

Decision(?)Making(?)

Robert R. Hoffman, *Institute for Human and Machine Cognition*
J. Frank Yates, *University of Michigan*

Consider, for a moment, “What are we making intelligent decision aids *for*?” Computers, including intelligent systems, assist human decision making in many ways.^{1–3} Decision aids can range from an online tool provided by

Consumer Reports to help people choose a refrigerator, to a large system for monitoring an industrial process. Group decision aids focus on supporting communication.^{4,5} Expert systems can be considered decision aids, as can systems that use statistical methods to assist in diagnostic procedures. Decision aids can provide information involved in deciding or information pertinent to evaluating states of the world. To support these situation assessment and monitoring activities, computers can integrate and display information and assist in replanning.⁶ When we deconstruct such generic tasks, most of the component tasks reduce to option generation, option selection, and outcome evaluation.⁷ For these, we can bring to bear various mathematical techniques such as utility analysis.⁸

But how well do such tools really help humans deal with the difficulties of deciding itself—for instance, adapting to changing circumstances or coping with situations that are both unfamiliar and infrequent?⁷ While several good analyses of decision-making situations and situational factors exist that can help shape decision-aid architectures and functionalities,^{9,10} has anyone done a corresponding analysis of what this thing called the *decision* is? After



discussing this question, we’ll sharpen the focus for new and potentially useful applications of intelligent systems technologies.

What’s a “decision”?

The word *decide* gets a great deal of mileage in English, being critical in the definitions of dozens of other concepts, including *arbitrate*, *conclude*, *convict*, *declare*, *define*, *disagree*, *intervene*, and *judge*. Another use is adverbial: the idea of acting “decisively,” achieving a final determination with clear and definitive intent. “Decide” also works as a transitive verb, meaning to influence or determine: for example, “This new development finally decided me,” or “The vote in New Hampshire often decides the outcome of the presidential election.”

The modern English word “decision” comes from the Sanskrit “/khidati/”—meaning to tear and then to Latin “caedere,” meaning to kill or cut down (as in battle), and then to “de + caedere” meaning to cut off from or to cut thoroughly.^{11–13} We see here a notion that the decision is a point, and a final point or action separating one thing from another (historically, a human head or limb from its body). All dictionary definitions preserve this sense of bringing a series of events, including a mental sequence, to a final, point-like conclusion.¹¹ For example, WordNet 1.7.1 (<http://wordnet.princeton.edu/online>) defines “decide” as “To make up one’s mind, to reach, make, or come to a decision about something.” While the circularity is obvious and troublesome, this essay starts with the idea that decisions are things that are “made.”

You can’t help but be struck by the irony that many scholarly books on judgment and decision making typically forego an attempt at defining this thing called “decision.”^{14–19} Only in one textbook²⁰ and one edited volume²¹ do we find the word “decision” even as an entry in the subject index. Only in a few sources do we find an explicit discussion of the difference between “decision” and “judgment.” (A “judgment” is an assessment or opinion as to what was, is, or will be the state of some decision-relevant aspect of the world—for example, whether a defendant committed a crime, whether a patient



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

Table 1. The received view of decision making, with examples.

| Case | Step 1 | Step 2 | Step 3 |
|---|--|--|--|
| General | Information that pertains to the decision is acquired. | Span of apprehension; the information is perceived and interpreted. | Commitment to action: the decision is made. |
| An everyday example | Dee arrives at the airport only to learn that her flight is delayed because of fog. | She apprehends or infers an impending threat—missing her meeting. She reasons that she should perhaps call ahead to Tom to ask for a delay in their meeting. | There is a mental act, not necessarily deliberative, and she begins to reach for her cell phone to make the call. |
| How can we get to the Moon? (known as the Apollo Mode Decision) ²⁴ | US President John F. Kennedy decided, for political reasons, to begin a space program with the goal of getting to the Moon in a certain time frame—before the Soviets. | This implied the threat of not actually achieving that goal. So, the decision problem involved committing to a program intended to avert that threat. In months of debate, reports, and meetings, the participants discussed and debated many alternative plans, sometimes in the presence of President Kennedy. | Lunar-orbit rendezvous was chosen as the best chance of winning the race to the Moon because it involved relatively little technological innovation compared to the alternative schemes. |
| Where should the atomic bomb be dropped? ²⁵ | The summer of 1945 was named as the most likely date when sufficient production would exist to make constructing an atomic bomb possible. A test in New Mexico was held six days after sufficient material was available for the first bomb. | Target selection began in the spring of 1945. Some important considerations were the aircraft's range, the desirability of visual bombing, probable weather, targets to produce the greatest military effect, target susceptibility, and, to determine an atomic bomb's effect, targets untouched by previous bombing. | It was decided that Hiroshima and Nagasaki met the criteria for the primary targets. |

