

# Novel Computational Techniques for Incident Reporting

From: AAAI Technical Report WS-00-03. Compilation copyright © 2000, AAAI (www.aaai.org). All rights reserved.

Chris Johnson

Department of Computing Science,  
University of Glasgow,  
Glasgow, G12 8QQ, Scotland.

[johnson@dcs.gla.ac.uk](mailto:johnson@dcs.gla.ac.uk), <http://www.dcs.gla.ac.uk/~johnson>

## Abstract

Incident reporting systems help users to provide information about potential safety hazards. They, therefore, represent an important subset of the wider range of applications that support process improvement. The following pages identify a range of novel computational techniques that can be used to address problems of existing reporting systems. In particular, it is argued that computer-assisted interviewing techniques, such as the familiar frame and script approaches, can guide the elicitation of incident reports. Probabilistic information retrieval systems reduce the classification problems that prevent attempts to index diverse reports in dynamic industries. Conversational case-based reasoning techniques can be used to avoid the problems of query formation that frustrate attempts to retrieve similar incidents. Finally, discourse-modeling techniques can be extended to represent the reasons why particular lessons have been learnt from particular incidents.

## A Brief Introduction to Incident Reporting

Incident reporting applications are an instance of the more general class of systems for organizational and inter-organizational learning (Weber et al, 1999). For example, the FAA's Aviation Safety Reporting System (ASRS) was set up to provide a confidential means of learning about aviation related incidents in the United States. It, therefore, provides a primary source of information about the causes of failure. This information has been used to generate Operational Bulletins that *alert* the aviation industry to "timely and important" concerns in recent submissions. The ASRS also helps to disseminate *best practice*. The lessons learned from previous failures are disseminated through the Callback publication. Approximately 85,000 copies of this are distributed to pilots, engineers, air traffic controllers etc. However, a number of problems limit the utility of such incident reporting systems.

1. It can be difficult to elicit information about previous incidents from the users that are involved in them. Many incident reporting forms provide a cursory overview of the events leading to failure and so

investigators have to visit contributors to identify missing information. This creates considerable logistical problems for the growing numbers of national and international reporting systems.

2. It can be difficult to correctly index and classify incident reports so that others can perform the statistical analyses that help to guide subsequent intervention. For example, there is a growing concern over the problems of Crew Resource Management (CRM) in many aviation incident reports (Johnson, 2000a). Unfortunately, incident reports are not routinely indexed in terms of specific problems such as CRM. As a result, analysts cannot query systems such as the ASRS to retrieve every CRM related incident over the last five years (Johnson, 2000b).
3. It can be difficult to correctly issue the queries that are needed to retrieve information about particular incidents. The importance of correct query formation is illustrated by the fact that the ASRS now has a cumulative total of more than 500,000 reports. It is impossible for individuals to manually search such collections to find the lessons that apply to their systems. It is also difficult for organizations to spot emerging trends amongst the mass of data that has been collected.
4. There is a danger that rather than learning the lessons of the past, organizations will simply use incident reports to find evidence that supports their existing preconceptions and biases (Johnson, 2000b). It is, therefore, important to explain why particular lessons can be drawn from particular incidents.

The following sections briefly describe a number of computational techniques that can be applied to avoid or mitigate the impact of these problems for incident reporting systems.

## **Problems of Eliciting Incident Information: Computer-assisted Interviewing**

The problems of eliciting information about previous incidents should not be underestimated. At present many systems rely upon confidential rather than anonymous reporting. For instance, the UK CIRAS rail reporting system sends an investigator out to conduct a follow-up interview in response to every report form that is submitted. Similarly, NASA personnel go back to the contributors of many ASRS submissions. This approach requires considerable resources. There must be enough trained analysts to elicit the necessary information during follow-up visits. Alternatively, it might be possible to recruit novel computational techniques to improve the quality of information that is initially contributed in response to an incident. These techniques might, therefore, reduce the expense associated with site visits. Equally importantly, they might also avoid the biases that affect follow-up interviews. A number of social concerns must affect contributors during safety-related discussions with external interviewers. Eliciting more information in the immediate aftermath of an incident also helps to reduce any delay between the contribution of a report and a follow-up interview.

