things you think with. In the Data/Frame Theory, frames are things that you think with but also things you think about. The Data/Frame Theory therefore blurs the border between phenomenological description and macrocognitive modeling.

We introduced the Data/Frame Theory by suggesting that when we try to make sense of events, we begin with some framework, however minimal. In the Cartesian view of things, sensory inputs (for example, a pattern of moving colored shapes) make contact with memory, lending them meaning in a process called perception ("it's a cat"). But there's a subsequent process, once called "apperception," which interprets the percept more broadly in terms of knowledge (for instance, "I like cats" or "cats can be a symbol for evil"). This is abductive inference, or something rather like it. So the challenge is, where do these frames come from in the first place? Here we see one of AI's outstanding problems, just as it has manifested in numerous views throughout psychology's history:[18] Any computational theory of how knowledge is formed as self-contained bundles should come with a full story about how these proposed "frame-ish" things are supposed to be created and what architectural assumptions underlie them. The AI systems Slate and Cyc both perform abductive reasoning to a plausible explanation using post-Minsky systems based on expressive logics. They both test their hypotheses by actively trying to refute them.

> We introduced the Data/Frame Theory by suggesting that when we try to make sense of events, we begin with some framework, however minimal.

The phenomenon of sensemaking ties also to the notion that frames are chunks of knowledge abstracted away from computational details—symbolic descriptions that are taken off the shelf and used to perceive things, and thereby constitute understanding. However, what we see in studies of sensemaking doesn't fit with this view of frames in three ways. First, understandings shift; frames get changed. They aren't just "taken off a shelf." Frames change as data are acquired (so this isn't just a matter of frame reuse).

Second, even though frames define what count as data (which could be interpreted as a Minskyian notion), as we said earlier, frames themselves actually shape the data (so this isn't just a matter of data primitives). For example, skilled weather forecasters don't passively rely on the data presented by computational aids. Many computer models exist for forecasting weather. Some are based on climate statistics, others on computational models of the atmosphere. Each of these has known biases—for example, a model's tendency to overforecast the depth of low-pressure systems as fronts pass over the Appalachian mountains and the lows reform over the US East Coast. Experienced forecasters take these biases into account and adjust their interpretations of the computer forecasts accordingly.

Finally, frames sometimes have a just-in-time quality. Rarely do decision makers simply identify a relevant mental model. Instead, they construct the frame from smaller sets of causal relationships.

Here too is a challenge for both cognitive science and AI: Any computational theory of how knowledge is formed as self-contained bundles has to come with a full story about how these proposed "frame-ish" things are supposed to be manipulated. In the first essay on sensemaking, we referred to the similarity between sensemaking and mental modeling.[1] Most discussions of psychological research on mental models focus on comparing student and expert mental models, the student's use of mental models to make (erroneous) inferences, and the issue of how to train students to move beyond their naive analogies.[19] We need richer accounts of how the structures are constructed and manipulated.[20] In contrast, work in intelligent systems, such as the Structure Mapping Engine,[21] has specified some process mechanisms.

If frames shape data, how do data mandate any operation on the frame? What defines how frames get changed? Does this require using other frames to govern the frame-changing process? If not, how is it done? If so, what distinguishes the frames being changed from the frames mandating changes? What process can we use to question or doubt a frame? If frames are the vehicle that supports sensemaking, then any doubt would seem to require us to use a doubting-frame to represent the alternative hypothesis (that the frame might be incorrect or faulty). What kind of relationship between frames does this imply?

The phenomenon of sensemaking ties also to the notion that frames organize the large-scale structure of inference making—they're recipes for solving problems. What we see in studies of sensemaking is that frames aren't recipes, although they do play roles in inference making. In AI, the lesson was that knowledge packets are great when you can get them exactly right, so that all you have to do is use them, but you almost never can. For Minskyian frames to be useful, they had to have lots of details, but that would render them of little use across contexts unless they had some sort of internal machinery. Either that, or the frames would have to be chunks that you could pick apart by invoking some external inference machinery. When a frame actually gets used, you must be able to take it apart into the basic facts that constitute it and be able to use the flexibility that this gives you, because the chances of a frame being a perfect fit to particular circumstances are close to zero. Furthermore, if you want to use a frame to tell you what to actually do in a particular circumstance, you need some way to connect the general "type" to the particular exigencies. Computational versions of frames are notoriously poor at this, whereas decision makers are skilled at using frames in these ways.[22]

A pertinent idea from sensemaking studies is that thinking about frames in terms of an either-or (large chunks, with internal inference machinery, versus small chunks, with external inference machinery) might not be the best way to proceed. Indeed, we might think of the progression from novice to expert as a process of learning whereby individual cases, or small, contextually bound understandings with specific inference possibilities attached to them, might develop into larger, more organized understandings. In some cases, these understandings would have inference opportunities bound to them; in other cases, they'd be subject to inference machinery external to them. A process akin to logical inference or proof would assemble multiple small pieces into useful larger pieces. A process akin to pattern recognition might pick up big packets all at once, whereas a process more like mutual matching and fitting might assemble small pieces. In other words, there might be multiple assembly processes. The former might be a mechanism to explain "recognition-primed decision making," when the expert goes from an immediate understanding of a situation to a course of action. The latter would constitute the cycles described in figure 1. This is reminiscent of Cyc, in which bundles of related assumptions and concepts ("microtheories") are logic-like in their small structure but frame-like in that they're large and specific to a topic or concept and the knowledge inside them is specific to them.