Table 2. Variations on the three-step scheme in various domains' literature.

| Literature domain | Step 1 | Step 2 | Step 3 |
|--|---|---|--------------------------|
| Decision aiding for process control | Situation assessment | Planning | Commitment ⁹ |
| Expert systems | Information acquisition | Modeling (via knowledge bases and decision trees) | Commitment ²⁶ |
| Naturalistic decision making ²⁷ | Situational awareness | Inference | Action |
| Problem solving | Staging activities, such as identifying problems and specifying goals | Apprehension, involving perceiving and comprehending activities such as the generation of possible solutions, implementation plans, and methods for evaluating outcomes | Selecting or choosing |

Table 3. Unpacking the everyday example.

| Case | Step 1. Information acquisition | Step 2. Perception and interpretation | Step 3. Commitment |
|----------------------|--|---|--|
| The example | Dee arrives at the airport only to learn that her flight is delayed because of fog. | She reasons that she should perhaps call ahead to Tom to ask for a delay in their meeting. | She begins to reach for her cell phone to make the call. |
| The example unpacked | After hearing of the fog delay, Dee looks out of the panoramic window at the concourse to perceive the weather directly. How foggy is it? Did she pause to deliberately decide, "Hey, I should look out the window"? Perhaps it was intuitive or automatic, perhaps not. | In thinking about the fog delay's implications, Dee mentally simulates possible futures. This takes into account the fact that she's hungry and she realizes that she has time to grab ... what? Perhaps a cup of coffee and a bagel. | While pushing the buttons on her cell phone, she's alert to signal strength ("What do I do if it's low?") and battery charge level ("What do I do if I forgot to recharge it last night?") |

has a particular disease, or whether product sales will exceed the break-even point.^{14,22,23}

The final-point notion

A decision can be defined as a commitment to a course of action having the intention of serving the interests and values of particular people.²³ A decision is generally understood as a mental event that occurs at a singular point in time—a psychological moment of choice—that leads immediately or directly to action (for example, push the button, or wait 30 seconds and then push

the button). Most views of decision making in the literatures of psychology and of judgment and decision making locate the notion of decision in the final stage of a stage-theoretic framework. Table 1 provides three specific examples. Table 2 presents variations on the three-step scheme at the core of a number of both normative and descriptive models of decision making.^{10,18}

The final-point notion allows us to say that decisions are "made," but even in apparently simple, clear-cut cases, the deciding process is much more than this.

Beyond three-step descriptions

The process of deciding (especially the big decisions as opposed to snap decisions) often entails making a number of component decisions. But more importantly, the decision process entails a host of significant cognitive, evaluative, and affective activities that are parallel and interactive. We illustrate this in table 3 by deconstructing one of table 1's examples.

While deciding involves acquiring information, that in and of itself might itself involve other decisions and deliberations

over future possibilities, preferences, options, and goals. This shows in the Apollo Mode Decision (see table 1), which began with other decisions as the informational starting point. While most three-step views regard decisions as culminations, decisions are often expressions of contingencies and anticipations of unfolding events and ways in which future events might be surprising. Consider the atomic bomb decision. An initial decision problem was where to drop the bomb, but “the” decision that came out of the process wasn’t just where to drop, but to try to develop a capability for generating accurate weather forecasts for Japan and then choose among targets and *then* decide *whether* to drop. Weather and weather-forecasting capabilities that would have to be created following the decision could not only affect the way the decision would be carried out but even forestall its implementation.

On this, three-step models are potentially misleading. You can always unpack any given three-step model into embedded three-steps, each having its own moment of choice. New decision problems constantly arise, either in the process of implementing a previous commitment or perhaps because that previous decision instigated new threats and opportunities. Often, there’s no single “end” point. Notice in the Apollo Mode and atomic bomb examples that a network of contingencies and interdependencies existed. You might say there was a series of decisions, each of which could be unpacked as we did for the everyday example. It’s perhaps possible to think of life as an endless chain of dominos—single commitments—each of which we can analyze, possibly to good effect. Despite the apparent popularity of three-step models, everyone would acknowledge that decisions are complex in this way. Thus, we could argue that in the table 3 reconstruction of the everyday example, we’ve merely decomposed each step into its own embedded three-step.