The problems of extracting information from domain experts has been addressed by work on knowledge elicitation in general and by computer-aided interviewing techniques in particular (Saris, 1991). These interviewing techniques, typically, rely upon frames and scripts that are selected in response to information from the user. For example, the user of an air traffic management system might first be prompted to provide information about the stage of flight in which an incident occurred. If it happened during landing then a script associated with that stage of flight would be selected. This might provide further prompts about the activities of arrivals and departures officers or about specific items of equipment, such as MSAW protection. These detailed questions would not be appropriate for incidents during other stages of flight, such as those filed during en route operations.

The relatively simple script-based techniques, described above, offer a number of further benefits. In particular, the use of computer-assisted interviewing can reduce the biases that stem from the different approaches that are used by many interviewers. Inter-analyst reliability is a continuing concern in many incident report systems (Johnson, 2000b). The scripts embodied in computer-assisted interviewing systems might also be tailored to elicit particular information about regulatory concerns. For instance, if previous accidents had indicated growing problems with workload distribution during certain team-based activities then scripts could be devised to specifically elicit information about these potential problems. Of

course, this analysis must be balanced against the obvious limitations of computer-based interviewing techniques (Saris, 1991). Further evidence is needed to determine whether the weaknesses of computers assisted interviewing in employment selection or the analysis of consumer behavior also apply to their application in incident reporting. Until this evidence is provided then there will continue to be significant concerns about the problems of bias that can be introduced during the elicitation of information about previous failures.

## **Problems of Indexing Dynamic Incidents: Probabilistic Information Retrieval**

A number of problems remain to be addressed once the information about an incident has been gathered. Perhaps the most important of these relates to the indexing of large-scale collections. At present, most successful incident reporting systems rely upon relational database technology. Each incident is classified according to a number of pre-determined fields. Queries can then be constructed, using languages such as SQL, to sort, filter and combine incident data according to the information contained in these fields. This approach has a number of consequences. In particular, it can lead to an extremely static classification system because there is often no way to automatically reclassify thousands of previous incidents if changes are made to a taxonomy. For instance, many existing schemes use Reason's (1990) GEMS taxonomy of human error to classify operator behavior in the lead-up to an incident. This taxonomy has recently been revised in a number of ways. However, few of these changes have been reflected in incident reporting systems because of the costs associated with manually analyzing and re-classifying existing records. This has profound consequences. As mentioned earlier, analysts are faced with retaining distinctions that may no longer reflect the way in which particular tasks or activities are organized. Alternatively, the problems of updating previous records can result in only a small portion of the incidents having values for the most recent set of fields.

In information retrieval, the concept of poor recall is used to describe situations in which only a small proportion of relevant documents are returned from a collection in response to a users' query. Conversely, poor precision results in many irrelevant documents being returned as potential hits. These concepts have particular importance for incident reporting schemes. If the fields in a relational scheme are not updated then queries about new concepts will often result in poor precision. Users will have to construct queries from existing values that do not accurately describe the concepts or classifications that they are interested in. Conversely, adopting more dynamic classifications in which new fields will only be maintained for subsequent reports will lead to poor recall. A highly precise set of incidents can be returned, for example in

response to queries about crew resource management. However, these will only represent the subset of all incidents that were indexed using the new data-model. There will be many other reports that were not classified in this way and hence will not be retrieved.

Information retrieval tools, such as Bruce Croft's Inquiry (Turtle and Croft, 1991), address the problems of precision and recall, mentioned above. They offer a range of techniques for exploiting semantic information about the relationships between the terms/phrases that appear in a document and the terms/phrases that appear in the users' query. Dictionaries can be pre-compiled to support query expansion. They can also exploit probabilistic information based on the relative frequencies of key terms. Retrieval engines can rank documents according to whether or not documents are thought to be relevant to a query. Relevance feedback techniques build on this process of assigning probabilities. The key point is that the compilation of dictionaries and concept formation can be altered over time. As a result, analysts need not rely upon the static classification of incidents that will eventually yield poor precision or poor recall.

