

The Application of User Modeling Techniques to Reason about the Human Contribution to Major Accidents

C.W. Johnson

Department of Computing Science, University of Glasgow, Glasgow, Scotland.
<http://www.dcs.gla.ac.uk/~johnson>

Abstract. Accident reports are important documents for the development of many interactive systems. They are a primary mechanism by which designers and regulators learn about human ‘error’ and systems ‘failure’. These documents do, however, suffer from a number of limitations. In particular, it can be difficult for readers to trace the ways in which particular individuals are influenced by, and in turn help to influence, their colleagues’ behavior. This paper argues that user modeling techniques can help to improve our understanding of operator interaction in the lead-up to accidents and incidents. This argument is illustrated by a “near miss” that occurred when two Air Traffic Controllers failed to coordinate the arrival and departure of aircraft at Heathrow Airport on the 27th August 1997.

1 Introduction

Human intervention has played a critical role in the causes of many major accidents. The officers and crew of the Herald of Free Enterprise set to sea with their bow doors open (Sheen, 1987). The pilot and co-pilot throttled back their one working engine rather than the engine that had failed prior to the Kegworth air crash (AAIB, 1990). The workers at the Bhopal chemical plant pumped Methyl-isocyanate into a leaking tank (Morehouse and Subamianam, 1986). In all of these accidents, users failed to predict the consequences of their actions because they could not gain an accurate indication about the state of their system. This raises two questions for the design of safety-critical, interactive systems. First, can we help operators to gain a more accurate impression of the state of their application? Second, can we then help users to make better predictions about the effects of their actions based upon an improved view of their application? This paper uses a near-miss incident to illustrate the importance of such situation awareness and anticipatory control for the design of safety-critical, interactive systems.

1.1 The Heathrow Case Study

The case study in this paper centres on the actions of two Air Traffic Controllers at Heathrow Airport in the United Kingdom during August 1997 (Air Accident Investigation Branch, 1997). One was responsible for marshalling arrivals whilst their colleague supervised departures. The Air Departures Controller was also undergoing training and was, therefore, being overseen by a

Mentor. This incident is appropriate because the incident typifies the complex communication failures that can occur between teams of users operating a number of inter-dependent control systems. It also illustrates the profound impact that particular expectations can have upon users' decision making behaviour. A Boeing 737 (SAB603) was attempting to land on Runway 27 Left (27L) when poor weather conditions forced the crew to initiate a missed approach. The Air Arrivals Controller was informed of this and they, in turn, informed the Air Departures Controller. This was standard procedure because the Arrivals Controller would have to alert SAB603 of any departing aircraft that might cause a hazard as they prepared for another landing. The Departures controller alerted their colleague that there was another flight, AFR813, about to depart. They also confirmed that this aircraft would be turned to the right. The Air Arrivals Controller, therefore, requested that SAB603 also turn right to maximise the separation from AFR813. On hearing this, the Mentor who was overseeing the Departures Controller, alerted the Arrivals Controller that another Boeing 757, BAW818, was also departing. The Arrivals Controller immediately ordered SAB603 to turn left while the Departures Controller ordered BAW818 to turn right. According to the Instrument Flight Rules applicable at Heathrow there should have been either 1,000 feet vertical separation or 2.5 nautical miles of horizontal separation between the aircraft. Subsequent calculations revealed that the minimum separation was 200 feet vertically and 0.16 nautical miles of horizontal separation.

1.2 User Modelling and Epistemic Logics

A number of user modelling techniques might be recruited to analyse the situation awareness problems described in the previous section. For example, SOAR provides automated means of analysing the operations interaction with their system. It can, however, be difficult to characterise the Controller's tasks in terms of problem solving activities. The maximisation of throughput and the maintenance of minimum separation criteria form part of several more complex co-ordination tasks. Other user modelling techniques, such as Interacting Cognitive Subsystems (ICS) (Barnard, 1985) and Executive-Process/Interactive Control (EPIC) (Kieras, Wood and Meyer, 1998), can be applied to focus more narrowly on the psychological precursors of the Controllers' interactions. Alternatively, it is possible to recruit stochastic models to represent and reason about the impact of risk and utility on individual decision making (Johnson, 1995a). Unfortunately, a number of problems frustrate these approaches to accident analysis. The limited evidence that is available after a major disaster can prevent investigators from developing detailed psychological models of the individuals that were involved in an accident. The AAIB only provides brief biographies and limited psychological 'profiles' of individual operators. Given this lack of detail, it is almost impossible for analysts to accurately identify the precursors of problems such as high workload. It is, therefore, important to identify some 'coarse grained' means of representing and reasoning about these cognitive factors that are described in accident reports. The following pages, therefore, exploit epistemic logics.

