ASSESSING EMERGENCY INTERFACE DESIGN

William C. Allen
Stanford University

INTRODUCTION

A functional model of beam dynamics has a theoretical basis in elasticity and mechanics. Similarly, a functional model of man/machine interaction must have an underlying theoretical basis for how the man and the machine respond to their environments. In most circumstances, control theory provides an adequate theoretical model for the mechanical system. Efforts to adapt this theory to the human operator have encountered severe limitations -- especially in unfamiliar or multi-attribute control environments (i.e., emergencies or strategic level decision making).

If we designed buildings using an elastieity model valid over a very narrow region, we would be careful not to employ this model as a predictor of system response outside that region (otherwise our building might fall down). We can criticize almost all existing interface design procedures because they assume (usually implicitly) that an optimal control model for operator response is valid over the entire range of control environments. Therefore, it should not be surprising that performance of these systems degrades seriously when exposed to situations where this assumption is invalid.

An important byproduct of this realization is that if an explicit model of human response is not incorporated into a design procedure, an implicit model will take its place. Clearly, something as critical as operator response should not be left to default modelling.

Cognitive psyehologists have developed a general theory for human information perception, storage, retrieval and manipulation based on the notion of schemata. One difficulty with modelling human response using schema theory is that there is little understanding of the underlying meehanisms of schema manipulation. However, in the same way that it is not necessary to understand the nuances of molecular interactions to make predictive theories about materials behavior, it can be hoped that we can address human response tendancies based on a macro level schema theory.

The next section highlights the fundamentals of schema theory necessary for this discussion. The third section addresses how the process of schema selection optimization appears to represent a reasonable descriptive model of operator response. In the fourth section, this model is expanded to provide an analytical approach to assessing interface design.

The result is a structured apprsaeh to addressing human decision making/action selection over a wide range of operating environments. Simulation can be used to illuminate the different paths the inferential decision maker can take given a particular display/control format and a given set of status/warning indications.

It is hoped that future work in this area will allow this modelling/simulation/analysis process to be codified. The result would be a general design procedure for emergency interfaces.

Introduction to Schemata

By constraining system design to include a "man in the loop", we are forced to recognize and address the issues of non-normative decision making typical of human response. Schema theory provides a robust mechanism through which we can address human information processing. Three structures must exist for schema theory to function:

1. Semantic knowledge structures. These are semantic networks between key variables in a frequently encountered phenomenon.

2. Episodic memory structures. The key features in an experience (as well as connectors between the features and the moderating schemata) are retained as episodic memory traces.

3. Schema selection optimization. Two levels of optimization must occur:

   --The relative costs of delays and potential perception errors must be rapidly balanced when selecting the schema to be used to structure exogenous data.

   --The costs of validating and determining the uniqueness of an "apparently valid" schema must be optimized before it is used to generate response scenarios.

Semantic and episodic knowledge structures are highly interactive and mutually supportive. Semantic memory provides the framework used to encode and recall episodic memory traces. Episodic memory traces provide experiential richness to a partially instantiated schema by allowing experience with "similar" or analogous situations to be brought into the evaluation.

With this approach, we would prediet that an operator would deal with limited sensory data by reaching into his episodic memory stores for "representative" values and constraints for unavailable variables. This ability to draw from past experiences provides added "context" when evaluating a situation. Conversely, faded episodic memory traces can be "reconstructed" using the implied contexts and constraints available through the semantic knowledge structure.

Optimizing Schema Selection

Limitations in short term memory dictate that people do not simply load all available sensory data into a memory buffer and perform an exhaustive sort of schemata to determine the best fit. Instead, we expect that a "reasonably valid" schema is rapidly chosen and used as a framework to organize the sensory data. If a selected schema fails to provide an adequate fit for subsequent sensory data, it is rejected in favor of a new schema.

From this, we would expect that inhibition of invalid schemata must play an important role in the initial schema selection process. Otherwise, we would be bombarded with potentially invalid schemata and have to eonciously sort through them all. Therefore, the attended features which appear central (and their apparent semantic connections) must serve to inhibit all schemata whose variables or structures do not permit such features.

Since this sort takes place very rapidly, a hierarchy of the variables in a schema must exist. This initial sort must compare the "central" sensory features to schemata with similar central variable constraints.