That might be the issue because, as in the vast chunk of the modern scientific literature on decision making, “the decision” would still be regarded as a point-like thing, a singular commitment, that marks the end of a sequence of clear-cut mental operations.

In *Communication and Group Decision Making*,²⁸ editors Randy Hirokawa and Marshall S. Poole ask “What, exactly, is a decision?”

Decisions are assumed to be discrete events, clearly distinguishable from other group activities Decision makers often can identify discrete decision points and feel a sense of completion at making a decision. These boundaries are not always as clear as they seem at first, however, and there is not always agreement on what events are involved in a given decision. Definitions of decision making episodes are ambiguous in several respects. (p. 9)

This volume comes from a different tradition than the psychology of judgment and decision making—specifically, the field of communications research. It also has a different emphasis: decisions’ social context and “embeddedness” rather than the mental events that are believed to underlie decision making.

The scholarly volume, *Decision Making in Action*,²⁹ comes from yet another perspective—the field of “naturalistic decision making.” This descriptive approach

Three-step models are potentially misleading. You can always unpack any given three-step into embedded three-steps, each having its own

focuses on how experts perform and reason in real-world, complex domains where decisions are often high-stakes, high-risk, and made under considerable time pressure. The seminal studies in this area led to dissatisfaction with the literature’s normative models. The authors latched onto a fact—that traditional decision research has been preoccupied with the issue of trade-offs:

The basic cause of the mismatch is that traditional decision research has invested most of its energy in only one part of decision making, which we shall refer to as the decision event. In this view, the crucial part of decision making occurs when the decision maker (generally a single individual) surveys a known and fixed set of alternatives, weighs the likely consequences of choosing each, and makes a choice ... The decision making activities suggested by [the naturalistic studies] offer few clean examples of decision events. (p. 5)

Verbs, not nouns

We assert the following premises:

1. Whenever we carve out a three-step model, each step itself will be some other mental event involving judging or deciding.
2. The rule, rather than the exception, is that the commitment to action involves a contingency pointing to possible worlds and future contingencies, with the intention of scaffolding a capability to recognize when to be surprised.

These premises basically say that any individual three-step is necessarily incomplete as a description of a deciding event, even though one might be able to identify a moment of choice and call it “the decision.”

We propose thinking of *decision making* in terms of *deciding*, regarding it as one of a number of macrocognitive processes that it supports and that support it.²⁷ This view hearkens to Franz Brentano’s ideas,³⁰ who regarded all mental representations as interactions between constantly fluxing memory and perceptual activity, which always have a judgmental aspect. From Brentano’s dynamic psychology comes the view that we should speak of mental phenomena using verbs rather than nouns—deciding, not the decision.

The richness of deciding

Deciding involves many factors beyond those notions that the main theories have captured.^{8,31} In the real world,

- deciding involves instantiating intentions and purposes;
- deciding is usually about causing good things to happen;
- actions are intended to bring about states of affairs that serve the interests and tastes of particular individuals or groups;
- commitments to act must be distinguished from action because, for various reasons, not all decisions are actually implemented; and
- choice among alternatives is never equivalent to choice among consequences, because alternatives rarely lead to single consequences.

A decider’s reasoning derives from a host of “deep” contributors that are important to understand in their own right—for example, unique personal experiences,

constitutional factors such as inherited dispositions or abilities, training, and culturally transmitted local customs. Aspects of decision making that have received the most attention in decision research are the evaluation of options, the anticipation of possibilities, judgment and reasoning biases, and the motivating values that particularize and specify individual decision-making episodes.

However, before decisions are “made” and while they are being “made,” the decider attends to one or more important preliminary issues. Empirical work in myriad fields (such as psychology, health care, management, finance, engineering, law, operations, anthropology, counseling, politics, and marketing) as well as analyses of many hundreds of incidents converge on a number of fundamental questions that arise in real-life decision problems. That’s why we use the phrase *cardinal decision issues* to describe them.³¹ We can characterize decision processes as the means by which these cardinal issues are addressed.