Epistemology, or the study of knowledge, has a history stretching back to the Ancient Greeks (Fagin, Halpern, Moses and Vardi 1995, Barwise and Perry, 1983). This work has produced a number of logics that can be used to represent changes in an individual's knowledge over time. Epistemic logics have a number of attractions for accident analysis. In particular, it is possible to build more complex user models out of epistemic formulae. For example, each of the cognitive

systems in the ICS approach contains an epistemic subsystem. The rest of this paper, therefore, demonstrates that epistemic logics can capture the cognitive observations that are embedded within the AAIB report

2 First Order Logic and Accident Investigations

Logic can be used to represent and reason about the flow of events leading to major incidents. The first stage in this approach is to model the agents and objects that were involved in the accident. For instance, the runways and air traffic control information systems are classified as inanimate objects. Agents include the Air Traffic Controllers as well as the Mentor. One of the benefits of this approach is that analysts can represent both agents and objects at different levels of abstraction. For instance, it is not always necessary to model all of the individuals involved in an accident. In our case study, the accident report does not distinguish between the individual officers onboard SAB 603, AFR 813 and BAW 818. For instance, the following proposition states that the crew of SAB 603 missed their approach to runway 27L.

$$\text{miss_approach}(\text{sab603}, 27L). \quad [1]$$

It is also possible to specify what a particular individual or group of individuals should do under such circumstances. For instance, the “missed approach procedure for aircraft approaching any of the East/West runways at Heathrow, published in the UK Aeronautical Information Publication was for aircraft to climb straight ahead to 3,000 feet...” (AAIB, 1998).

$$\text{miss_approach}(\text{Plane}, \text{Runway}) \wedge \text{east_west}(\text{Runway}) \Rightarrow \text{climb}(\text{Plane}, 3000). \quad [2]$$

This formalisation process is important because proof techniques can be applied to determine whether or not the crew responded appropriately given the context of operation. For example, the following rule states that if we know P and that if P is true then Q is true then it is safe to conclude Q :

$$P, P \Rightarrow Q \vdash Q \quad [3]$$

In the context of the previous discussion, we can apply this inference rule to assess whether or not the crew of SAB 603 correctly followed the missed approach procedure. In order to do this, we must first determine whether the antecedent, P , of the implication actually holds. In other words, did the aircraft miss its approach and was 27L an East-West runway. We already know from formula [1] that the aircraft missed its approach. It is also possible to confirm that 27L was an West-East runway (AAIB, 1998). It is, therefore, possible to conclude that the crew did satisfy the relevant operating procedures of [2] when they climbed to 3,000 feet: “they complied with the laid down procedure which was to ‘Climb straight ahead to 3,000 feet then as directed by the Air Traffic Control’” (AAIB, 1998). This might seem relatively obvious, however, if further enquiries revealed that 27L was not an East-West runway then the procedures of [2] need not have applied and, in turn, the crews actions may not have been appropriate.

As mentioned, logic can be used to focus in upon critical incidents in the course of an accident. For instance, it is important that the Arrivals Officer informs the Departures Officer if an aircraft misses its approach because the aircraft could cross the flight paths of aircraft that are departing from another runway. The following formulae not only illustrates how this requirement can be represented, it also shows how logic can explicitly represent implicit assumptions about an operator's behaviour. In this case, we specify that the Arrivals officer must first observe the missed approach if they are to warn their colleague:

$$\text{miss_approach}(\text{Plane}, \text{Runway}) \wedge \text{confirm}(\text{atc_arrivals}, \text{Plane}, \text{missed_approach}) \Rightarrow \text{inform_miss_approach}(\text{atc_arrivals}, \text{atc_departures}, \text{Plane}, \text{Runway}). \quad [4]$$

We can again check to determine whether these antecedents did, indeed, hold during the case study incident. We already know from [1] that the approach was missed and we know that "Air Arrivals had seen the aircraft going around and had acknowledged the manoeuvre by radio to the crew" (AAIB, 1998):

$$\text{confirm}(\text{atc_arrivals}, \text{sab603}, \text{missed_approach}) \quad [5]$$