Based on these (and other) considerations, a "branch and bound" optimization strategy appears to be an appropriate model for the initial schema selection optimization proeess. This approach allows'schema selection to be dependent on both previously attended information and prior decision path. Experiential and analytical heuristics are used to estimate the schema which "appears" to contain the greatest potential for providing a solution. Other heuristics are used to estimate the uniqueness of the solutions resulting from the use of that schema.

Before modifying the branch and bound approach to accommodate schema optimization, it is worthwhile to review the steps in the general algorithm:

1. Partition the solution place into mutually exclusive, collectively exhaustive sub-spaces.

2. Develop upper bounds (assuming a maximization problem) for each of the sub-spaces.

3. Test each upper bound solution to determine if it is feasible.

4. If one (or more) feasible solution exists, make the highest of these the incumbent solution.

5. If any of the following are true, the sub-space is considered fathomed:

   -- subspaces whose upper bound solution is also feasible
   -- sub-spaces with upper'bounds lower than the incumbent solution
   -- sub-spaces containing no feasible solutions

   Fathomed sub-spaces can be removed from further consideration.

6. Take either the highest (the most promising) or the most-recent of the unfathomed sub-spaces (depending on the branching rule employed) and partition it into smaller sub-spaces.

7. Go back to step 2 and continue until all subspaces are fathomed-- the incumbent solution at this point will be an optimal solution.

The schema selection algorithm would follow the sane general pattern:

1. The "solution space" is partitioned by activating schemata with appropriately eonstrained "central" variables and inhibiting inappropriate schemata. Within these overlapping sub-spaces simple cause-effect relationships exist between focal and peripheral components or sub-systems.

2. Constraints provided by episodic memory are used to provide "quick access" bounds on the "goodness of fit" of all active'schemata.

3. The best "quick fit" would be selected as a candidate schema. Using available features to partially instantiate the candidate would form a test of it's feasibility (i.e., if structural or variable constraints were violated, the schema would be fathomed and removed from consideration.

4. The "stopping rule" for the schema search algorithm would be based on either the difference between the "goodness of fit" of the partially instantiated candidate schema and other active schemata or the difference between the episodic associations activated by the partially instantiated candidate and those activated by the available features (i.e., improved "goodness of fit" of the partially instantiated schema inhibits activation of alternate schemata while poor "goodness of fit" increases their aetivation).

When applied to schema selection, branching difficulty can be considered the number of hypotheses which are "candidates" for interpreting the data. If only one hypothesis receives activation from the available (or attended) features, little or no cognitive effort is required. These unambiguous hypothesis selections should appear "automatic" (i.e., require no cognitive effort).

Conversely, if multiple hypotheses(or no hypothesis) receive activation, conscious hypothesis selection must occur since some means must exist for organizing the available data. For example, if conscious hypothesis selection does not take place in an ambiguous situation, the operator will be unable to make "sense" of it and will take no action.

This discussion implies that ambiguous situations will be predominated by "conceptually driven processing" while unambiguous situations will be predominated by "data driven processing".' This points out a significant difference between designing routine and emergency interfaces:

-- Routine situations can be effectively controlled by providing adequate sensory input (data driven processing predominates).
-- Emergency situations require simulation of appropriate perceptions of the situation to allow adequate response (conceptually driven processing predominates).

Since initial schema selection is used to encode available data, while action selection involves an irrevocable allocation of resources, we would expect additional schema optimization testing before action is allowed. Although the incumbent schema can be considered reasonable and feasible after it's acceptance as a sensory data framework, it's uniqueness remains to be resolved.

In unambiguous situations, this fathoming of alternate candidates is trival-- there are no alternate candidates. Therefore, we can expect that these situations can lead to "slips" when relatively rare alternatives are the "correct" choice. For example, if an indicator "always" implies a particular action during normal system operation, we can expect "intrusions" of that action during emergency operations where the implications of that indicator may be more ambiguous.