Need

Why are we (not) deciding anything at all? This issue is about whether and how decision problems are recognized in the first place, how people come to recognize that existing or developing circumstances constitute threats or opportunities. In the former case, unless something is done, those people will be harmed, and in the latter, they will miss out on the chance to improve their situation. Deciding, in this view, is about arriving at commitments to actions that are intended to meet impending threats or opportunities. This issue strongly links decision making to notions of vigilance, problem-finding, and recognition-primed decision making.^{27,32}

Mode

Who will decide, and how will they approach that task? How will those individuals address the other cardinal issues that must be resolved? A major part of the “who” question concerns whether to defer authority and to whom authority is deferred. The “how” question of the mode issue is about the nuts and bolts of how deciders carry out their work. The decision-making literature discusses several broad categories of possibilities, including (but not limited to) analytic, rule-based, automatic, and so-called intuitive decision making. Another aspect

of mode is deliberation over whether to seek opinions, and from whom. Deciding often benefits from opinion-seeking,^{20,23} yet relatively little is known about how people evaluate and aggregate the verbal opinions and evaluations they receive.

Investment

What kinds and amounts of resources will we invest in the process of deciding? This issue is about how and how well we determine whether the investment of resources in the process of deciding—for example, time, expertise, or tools—is appropriate, neither too little nor too much. Two important considerations are as yet unstudied systematically:

- The evaluation of resources according to resource categories. Resources will rarely all be of the same type (for instance, human

We propose thinking of decision making in terms of “deciding,” regarding it as one of a number of macrocognitive processes that it supports and that support

resources, time, and material), and evaluations of investment will differ for such categories.

- The strategy of minimizing the costs of deciding by deliberately engaging “direct,” nonanalytic modes of deciding (such as intuitive decision making).³³

Options

What are the different actions we could potentially take to deal with this problem? Evaluating options, especially in terms of their costs and benefits, has been a central focus in the field of judgment and decision making. But this isn’t the activity we refer to here. The options issue is about how people come to apprehend some prospective solutions to their decision problems but somehow never even recognize the existence of others. The issue’s significance is implicit in the truism that you can’t choose an alternative you don’t know about. It’s essential

to recognize that expert navigation of the options issue isn’t about increasing the number of alternatives considered (a false assumption that many scholars make when asserting that people are limited in their ability to consider multiple options). The ideal “option consideration set” for a given problem consists of only a single alternative—the best one. Recognizing, let alone deliberating, over other options is often wasteful, requiring the decision maker to expend precious time and resources vetting alternatives that ultimately will (or should) be rejected. But recent work has demonstrated that deliberating over large consideration sets can do more than simply waste time. It can also exact significant psychological costs, such as turmoil over the possibility of failing to pick the best alternative.³⁴

Possibilities

What are the various things that could happen if we took that action, and which ones do we care about? This issue also involves the macrocognitive functions of sensemaking and mental projection—recognizing outcomes of prospective actions that are capable of occurring, and which would matter greatly if they were to occur. The concern isn’t with whether those outcomes will or would occur, only whether they could.

Judgment

Which of the things that we care about actually would happen if we took that action? Although this issue might logically follow the possibilities issue, in the real world the fuller process is macrocognitive. If a decider recognizes (accurately or otherwise) that some decision-relevant event can happen, there must then be a judgment as to whether it would happen. One of us (Frank Yates)²⁰ observed that there are two classes of judgment processes: formalistic and substantive. Formalistic procedures are exemplified by the application of rules such as those in probability theory or regression analysis. Significantly, such rules are largely indifferent to the content of judgment problems. Quite the opposite is true of substantive procedures, which entail the attempt to envision how the world would (or wouldn’t) literally create the event in question. Recognition-primed decision making²⁹ is one such substantive procedure. There are many indications in the literature that people resort to formalistic procedures only when they can’t

use substantive ones, which seem to be more “natural.”²⁰

Value

How much would anyone really care—positively or negatively—if this particular outcome happened? The value issue is a special case of the judgment issue, albeit an exceptionally important case. That’s because it centers on what makes decision problems so distinctive and difficult—individual differences in what people like and dislike. But what is a “good” decision depends on the parties involved. And this gets back to the core idea that deciding involves a commitment to actions that are intended to result in outcomes that are satisfying to particular people. In order for a decider to pursue actions that promise outcomes that the intended beneficiaries find satisfying—the goal of any decision-making effort—the decider must know those persons’ tastes. That is, a decider must make judgments about how other people feel about things.