### **Problems of Complex Query Formation: Conversational Case-based Reasoning**

The problems of poor precision and recall do not simply relate to the techniques that are used to index incident reports. They also stem from the users' inability to correctly form a query that will return the anticipated results. For example, if the user were interested in finding information about CRM then a query on "workload AND stress" might achieve relatively good recall but poor precision. Many CRM-related would be returned, as would a large number of incidents involving other forms of stress. This might be perfectly acceptable if the user had planned to manually filter these additional reports. However, in many cases users are forced to perform this filtering not because they choose to do it but because they simply cannot refine a query to return their desired results. The problems of query formation are also compounded by the languages that users, typically, must exploit to search the relational systems that support most incident reporting schemes. It can be extremely difficult for even highly skilled analysts to correctly form the queries that are supported by languages such as SQL.

We have explored the use of case-based reasoning techniques as a means of avoiding some of the limitations of relational systems. In particular, the US Navy's Conversational Decision Aids Environment (NaCoDAE) (Aha, Breslow and Munoz-Avila, 2000) has been used to store and retrieve subsets of the ASRS collection (Johnson, 2000b). This system can be used to group similar cases according to the answers that users provide to a number of predetermined questions. The conversational style of this

interface helps to ensure that users never have to master the syntax of languages such as SQL. The user initially types a free text query. The system responds with a list of previous failures that are assumed to have matched this initial query. They are also presented with a list of questions that the system considers will help the user in refining their search. By answering these questions, the user not only affects the matching cases that are displayed but also helps to filter the list of further questions. The ability to guide users in this way is a significant consideration. Many of the engineers and managers who are best placed to apply particular lessons do not have the motivation or necessary training to learn the more advanced features of relational query formation.

An important benefit of the conversational approach to case-based reasoning is that users get a considerable amount of feedback about the precision of their interaction. The list of matched cases is updated each time they provide an answer about the incidents that they are looking for. Information is also displayed about the degree to which each case matches the answers that they have given. As a result, poor precision is indicated by a large number of cases with a low matching score. This encourages users to refine their search by answering additional questions. Poor recall is less easy to avoid with conversation case-based systems because the answers to each question tend to be hard-coded with each incident. As a result, they tend to suffer from some of the problems associated with the static classifications of relational systems. As a result, we are currently working to combine the flexibility of probabilistic information retrieval systems with the guided exploration supported by systems such as NaCoDAE.

### **Problems of Inter-Analyst Reliability: Computational Models of Argumentation**

Lekberg's (1997) work for the Swedish Nuclear Power Inspectorate illustrates the problems that arise in analyzing the contributions to incident reporting systems. She demonstrates that there are fundamental biases in the way that different experts analyze particular incidents. Previous training and expertise affect an engineer's interpretation of causal events. Inter-analyst biases have a significant effect on the conclusions that are drawn about particular incidents. This, in turn, has a significant impact on the lessons that are drawn from previous failures.

The problems of interpreting the causes of an incident can be approached in a number of ways. For example, classification rules can be used to guide a causal analysis. Investigators must apply these rules to select a small number of root causes from an approved taxonomy. Individual biases are reduced because the analysis procedures force analysts to explicitly consider a wide range of latent and catalytic factors, including human error,

system failure, managerial weakness, environmental factors etc. Unfortunately, this approach suffers from the problems of static classification schemes mentioned in previous sections. If new causes are identified then analysis procedures must be revised. If existing records are not updated then this will lead to poor recall during subsequent retrieval tasks.