This approach asserts that the schema will remain unchanged unless contradictory evidence is found. Since evidence is asserted to be requested for the purpose of confkrming the current hypothesis, we can expect that an incorrect hypothesis formed on the basis of weak evidence will be more resistant to change based on new, better data than a hypothesis formed on the basis of the new data alone. For example, when critical data is unavailable, several schemata might be selected with equal validity. Exposure to a few non-critical features contradietory to the intial schema selection will likely result in the operator "explaining away" the new features rather than changing schemata,

Assessing Interface Design

Although schema theory provides valuable insights, it lacks the computational structure necessary for quantitative analysis. The branch and bound model for schema manipulation provides the needed structure but does not include a systematic means for handling uncertainty in our knewledge of the interactions between model parameters or uncertain outcomes.

This shortcoming can be resolved (to some extent) by adopting a modified decision analysis approach to interface design assessment. Since the structural models used in decision analysis share with branch and bound a sequential tree-like structure, the models should be compatible.

Figure 1 depicts a "status annunciating" display configuration for a simple oil transport/storage system. Thfs system assumes that the oil source (i.e., tankers) is to be unloaded as quickly as possible and that excess pumping capacity is diverted to temporary oil storage tanks. The valves can be open or closed, pressures can be zero, low, medium or high and the pump can be set at off, slow or fast. If problems develop, they will be annunciated by the lights at the bottom of the display.

Figure 2 is the abbreviated decision tree which represents a simplified structural. model for an inferential decision maker. Note that the structural model follows the same general format outlined in the discussion of schema theory but excludes all functional details. Figure 3 is a truncated decision tree which results when this model is applied to the example display format.

The dependence of decision variables on attended information and the heuristic nature of the hypothesis formation process causes the "decision" variables to retain a probabilistic nature. The probabflity associated with the selection of a candidate "decision" depends on the "degree of association" existing between the candidate and the hypothesis given a particular generating scenario. This approach is consistent with the aetivation model for schema moderated behavior presented in the previous section.

At this level of sophistication in the inferential decision model, the assignment of probabilities to the branches of a decision node would be entirely subjective. "Reasonable" values for these probabilities would be derived from experience with similar systems, results of pencil and paper or simulator studies, or from "engineering judgement".