Trade-offs

All our prospective actions have both strengths and weaknesses. So how should we make the trade-offs that are required to settle on the action we’ll actually pursue? This issue concerns the fact that in many (most?) real-world deciding situations, deciders eventually arrive at this reality: Every alternative has drawbacks. Expected utility theory, the point of reference for the field, is at heart about the trading-off of outcome value and uncertainty.^{20,35} Decision research has been dominated by questions about deviations of people’s actual decision behavior from what is predicted or prescribed by rules such as the expected utility, additive utility, and discounting models. This perspective reflects a narrow, idealistic, rationalistic conception of how people deal with trade-offs in real life. It’s also a reflection of the reductive tendency.³⁶ The dominant models presume a “pick among these” stance by the decider. Evidence shows that a major tactic deciders use is to transform trade-off problems into options problems.³⁷ Specifically, people sometimes seek to avoid having to make an onerous trade-off altogether by finding or creating a new alternative that makes the trade-off unnecessary. In many cases, deciders have reached their decisions saying, for example, “So this is what we are going to do,” once they’ve resolved the trade-offs issue. But that’s

never the end of things. In the “aftermath,” often before anything else occurs, acceptability and implementation issues can take center stage—for example, how various other parties feel about how the decision was made.

Acceptability

How can we get the other stakeholders to agree to this decision procedure? In most high-stakes situations, the decider isn’t a free agent but must contend with many stakeholders’ sentiments concerning what’s decided, how it’s decided, and how it’s implemented. Negotiations are the most familiar context where the acceptability issue figures significantly, but the acceptability issue assumes significance beyond the realm of formal negotiations. American automakers have lost several major lawsuits

Disasters with respect to implementation generally result from the prior mishandling of one or more of the other cardinal issues when the original decision was being deliberated.

because they mishandled the acceptability issue in design decisions. In one prominent case, jurors were repelled by testimony that the decision to cut costs on certain features rested partly on a decision analysis in which a dollar figure (based on actuarial records) was attached to lives that might be lost in accidents linked to those features. The jurors responded by forcing the company to pay billions in punitive damages.³⁸

Implementation

That’s what we decided to do. Now, how can we get it done? Or can we get it done, after all? While a decision aid might seek to get people to a point of commitment, events follow the commitment. As we suggested earlier in unpacking the “commitment to action,” even acting has an element of deciding because it can and often does involve contingency. A commitment to act

doesn’t necessarily have action as its primary functionality. Rather, it’s a resolution to accept a particular understanding in the hope that the understanding will serve to help the decider know when to be surprised after the action has commenced and the anticipated contingencies play themselves out, or not. The implementation issue is particularly important in decision situations where the selected alternative entails a nontrivial “project” that must be executed, as opposed to a single action that’s virtually synonymous with the decision itself (for example, the final-point notion). Sometimes, a project proves to be difficult or even impossible to actually conduct and entails other decisions, even waves of decisions (as in the atomic bomb example). Disasters with respect to the implementation issue generally result from the prior mishandling of one or more of the other cardinal issues when the original decision was being deliberated—for instance, overlooking various implementation barriers and therefore failing to see them as possibilities. There’s been little systematic research aimed at understanding how people address the implementation issue (as an exception, see Utpal Dholakia and Richard Bagozzi’s work³⁹).

Implications for intelligent systems

The modeling component of many decision aids basically involves taking input data; creating tabular representations of entities, attributes, and weightings; and performing mathematical operations predicated on notions of decision analysis.³ Decision analysis has many benefits. It offers comfortable means to describe decision making in terms of choice among probability distributions. It offers techniques to mathematically specify preferences, derive and evaluate probabilities, and work on equations that balance gain and risk. It provides mathematical methods to achieve consistency by rules of logic. This approach prescribes a decision process that involves identifying promising prospective courses of action and their potential significant consequences (step 1), assessing the utility of those consequences and evaluating the likelihoods of all the recognized potential outcomes (step 2), and then selecting the alternative that’s indicated to be best according to a “rational” decision rule (step 3).

Assuming this is what deciding is, then surely people must need help with these

things.⁸ Over the years, this view has had a substantial influence on the character of the literature on human biases and limitations. But the promise has not caught up to the reality:

Behavior-focused decision aids have had little documented success ... decision quality entails myriad diverse facets ... yet the typical decision aid (and its theoretical underpinning) is predicated on a narrow conception. ... Deciders therefore often ignore such aids because they appear irrelevant to significant decider concerns. And when deciders do try the aids, the results disappoint them because the aids leave untouched the quality dimensions that matter to them.³ (p. 13)

The benefits of the system must be apparent ... The degree to which the judge will be held responsible for the judgment must be made evident, and the quality of the information the system is supplying must be indicated. In short, unless the judge sees a need to bother with the support system, the work put into its design will be of no avail.⁴⁰ (p. 124)

