Position Paper for NSF Extreme Events Decision Making Workshop

Alex Kirlik, Georgia Institute of Technology

My research involves understanding and supporting the judgment and decision making of experienced professionals in complex sociotechnical systems. Examples we have studied range from a combat information center of an Aegis cruiser, to a commercial airline flight deck, to restaurant short order cooking. Cases such as these may provide a stepping stone toward the study of extreme events, due to their dynamism, interactivity, ambiguity and high stakes. I will summarize a few major points that my cognitive engineering colleagues and I have learned about researching these domains.

Get your hands dirty

Few if any cognitive engineers take decision research done in a domain such as aviation or health care seriously anymore unless the work was done with the benefit of substantial domain expertise. The most productive research efforts are conducted by problem-oriented teams, comprised of a mix of domain professionals and behavioral or decision scientists or cognitive engineers.

Decisions happen IN the domain

Hilary Putnam once chastised his fellow philosophers of science for doing meta-mathematics when he thought that they should be doing mathematics (of physical theory for example). Putnam noted that philosophical issues were being expressed and moves being made directly within the language of the scientific domain, and philosophers who worked as a level above the domain were missing out on the key action. Similar comments can and should be made about some judgment and decision research, or at least that research that aims for direct practical relevance. The more successful JDM efforts in sociotechnical systems tend to be those in which the decisions studied and supported are actual domain (i.e., medical, health care) decisions and not decisions that stand above the domain in some way (meta-decisions). The relationship between decision researcher and domain expert should not be hierarchical but instead complementary. One must resist the tendency to let the blunt tools of decision analysis guide one toward formulating the decision problem in a more abstract manner than is ever really

experienced by domain practitioners. When this happens, there is a strong likelihood that the research is moving toward irrelevance.

Work with domain models

A decision tree is a very blunt tool for describing many of the environmental structures to which professionals in sociotechnical systems are adapted. Asaf Degani, Mike Shafto and I did a study of the decision making associated with pilots' use of the autoflight system in commercial aircraft, logging more than 100 hours in the cockpit jumpseat. To make sense of these data we soon realized that we required a detailed model of the autoflight control system itself, in order to describe exactly what pilots were being asked to adapt to and what problems they were confronting in doing so. The modeling framework we eventually adopted was taken from the control engineers involved in designing the automation rather than from decision research. Once we had a good representation of the controlled system (which was complex -- there are dozens of ways of configuring an autoflight system in the aircraft studied) we were able to describe the structure of pilot decision data in a relatively simple manner. We used canonical correlation to identify preferred configurations of the autoflight system as a function of clusters of system state variables, obtaining a relatively clean mapping of pilots' decision policies in one pass of the analysis. If one did not have a good representation of the global environmental structure that was being locally reflected in pilot behavior I suspect one might end up trying to create a piecewise approximation to these decision policies by presenting pilots with dozens or perhaps hundreds of forced-choice decision scenarios.

Preparation and Deliberation

The unit of analysis for studying the decision making of experienced professionals must span the boundary between preparation and deliberation. There is a natural tendency for decision research to focus more on the latter. However, many organizations try to handle the lion's share of their decision problems by identifying procedures, response plans, rules of engagement, and the like. These organizations recognize that off-line decision making is sometimes preferred as it is sheltered from time pressure and various other stressors, even if it sacrifices access to some information. The design of a decision methodology for handling a complex decision scenario must address the question of how much, or which aspects of, decision making should be

"precomputed" and what should be instead done "on line." Many organizations, of course, have a strong preference for precomputation, and an important research question would seem to be how to make decision makers aware of the boundary conditions and any other limitations of their rules and standard operating procedures.

Dynamism

A decision not to act in a dynamic environment is no more passive than a decision to act, as it merely represents a preference for one particular (default, as it were) evolution of the system over another. Ideally, the decision maker should be concerned solely with achieving a desirable evolution of the system and not the question of how much in the way of action is necessary to achieve it. Yet, individuals and organizations can differ strongly in their attitudes toward what they somewhat inappropriately conceive of as "inaction versus action." Any systematic tendencies for one over the other therefore have the potential of restricting the possible evolution of the system in ways that may be inconsistent with the individual's or organization's goals. There are of course many other issues that arise in dynamic contexts.

The multiplicity of action

In contexts such as the operating room, the flight deck, and the combat information center, we have been struck by the frequency with which action is taken for reasons other than maximizing some kind of utility. Some actions of course are taken in an effort to sample information to reduce uncertainty. Other actions, however, have a cognitive rather than performatory role that is less direct. Some actions seem to be taken for the purpose of trying out a reduced version of a potential decision solution. Others seem to be taken to keep the evolution of the system within a restricted portion of the state space in which the decision organization has past experience. Others are taken merely to determine if they have any effects at all, and to identify what these might be. Modeling tools enabling a richer description of the various roles of action are needed.

Suggested reference

Kirlik and Bisantz, (1999), Cognition in human-machine systems: Experiential and environmental aspects of adaptation. In P.A. Hancock (Ed.), *Handbook of Perception and Cognition: Human Performance and Ergonomics*. New York: Academic Press.