We relied on the frame concept in the Data/Frame Theory as a metaphor to bootstrap a discussion of how people create, use, and manipulate organizing structures. We do not offer any clear path to a computational theory of how "frame-ish" things are created or manipulated. Our main goal in discussing the Data/Frame Theory is to point to empirical studies of how domain practitioners make decisions in complex, real-world contexts and then to mine these results for ideas that might invigorate and inform work on these fundamental issues. ■

## Acknowledgments

Robert Hoffman's work on this essay was supported through his participation in the Advanced Decision Architectures Collaborative Technology Alliance, sponsored by the US Army Research Laboratory under cooperative agreement DAAD19-01-2-0009.

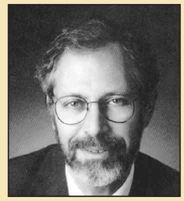

**Gary Klein** is chief scientist in the Klein Associates Division of Applied Research Associates. Contact him at Klein Associates, 1750 Commerce Center Blvd. N., Fairborn, OH 45324; gary@decisionmaking.com.

**Brian Moon** is a research associate in the Klein Associates Division of Applied Research Associates. Contact him at Klein Associates, 1750 Commerce Center Blvd. N., Fairborn, OH 45324; brian@decisionmaking.com.

**Robert R. Hoffman** is a senior research scientist at the Institute for Human and Machine Cognition. Contact him at IHMC, 40 So. Alcaniz St., Pensacola, FL 32502-6008; rhoffman@ihmc.us.

## References

1. G. Klein, B. Moon, and R.R. Hoffman, "Making Sense of Sensemaking 1: Alternative Perspectives," *Intelligent Systems,* vol. 21, no. 4, 2006, pp. 70–73.
2. G. Klein et al., "A Data/Frame Theory of Sensemaking," to be published in *Expertise Out of Context: Proc. 6th Int'l Conf. Naturalistic Decision Making*, R.R. Hoffman, ed., Lawrence Erlbaum Associates, 2006.
3. K.E. Weick, *Sensemaking in Organizations*, Sage Publications, 1995.
4. C.A. Chinn and W.F. Brewer, "The Role of Anomalous Data in Knowledge Acquisition: A Theoretical Framework and Implications for Science Instruction," *Rev. Educational Research*, 1993, vol. 63, pp. 1–49.
5. P.J. Feltovich, R.J. Spiro, and R.L. Coulson, "Issues of Expert Flexibility in Contexts Characterized by Complexity and Change," *Expertise in Context: Human and Machine*, P.J. Feltovich, K.M. Ford, and R.R. Hoffman, eds., AAAI/MIT Press, 1997, pp. 125–146.
6. G.A. Klein and B.W. Crandall, "The Role of Mental Simulation in Naturalistic Decision Making," *Local Applications of the Ecological Approach to Human-Machine Systems*, vol. 2, P. Hancock eds., Lawrence Erlbaum Associates, 1995, pp. 324–358.
7. P.J. Feltovich, R.R. Hoffman, and D. Woods, "Keeping It Too Simple: How the Reductive Tendency Affects Cognitive Engineering," *IEEE Intelligent Systems*, May/June 2004, pp. 90–95.
8. R. Heuer, *The Psychology of Intelligence Analysis*, tech. report, CIA Center for the Study of Intelligence, 1999.

9. J.W. Rudolph, "Into the Big Muddy and Out Again," doctoral dissertation, Boston College, 2003; http://escholarship.bc.edu/dissertations/AAI3103269.

10. W.R. Sieck et al., "Basic Questioning Strategies for Making Sense of a Surprise: The Roles of Training, Experience, and Expertise," *Proc. 26th Ann. Conf. Cognitive Science Soc.* (CogSci 04), Lawrence Erlbaum Associates, 2004; www.cogsci.northwestern.edu/ cogsci2004/ma/ma305.pdf.

11. M.S. Cohen et al., "Dialogue as the Medium for Critical Thinking Training," *Proc. 6th Int'l Conf. Naturalistic Decision Making*, Lawrence Erlbaum Associates, 2003.

12. E. Salas et al., "Myths to Avoid about Crew Resource Management Training," *Ergonomics in Design*, Fall 2002, pp. 20–24.

13. R.R. Hoffman and P.J. Hayes, "The Pleasure Principle," *IEEE Intelligent Systems*, Jan./Feb. 2004, pp. 86–89.

14. C.R. Mynatt, M.E. Doherty, and R.D. Tweney, "Consequences of Confirmation and Disconfirmation in a Simulated Research Environment," *Quarterly J. Experimental Psychology*, vol. 30, 1978, pp. 395–406.

15. P.C. Wason, "On the Failure to Eliminate Hypotheses in a Conceptual Task," *Quarterly J. Experimental Psychology*, vol. 12, 1960, pp. 129–140.

16. R.R. Hoffman, G. Trafton, and P. Roebber, *Minding the Weather: How Expert Forecasters Think*, MIT Press, 2006.

17. M. Minsky, "A Framework for Representing Knowledge," *The Psychology of Computer Vision*, P.H. Winston, ed., McGraw-Hill, 1975, pp. 211–277.

18. H.L. Dreyfus, "A Framework for Misrepresenting Knowledge," *Philosophical Perspectives in Artificial Intelligence*, M. Ringle, ed., Humanities Press, 1979, pp. 110–123.

19. D. Gentner and A. Stevens, eds., *Mental Models*, Lawrence Erlbaum Associates, 1983.

20. J. Greeno, "Conceptual Entities," *Mental Models*, D. Gentner and A. Stevens, eds., Lawrence Erlbaum Associates, 1983, pp. 227–252.

21. B. Falkenhainer, K.D. Forbus, and D. Gentner, "The Structure-Mapping Engine: Algorithm and Examples," *Artificial Intelligence*, vol. 41, 1990, pp. 1–63.

22. G. Klein, *Sources of Power: How People Make Decisions*, MIT Press, 1998.

For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.