A case in point is the new decision aids that assist intelligence analysts. These take the analysts away from meaningful study of intelligence information and force them to engage in evaluations of probabilities and hypotheses. They do enough of that as is. The make-work and overhead necessitated by the new systems sometimes outweigh any value added.⁴¹

In the case of expert systems (regarded as decision aids), it was clear during so-called first-generation work that the structures of knowledge bases and brittle (or context-insensitive) procedural rules didn't capture the subtleties of expert deciding.²⁶ Likewise for aids based on decision analysis, there's been little if any concern about some of the tough and crucial aspects of deciding. These include determining whether there's a significant decision problem to solve in the first place, developing promising alternatives, envisioning nonobvious but critical potential side effects of alternatives, and discerning how key parties truly feel about possible outcomes of selected options as well as the decision process itself. When decision analyses acknowledge such matters at all, they assume that they've been addressed outside the analyses per se. But there is reason to believe that it's precisely these other tough and crucial aspects of deciding that often spell the difference between effective and ineffective deciding and thus are ones that people need help dealing with.

Table 4. Some possibilities for intelligent decision aids.


| Cardinal issue | Intelligent systems might ... |
|----------------|---|
| Need | Help people monitor and recognize threats and opportunities that warrant efforts to make decisions that address them (for example, track trends in consumer needs and tastes for services). |
| Mode: Who? | Help people deliberate on whether and to whom authority should be delegated (for example, assess product proposal review skill requirements). |
| Mode: How? | Help people determine how well alternative procedures for addressing various other cardinal issues are suited to present circumstances (for example, monitor best practices databases for new product appraisal techniques). |
| Investment | Help people monitor and minimize decision process costs without jeopardizing other quality dimensions (for example, track trends in product appraisal expenses). |
| Options | Help people scan for and filter existing alternatives and organize their efforts to create new ones (for example, recommend participants in product creativity exercises according to their track records and personal characteristics, such as intellectual diversity). |
| Possibilities | Help people envision nonobvious but real and significant potential consequences of alternatives under consideration, such as "side effects" distinct from intended effects (for example, identify ways that potential new products might be misused, resulting in product liability claims). |
| Judgment | Help people anticipate the actual states of decision-relevant events and conditions (for example, structure the deliberations of supply chain buyers to best exploit their expertise in predicting future product component price changes). |
| Value | Help people assess how the various parties to a decision feel about potential consequences (for example, administer new product preference assessment and forecasting exercises under realistic product use simulation scenarios). |
| Trade-offs | Help people decide how to deal with trade-offs, including possibly obviating the need for trade-offs by identifying better alternatives or transforming trade-offs into opportunities (for example, manage interactive routines in which consumers can test their beliefs about new product feature trade-offs, such as style, convenience, and price, in vivid simulations). |
| Acceptability | Help people anticipate how various stakeholders will regard a prospective decision or the process used to make it and craft ways of achieving their acceptance (for example, guide a review of parties who have a stake in the introduction of a new product class and routines for negotiating with them). |
| Implementation | Help people anticipate common impediments to decision implementation before finalizing decisions, and respond quickly and effectively to impediments that occasionally arise despite those efforts (for example, guide impediment-anticipation exercises intended to surface new product manufacturing glitches after volume or scale ramp-up). |

At one level, this is all terribly disappointing. But the perspective afforded by the cardinal decision issues may broaden the horizons for intelligent decision support systems. That perspective would point toward concrete ways that a new generation of systems might complement and extend human capabilities in ways that could result in decision processes that add significant, demonstrable value. Table 4 takes the cardinal issues one at a time and suggests the kinds of enhancements that intelligent technologies might make to how people ordinarily address those issues. For concreteness, we illustrate the ideas with a running product development

illustration. Of course, not all of these suggestions are entirely new. For example, Lee Beach suggested a form of aid that would help decision makers construct explanatory narratives (that is, representations of their mental models).⁴⁰

When the expert blacksmith hammers away at the anvil, he can make precise impacts time and time again, at just the right spot. Yet, measurement of the movements shows that the strokes are never exactly the same.⁴² When we trace

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.