The integration of probabilistic information retrieval techniques and case-based reasoning techniques provide a further solution to these problems. Our previous use of NaCoDAE has implemented an analysis procedure similar to that described in the previous paragraph (Johnson, 2000b). The problems of extending causal taxonomies are avoided by using Inquiry to identify previous cases that might be affected by any subsequent reclassification. However, this approach does not provide a panacea. In particular, simply knowing that an incident was caused by a perceptual error on the part of the operator or by a particular sub-component failure can provide only limited benefits. Often it is important to explicitly know the reasons WHY an analyst reached such a conclusion. Without this rationale, it can be difficult to explain why particular solutions might be used to effectively avoid future failures.

A number of techniques have been developed to explicitly model argumentation structures (Buckingham Shum, MacLean, Bellotti and Hammond, 1997). Much of this work has focussed upon using computer-supported hypertext tools to denote alternative lines of reasoning. A graphical format is used to denote the arguments for and against particular decisions. This helps other analysts to see the reasons why a conclusion was reached. For example, Figure 1 uses a Conclusion, Analysis and

Evidence diagram to explicitly represent the causes of a collision off the Great Barrier Reef that was published by the Australian Maritime Incident Investigation Unit (MIIU). They concluded that the crew made several human 'errors'. These mistakes included their failure to complete adequate contingency and passage planning. This analysis is supported by evidence that the crew failed to identify the waters off Heath Reef as being restricted for deep draught vessels. The human errors also included a lack of awareness about the other traffic on the reef. This is supported by evidence that both the Fourth Officer and the Commander assumed that the River Embley was some 2.5 miles away when they were, in fact, much closer. As can be seen, CAE diagrams label each conclusion about the causes of an incident with a 'C'. Lines of analysis are labeled using an 'A' and an 'E' denotes the evidence that supports these arguments.

The explicit representation of the arguments that support particular conclusions provides a form of transparency that encourages peer review. Comparing the arguments that are used by different analysts as they investigate different incidents can help to identify potential biases. Computational support offers a number of further benefits. For example, problems can arise when analysts use the same evidence to both weaken and support a single line of reasoning. Alternatively, the same line of reasoning can be used to both support and weaken a conclusion. In either case, it is possible to detect and flag these potential inconsistencies before an incident is entered into a collection. This encourages the inter-analyst reliability that has been identified as a critical problem in previous incident reporting systems (Lekberg, 1997).

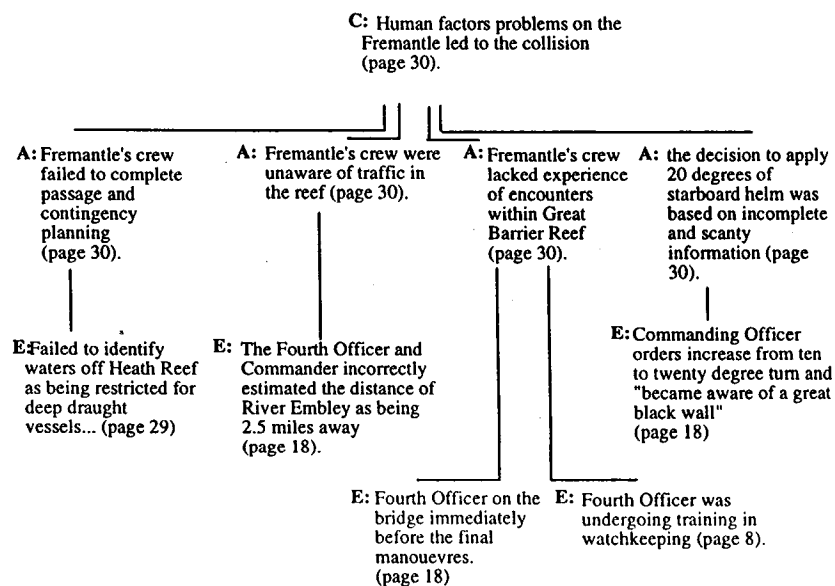


Fig. 1. Conclusion, Analysis, Evidence (CAE) Diagram for the Fremantle Collision

## Conclusions and Further Work

Incident reporting systems help users to provide information about potential safety hazards. They, therefore, represent an important subset of the wider range of applications that support process improvement. However, incident reporting systems have not achieved a "quantum leap" in organizational learning or safety culture. This paper has, therefore, identified a range of novel computational techniques that can be used to address problems of existing reporting systems. In particular it is argued that:

1. Computer-aided interview techniques can be used to prompt contributors for information about the latent and catalytic causes as well as the detection factors and mitigating circumstances that characterize particular incidents. This reduces the need for costly site visits and can reduce the social pressures that arise when contributors must answer the detailed questions of external interviewers.
2. Probabilistic information retrieval engines enable analysts to revise concept structures and thesauri as a means of introducing new indexing features into incident collections. This avoids the problems of poor precision and low recall that stem from the difficulties of maintaining static taxonomies of the causal factors behind common incidents.
3. Conversational case-based reasoning techniques provide users with questions that help to guide their search amongst many thousands of similar incidents. This avoids some of the problems associated with traditional query formation in large-scale incident reporting systems.
4. Techniques from computational models of argumentation can be used to explicitly represent the justification for particular causal classifications of incidents. This improves inter-analyst reliability by helping others to see the evidence that supports findings about previous incidents.

We have undertaken a number of feasibility studies to demonstrate the utility of the approaches that are mentioned in this paper (Johnson, 2000b). It is important to emphasize, however, we have not fully explained all of the opportunities that are created by the novel application of these technologies. For instance, many case-based reasoning tools have been adapted to provide information about particular actions that ought to be taken in response to certain cases. Most typically this centers on repair actions following fault diagnoses. These techniques could easily be extended to record the corrective actions that are taken in response to particular safety-related incidents. More work is urgently required to exploit the opportunities

that are created by these approaches. In the meantime, however, existing reporting schemes continue to receive thousands of submissions that cannot be effectively analyzed and which suffer from a range of biases both in the elicitation and the interpretation of the information that they contain.

**Acknowledgments.** Thanks are due to the members of Glasgow Accident Analysis Group. This work has been funded in part by the UK Engineering and Physical Sciences Research Council.

## References

- Aha, D., Breslow, L.A., and Munoz-Avila, H., 2001. Conversational Case-Based Reasoning. *Journal of Applied Intelligence* (to appear).
- Buckingham Shum, S., MacLean, A., Bellotti, V. and Hammond N., 1997. *Graphical Argumentation and Design Cognition*. *Human Computer Interaction*, (12)3:267-300.
- Busse, D.K. and Wright, D. 1999. Identification and Analysis of Incidents in Complex, Medical Environments. In Johnson, C. (ed.) *Proceedings of the First Workshop on Human Error and Clinical Systems*. Glasgow: Glasgow Accident Analysis Group G99-1.
- Johnson, C.W., 2000a, Reasons for the Failure of CRM Training. To appear in the *International Conference on Human Computer Interaction in Aerospace Applications*.
- Johnson, C.W., 2000b. Using Incident Reporting to Combat Human Error. In G. Cockton (ed.) *People and Computers XIV: Proceedings of HCI 2000*. Berlin: Springer Verlag.
- Lekberg, A., 1997. Different Approaches to Accident Investigation: How the Analyst Makes the Difference. In *Proceedings of the 15<sup>th</sup> International Systems Safety Conference*. Sterling, VA: International Systems Safety Society.
- Reason, J. 1990, *Human Error*, Cambridge: Cambridge University Press.
- Saris, W.E., 1991, *Computer-assisted Interviewing*. Newbury Park: Sage.
- Turtle, H.R., and Croft, W.B., 1991, Evaluation of an inference network-based retrieval model. *ACM Transactions on Information Systems*, 9(3):187-222.
- Weber, R., Aha, D.W., & Becerra-Fernandez, I. (2000). *Categorizing Intelligent Lessons Learned Systems*. This volume.