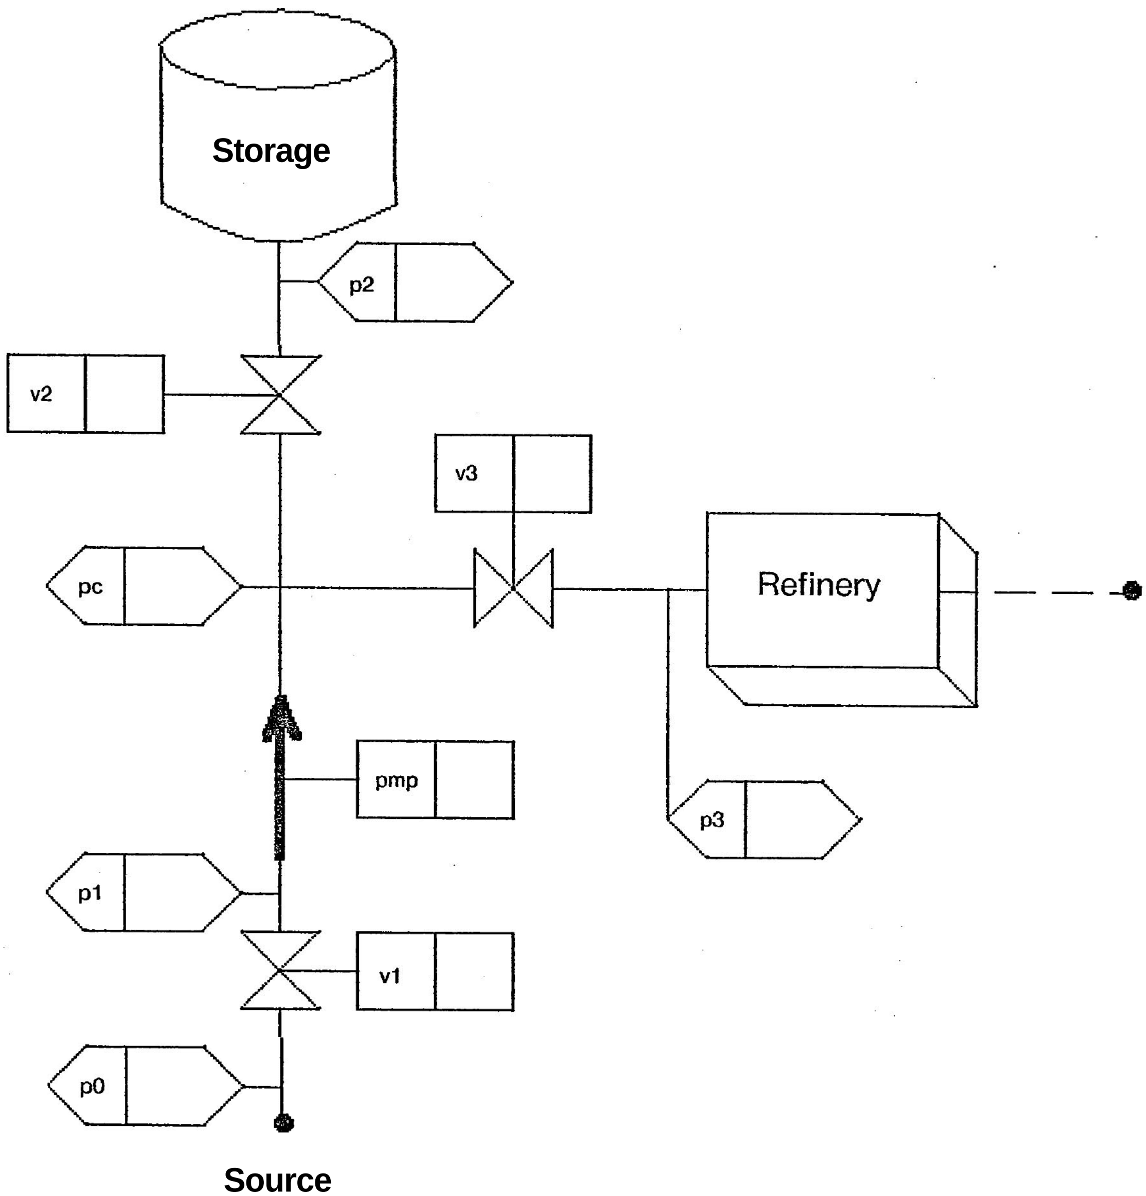

WARNINGS

Source Disruption	Pump Breaker Tripped	Excessive Tank Pressure	Insufficient Production Pressure