J. Frank Yates is a professor in the Department of Psychology at the University of Michigan, where he serves as acting director of the Afro-American Studies Program and director of the Coalition for the Use of Learning Skills. Contact him at 3038 East Hall, Univ. of Michigan, Ann Arbor, MI 48109; jyates@umich.edu.

the history of a decision process, it's always possible to identify one or more moments of choice. We can then describe history in terms of causal steps leading up to that moment, creating a simple causal model that might then be amenable to specification in terms of rules. But when we look at deciding as it occurs, a different picture emerges. Like the blacksmith's process, people can reach moments of commitment that signal their occurrence clearly but are never achieved by following precisely the same path. People are not engaging a cause-effect chain or a rule-based process. They're navigating a space of constraints and issues, involving contingencies and contextual dependencies. Capturing such dynamics and interactions in ways that avoid making causal-chain theories is always a challenge.⁴³ Punctuated histories are what falls out as a result of our telling stories. Those who would create intelligent decision architectures might benefit from considering a macrocognitive view of deciding, one that's significantly richer than the domino three-step. To make intelligent decision aids that are maximally useful, designers might focus on trying to enhance consequential elements of the entire decision process, not just what occurs in the analysis of trade-offs to culminate in a single moment of commitment. ■

Acknowledgments

We thank John Flack for his comments on an early draft. Robert Hoffman's contribution was through participation in the Advanced Decision Architectures Collaborative Technology Alliance, sponsored by the US Army Research Laboratory under Cooperative Agreement DAAD19-01-2-0009.

References

1. G.E.G. Beroggi and W.A. Wallace, "Closing the Gap: Transit Control for Hazardous Material Flow," *J. Hazardous Materials*, vol. 27, 1991, pp. 61–75.
2. G. Wright and F. Bolger, eds., *Expertise and Decision Support*, Plenum, 1992.
3. J.F. Yates, E.S. Veinott, and A.L. Patalino, "Hard Decisions, Bad Decisions: On Decision Quality and Decision Aiding," *Emerging Perspectives on Judgment and Decision Research*, S.L. Schneider and J. Shanteau, eds., Cambridge Univ. Press, 2003, pp. 13–63.
4. P. Reagan-Cirincione and J. Rohmbaugh, "Decision Conferencing," *Expertise and Decision Support*, G. Wright and F. Bolger, eds., Plenum, 1992, pp. 181–202.
5. G. Rowe, "Perspectives on Expertise in the Aggregation of Judgment," *Expertise and Decision Support*, G. Wright and F. Bolger, eds., Plenum, 1992, pp. 155–180.
6. K. Hammond, "Explaining and Repairing Plans That Fail," *Artificial Intelligence*, vol. 45, 1990, pp. 173–228.
7. W.B. Rouse, "Design and Evaluation of Computer-Based Decision Aids," *Human-Computer Interaction*, G. Salvendy, ed., Elsevier, 1984, pp. 229–246.
8. B. Fischhoff, "Decision Making in Complex Systems," *Intelligent Decision Support in Process Environments*, E. Hollgagel, G. Mancini, and D. Woods, eds., Springer-Verlag, 1986, pp. 61–85.
9. G. Johanssen, "Architecture of Man-Machine Decision Making Systems," E. Hollgagel, G. Mancini, and D. Woods, eds., *Intelligent Decision Support in Process Environments*, Springer Verlag, 1986, pp. 327–339.
10. F.D. Rigby, "Heuristic Analysis of Decision Situations," *Human Judgments and Optimality*, M.W. Shelly and G.L. Bryan, eds., John Wiley & Sons, 1964, pp. 37–44.
11. *Webster's New Collegiate Dictionary*, G. & C. Merriam, 1979.
12. E. Partridge, *Origins: A Short Etymological Dictionary of Modern English*, Macmillan, 1958.
13. W. Smith, *Latin-English Dictionary*, John Murray, 1933.
14. H.R. Arkes and K.R. Hammond, eds., *Judgment and Decision Making: An Interdisciplinary Reader*, Cambridge Univ. Press, 1986.
15. P. Juslin and H. Montgomery, eds., *Judgment and Decision Making*, Lawrence Erlbaum, 1999.
16. D. Kahneman and A. Tversky, eds., *Choices, Values, and Frames*, Cambridge Univ. Press, 2000.
17. S. Plous, *The Psychology of Judgment and Decision Making*, McGraw-Hill, 1993.
18. K. Smith, J. Shanteau, and P. Johnson, eds., *Psychological Investigations of Competence in Decision Making*, Cambridge Univ. Press, 2004.
19. R.J. Sternberg and P.A. Frensch, eds., *Complex Problem Solving*, Lawrence Erlbaum, 1991.
20. J.F. Yates, *Judgment and Decision Making*, Prentice Hall, 1990.
21. O. Svenson and A.J. Maule, eds., *Time Pressure and Stress in Human Judgment and Decision Making*, Plenum Press, 1993.
22. K.R. Hammond, G.H. McClelland, and J. Mumpower, *Human Judgment and Decision Making: Theories, Methods, and Procedures*, Praeger, 1980.
23. J.F. Yates and M.D. Tschirhart, "Decision Making Expertise," *Cambridge Handbook on Expertise and Expert Performance*, A. Ericsson et al., eds., Cambridge Univ. Press, to be published in 2006.
24. P.E. Mack, "The Apollo Mode Decision," 1997; <http://people.clemson.edu/~pammack/apmode.htm>.
25. "Manhattan Project History: The Atomic Bombings of Hiroshima and Nagasaki by The Manhattan Engineer District, June 29, 1946," Manhattan Project Preservation Assoc., 17 Oct. 2003, www.childrenofthemanhattanproject.org/HISTORY/H-05.htm.
26. J. Gammack, "Knowledge Engineering Issues for Decision Support," *Expertise and Decision Support*, G. Wright and F. Bolger, eds., Plenum, 1992, pp. 203–226.
27. G. Klein et al., "Macrocognition," *IEEE Intelligent Systems*, May/June 2003, pp. 81–85.
28. R.Y. Hirokawa and M.S. Poole, eds., *Communication and Group Decision Making*, Sage Publishing, 1996.

