

Adaptive Mixed-Initiative Systems for Decision-Theoretic Crisis Response

Melinda T. Gervasio

Institute for the Study of Learning and Expertise
2164 Staunton Court, Palo Alto, California 94306
gervasio@isle.org

A *crisis* is characterized by three elements: threat, urgency, and uncertainty (Rosenthal et al. 1989). In a crisis, an agent risks losing something of great value. The agent must take action quickly because without intervention, this undesirable situation will be realized in a limited amount of time. Different actions will generally incur different costs and have different outcomes under different circumstances. The agent's decision-making process is further complicated by various sources of uncertainty—noisy information, exogenous events, actions with nondeterministic outcomes, etc. Information-gathering actions may reduce uncertainty, but they incur additional costs, not the least of which is the time so ill-afforded in crisis domains.

Consider an incident involving the spill of some hazardous material in a factory. Unabated, the spill may damage the floors of the building, release toxic fumes that threaten the health of the factory workers, and seep into waterways to pollute the environment. Rapid response is necessary to contain the spill and minimize its harmful effects. Because of limited resources, the response team will not be able to completely avoid all the negative consequences. Their task is complicated by uncertainty—incomplete information about the involved material, unpredictable weather, inaccurate information about the building and its occupants, and so on. Gathering more information may reduce this uncertainty, but in the time this takes, the magnitude and the scope of the problem presented by the spill will continue to increase.

Decision theory seems a natural approach, given the characteristics of crisis. However, decision-theoretic techniques are particularly difficult to implement for crisis domains for the following reasons. Decision-theoretic methods require the development of complete models, a time-consuming task ill-afforded in crisis domains with their characteristic of urgency. The limited amount of time for decision-making also constrains the development of alternative courses of action and their subsequent consideration. Furthermore, the un-

certainty inherent in crisis domains makes the enumeration of alternatives as well as the elicitation of probability models more difficult.

However, decision-theoretic methods remain attractive for crisis domains because they provide well-founded principles for making critical decisions. They force an agent to explicitly enumerate the values of importance, the alternative courses of action, and the uncertainty of the available information. This allows rational decision-making with clear justifications in high-stress and often volatile situations.

The key for crisis response is to be able to rapidly develop adequate decision models and use them efficiently. Towards this objective, we propose the use of machine learning techniques to learn user preferences¹ in crisis response. By learning to correctly anticipate a user's decisions, such an *adaptive* assistant can facilitate more rapid crisis response.

The planning process involved during crisis response—and indeed any decision-making process—may be viewed as a search through a space of possible solutions for one that meets particular criteria. Machine learning techniques have been applied in automated planning to speed up this search by learning macro-operators or control rules that result in following more correct and shorter paths to acceptable solutions (e.g., Fikes et al. 1972; Minton 1988). When the criteria for acceptability become difficult to specify or when they become user-dependent, learning can also be used to acquire user-specific rules that lead to preferred solutions more quickly (e.g., Langley 1997). Decision-theoretic planning presents additional opportunities for the application of machine learning to achieve similar efficiency gains.

For example, an early step in decision analysis is the development of a decision basis—the enumeration of

¹ *Preferences* in the general sense, not only the decision-theoretic sense of relative valuation of possible outcomes.

alternative actions, sources of information, and possible outcomes. Consider the situation where the computer assistant presents the human user with an initial guess as to the relevant factors and the user then interacts with the assistant to add new elements and delete undesirable ones. The user presumably makes these modifications according to a certain, though often unstated, set of criteria reflecting the user's judgments regarding the most important factors to consider in the decision. If the assistant can extract these preferences from its interactions with the user, then in future problems the assistant should be able to present better initial candidate solutions that more closely satisfy the user's criteria, thereby reducing the modifications the user needs to make.

Similarly, in the development of probability and utility models, there may exist some initial set of default values that the user can manipulate to construct an appropriate model. The assistant can help by predicting the adjustments the user will make and suggesting those to the user. By learning a more accurate user model, the assistant should be able to make suggestions more in tune with the user's individual preferences. For example, the assistant may learn that the user tends to be more confident about particular sources of information, and that the user is always concerned about particular outcomes but is indifferent about others. This knowledge allows the assistant to make suggested modifications the user is more likely to accept, again speeding up the decision-making process.

For a machine learning approach to be effective, however, certain assumptions must hold. First, there must be generalizable (i.e., learnable) aspects of the user's behavior across problems. If every problem has a unique decision basis, or if users behave erratically, learning will be unlikely to help. Second, the users must solve a sufficient number of problems to allow the assistant to learn. The greater the interaction between the assistant and the user, the better the assistant's chances at developing a good user model. Third, the preferences being learned should not be easily obtainable from the users. If they are, then there is no point in expending the effort to learn them, approximately at best.

While crises such as the Cuban Missile Crisis, the Tylenol scare, the collapse of the Asian economy, as well as many personal crises may be outside the realm of this approach, there are many interesting, real-world domains with enough of the necessary structure to make this approach viable. For example, in the hazardous materials domain, the available courses of action are well-defined, the set of factors that can affect an incident is known, and there exist basic guidelines

for developing an appropriate response (NAERG96). Adaptive assistants can facilitate the development of specific responses tailored to particular response teams and within the boundaries of these constraints and guidelines. For example, they can learn that weather information from yesterday's newspaper is sufficiently reliable, that if a chemical factory is involved it will provide additional resources in the form of equipment and experts, and that a medical emergency team should always be put on standby. Military operations and disaster relief operations are similarly governed by standard operating procedures that make the use of adaptive assistants feasible.

With the realization that humans are often irreplaceable in the decision-making process, not to mention also reluctant to relinquish complete control to a machine, there has been a trend away from autonomous AI approaches toward mixed-initiative systems that allow humans to retain control of the problem-solving process. In the proposed mixed-initiative system for crisis response, control of the decision-making process lies ultimately with the human user. The role of the adaptive assistant is to facilitate the user in the decision-making process by modifying its behavior through the user-specific preferences it learns from its interactions with the user.

Acknowledgments This research is supported by the Office of Naval Research under Grant N000014-96-1-1221, and benefited from discussions with Wayne Iba, Pat Langley, and the Organizational Dynamics Center Group at Stanford University.

References

- Fikes, R. E.; Hart P. E.; and Nilsson, N. J. 1972. Learning and Executing Generalized Robot Plans. *Artificial Intelligence* 3(4): 251-288.
- Langley, P. 1997. Machine Learning for Adaptive User Interfaces. In *Proceedings of the Twenty-First German Annual Conference on Artificial Intelligence*.
- Minton, S. 1988. *Learning Search Control Knowledge: An Explanation-Based Approach*. Norwell, MA: Kluwer Academic.
- Rosenthal, U.; Charles, M.; 't Hart, P.; Kouzmin, A.; and Jarman, A. Case studies to theory and recommendations: A concluding analysis. In Rosenthal, U.; Charles, M.; and 't Hart, P. eds. 1989. *Coping with Crises*.
- Transport Canada, the U.S. Department of Transportation, and the Secretariat of Communications and Transportation of Mexico. 1996 *North American Emergency Response Guidebook* (NAERG96).