Status Warning Annunciator Display

Figure 1

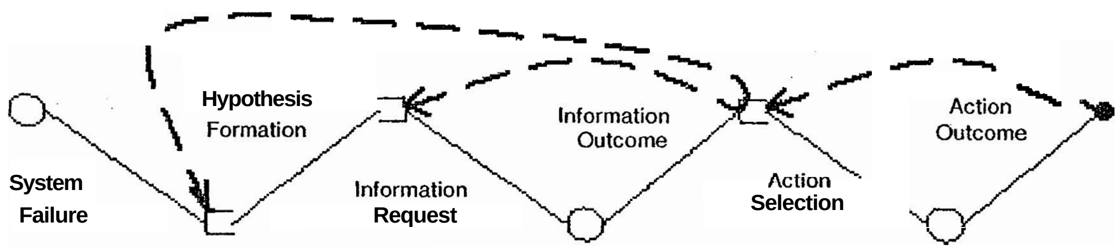

Abbreviated Inferential Decision Tree

Figure 2

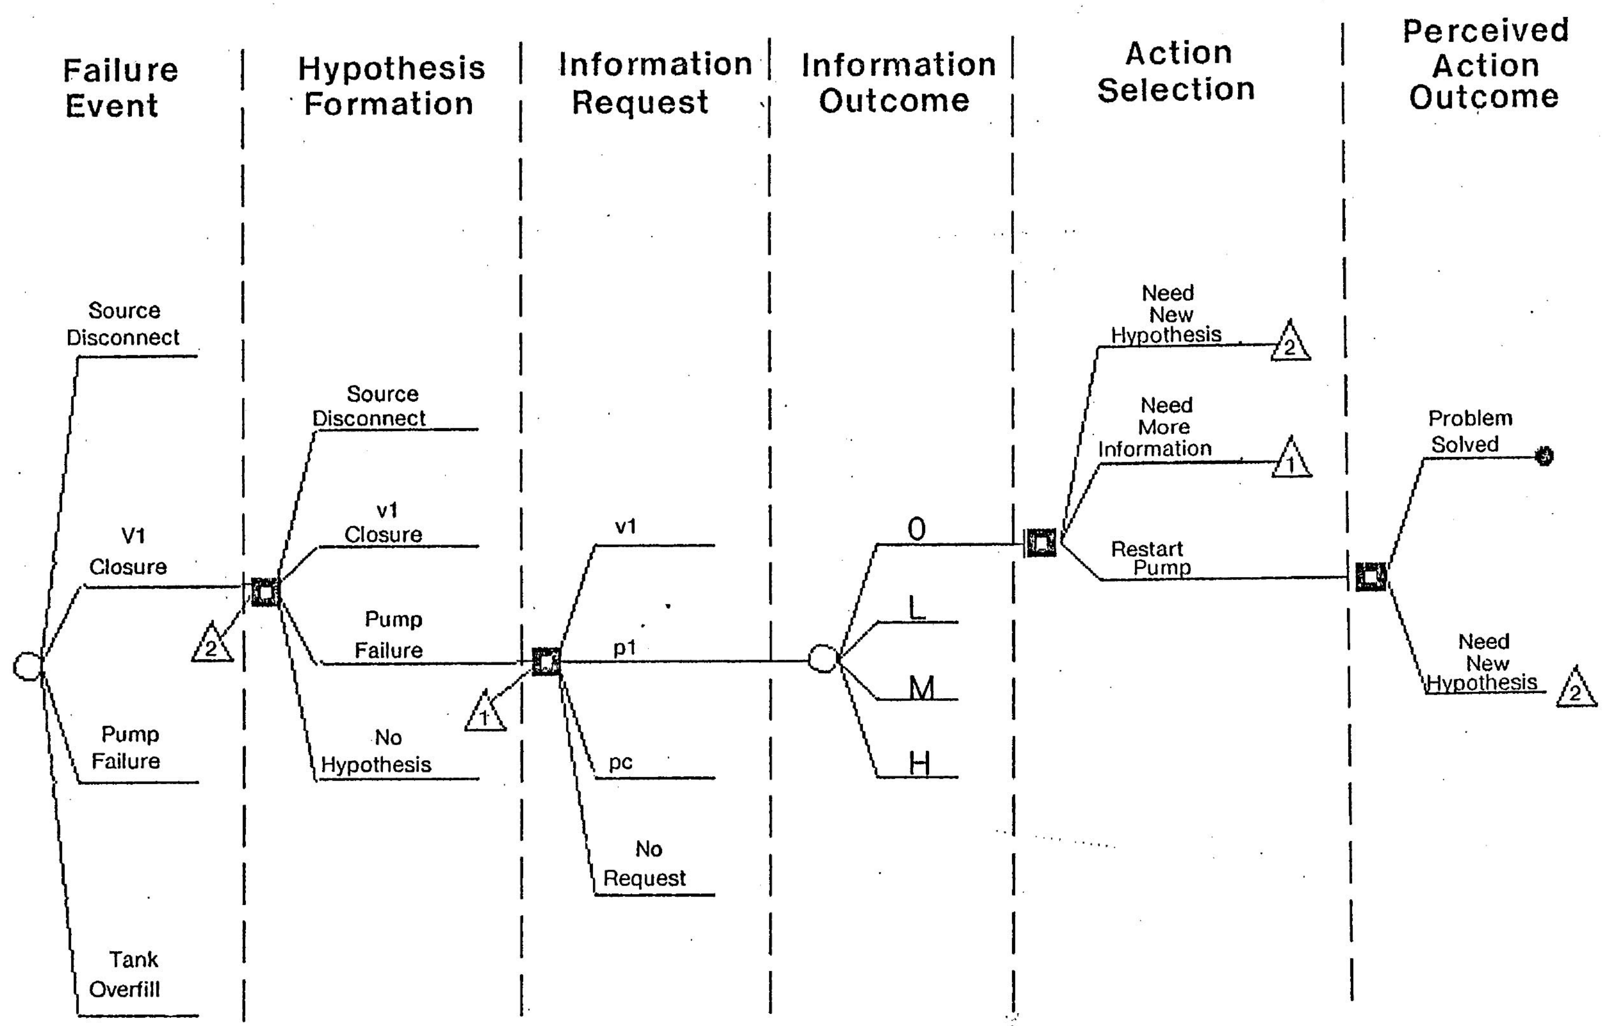

Truncated Inferential Decision Tree
Figure 3

Additional detail can be incorporated into the model by including an "attended features" branch prior to schema selection (as schema theory would dictate). In this way, probability assignments for schema selection would be conditional on the attended features and/or the contents of short-term memory immediately prior to schema selection.