29. G. Klein et al., eds., *Decision Making in Action: Models and Methods*, Ablex Publishing, 1993.
30. F.C. Brentano, *Psychology from an Empirical Standpoint*, A.C. Rancurello, trans., Humanities Press, 1874/1973.
31. J.F. Yates, *Decision Management*, Jossey-Bass, 2003.
32. S.M. Rostan, "Problem Finding, Problem Solving, and Cognitive Controls: An Empirical Investigation of Critically Acclaimed Productivity," *Creativity Research J.*, vol. 7, 1994, pp. 97–110.
33. G. Klein, *Intuition at Work: Why Developing Your Gut Instincts Will Make You Better at What You Do*, Doubleday, 2003.
34. B. Schwartz et al., "Maximizing versus Satisficing: Happiness Is a Matter of Choice," *J. Personality and Social Psychology*, vol. 83, no. 5, 2002, pp. 1178–1197.
35. R.L. Keeney and H. Raiffa, *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*, John Wiley, 1976.
36. P.J. Feltovich, R.R. Hoffman, and D. Woods, "Keeping It Too Simple: How the Reductive Tendency Affects Cognitive Engineering," *IEEE Intelligent Systems*, vol. 19, no. 3, 2004, pp. 90–95.
37. E. Shafir, I. Simonson, and A. Tversky, "Reason-Based Choice," *Cognition*, vol. 49, 1993, pp. 11–36.
38. J.L. Fix, "Memos Key in \$4.9-Billion Verdict," *Detroit Free Press*, 13 July 1999, p. A1.
39. U.M. Dholakia and R.P. Bagozzi, "Mustering Motivation to Enact Decisions: How Decision Process Characteristics Influence Goal Realization," *J. Behavioral Decision Making*, vol. 15, no. 3, 2002, pp. 167–188.
40. L.R. Beach, "Epistemic Strategies: Causal Thinking in Expert and Nonexpert Judgment," *Expertise and Decision Support*, G. Wright and F. Bolger, eds., Plenum, 1992, pp. 107–127.
41. B.M. Moon and R.R. Hoffman, "How Might 'Transformational' Technologies and Concepts Be Barriers to Sensemaking in Intelligence Analysis?" presentation at the 7th Int'l Conf. Naturalistic Decision Making, 2005; available as a report from the Inst. for Human and Machine Cognition, www.ihmc.us.
42. N.A. Bernstein, "On Dexterity and Its Development," *Dexterity and Its Development*, M.L. Latash and M.T. Turvey, eds., Lawrence Erlbaum, 1996, pp. 1–246.
43. R.J. Jagacinski and J.M. Flach, *Control Theory for Humans: Quantitative Approaches to Modeling Performance*, Lawrence Erlbaum, 2003.

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

Who sets computer industry standards?

802.11

firewire

gigabit Ethernet

Together with the IEEE Computer Society, **you do.**

Join a standards working group at www.computer.org/standards/