Although this would reduce the subjectivity of the schema selection probability estimates, it greatly increases the computational difficulty of the problem by necessitating the inclusion of a significant number of "attended feature" branches into the model(the number of branches would be approximately the number of observable features chosen the size of working memory at a time -- a huge number for any realistic situation.

To a large extent, this difficulty can be overcome by eliminating trivial or redundant branches. However, a more likely solution would be expected to lie in adopting a simulation approach to enumerating the effects of various attended feature combinations. Regardless, a great deal of analysis would be required to assess the impact of various attended feature combinations on schema selection.

In the simple model presented, action selection and information request probability estimates would be based on experience, intuition or the results of simple experiments. However, an expanded model could include conditional branches to account for the predictable effects of requested information on schema activation and inhibition. Additional conditional branches could be included to account for the effects of differing schema fathoming or stopping rule strategies on information requests and action selection.

For example, time pressure, experience, environmental distractions and motivation can all have a reasonably predictable influence on these parameters. Conditional branches can be included to allow for these variations. As system reliability is allowed to degrade, the dynamic nature of the task is allowed to increase or time sharing actigity is increased, additional conditional branches would have to be included to account for the effects.

It is important to make a distinction between computational and theoretical modelling difficulty. The advantage of a structural model is that explicit assumptions are made about the path taken to reach a particular decision or event. The constraints imposed by these preconditions makes it much easier to predict what will happen at that particular point in time. The difficulty lies in the computational burden of going through all the possibilities and estimating what will happen at each.

Therefore, this approach reduces the theoretical complexity of the analysis at the cost of increased computational difficulty. Fortunately, it is far easier to develop computational short-cuts and streamlined algorithm's to solve these computationally more complex problems than it is to develop a comprehensive theory about human information processing.

For example, a progressive model building/sensitivity analysis approach might allow significant ''pruning'' of redundant or trivial branches before they have to be explicitly evaluated. This would be accomplished by developing a simple structural model, making "ball park" estimates for the possible branches and their probabilities and performing a sensitivity analysis to determine which branches appear to be most important. Less important variables would be set at "nominal values" while the complexity of the model would be increased for the critical variables. Oddly enough, this is exactly the process that schema theory predicts takes place in human information processing.

Simulation can be used to enumerate decision and event chains which present a significant hazard (hazard is defined as the probability of a decision/event chain times the "cost" of the outcome). Additional modelling effort can be given to addressing the interaction of the variables in decisionlevent chains representing the greatest hazard.

At the very least, this approach has the benefit of explicitly enumerating the man/machine/environment interaction assumptions necessary to design an interface. The combination of formalized structure and simulation aids the designer to uncover unintuitive'or insideous sources of operator or system error. Knowledge of biases in human information processing can be used to ferret out potential "slips" or inferential errors which would·mot 'be illustrated by an optimal control or sequential analysis approach. Further, the structure of this approach makes it less likely that a designer will assume his own response biases in assessing potential operator response.

The major benefit of an approach of this kind is it's usefulness in overcoming one of the greatest difficulties (and dangers) in a priori caution and warning system design -- in order for problems to be clearly and unambiguously annunciated, the possibility of operator uncertainty at all potential decision points must be foreseen by the designer.

Prospects for the Future

It is entirely possible that cross-situational consistency in certain categories of response can be uncovered by a structured approach sueh as the one proposed. If these consistencies can be codified into either a simulation or rule based system, it may be possible to evaluate interfaces in a fraction of the time necessary to "start from scratch".

Further, if a system of evaluation can be codified, it can be used to optimize the design of the interface. This presents the possibility of elevating interface design from a "satisfycing" approach to one where the designer can quantitatively estimate how far the proposed system deviates from "optimal". This would allow him to perform "value of information" studies to judge how much additional expense or research effort is warranted in attempting to improve the design.