Given this series of behavioural requirements on the operators, it is possible to reconstruct the ways in which events unfolded during the particular incident as follows:

$$\text{miss_approach}(\text{SAB603}, 271) \wedge \text{confirm}(\text{atc_arrivals}, \text{sab603}, \text{missed_approach}) \Rightarrow \text{inform_miss_approach}(\text{atc_arrivals}, \text{atc_departures}, \text{sab603}, 271).$$

$$\text{(Instantiation of SAB603 for Plane and 271 for Runway in [4])} \quad [6]$$

$$\text{inform_miss_approach}(\text{atc_arrivals}, \text{atc_departures}, \text{sab603}, 271).$$

$$\text{(Modus ponens, application of [3] to [6] given [1] and [5])} \quad [7]$$

This line of reasoning leads to yet further requirements on the operators; "Air Arrivals and Air Departures are to co-ordinate with each other to establish separation between the "go-around (aircraft)" and any conflicting departing traffic" (AAIB, 1998). The following clause focuses on the Departures Officer and specifies that if anyone informs them of a missed approach then they must check for any conflicts with other aircraft. Additional clauses can be added to such formulae to represent subsequent instructions to alter the height or direction of those aircraft:

$$\text{inform_miss_approach}(\text{Anyone}, \text{atc_departures}, \text{Plane}, \text{Runway}) \Rightarrow \text{conflicting_traffic}(\text{atc_departures}, \text{Plane}, \text{Runway}, \text{Conflicts}) \quad [8].$$

As before, we can apply the reasoning techniques of first order logic to formalise the changing requirements that must be satisfied by the operators of safety critical control systems:

$\text{inform_miss_approach}(\text{atc_arrivals}, \text{atc_departures}, \text{sab603}, 271) \Rightarrow$
 $\text{conflicting_traffic}(\text{atc_departures}, \text{sab603}, 271, \text{baw818}).$ [9]

$\text{inform_miss_approach}(\text{atc_arrivals}, \text{atc_departures}, \text{sab603}, 271) \Rightarrow$
 $\text{conflicting_traffic}(\text{atc_departures}, \text{sab603}, 271, \text{afr813}).$ [10]

(Instantiation of atc_arrivals for Anyone, sab603 for Plane, 271 for Runway etc in [8])

$\text{conflicting_traffic}(\text{sab603}, 271, \text{baw818}) \wedge \text{conflicting_traffic}(\text{sab603}, 271, \text{afr813}).$

(Modus ponens, application of [3] to [9] given [8]) [11]

This final inference brings us to the heart of the accident because it captures the assumption that both BAW 818 and AFR 813 were in potential conflict with the revised flight path of SAB 603. Unfortunately, the MATS regulations cited earlier rely upon a subjective definition of what constitutes conflicting departing traffic. “However, the actions of the Air Departures was based on the supposition that Air Arrivals would turn SAB 603 to the left. She did not consider that BAW 818 was a confliction and therefore did not inform Air Arrivals of its departure” (AAIB, 1998). Unfortunately, if we want to represent and reason about such differences in the knowledge and beliefs of agents we must move beyond the relatively simple first order logics used in the previous formulae to higher order, epistemic logics.

3 Epistemic Logics and User Modelling in Accident Investigations

The Air Departures Controller did not consider that BAW 818 was in potential conflict with SAB 603 as it missed its approach. This can be represented by the following negation:

$\text{not}(\text{conflicting_traffic}(\text{sab603}, 271, \text{baw818}))$ [12]

From a logic perspective, this leads to a contradiction because we have already stated in [11] that both BAW 818 and AFR 813 were, in fact, in conflict with SAB 603. We cannot in first order logic allow the same fact to be both true and not true at the same time:

$\text{conflicting_traffic}(\text{sab603}, 271, \text{baw818}) \wedge \text{not conflicting_traffic}(\text{sab603}, 271, \text{baw818})$

(\wedge introduction, [12] and [11]) [13]

Epistemic logics provide means of avoiding such contradictions. For example, we can introduce a modal operator to express the fact that the ATC Departures Officer knows that BAW 818 does not conflict with SAB 603:

$\text{Kn}_{\text{atc_departures}}(\text{not conflicting_traffic}(\text{sab603}, 271, \text{baw818})).$ [14]

This does not lead to a contradiction. It is perfectly possible for other operators, such as the Arrivals Officer, to believe that BAW 818 does conflict with the going around aircraft at the same time that the Departure Officer knows that it does not create a conflict:

$$\text{Kn}_{\text{atc_departures}}(\text{not conflicting_traffic}(\text{sab603}, 271, \text{baw818})) \wedge \text{Kn}_{\text{atc_arrivals}}(\text{conflicting_traffic}(\text{sab603}, 271, \text{baw818})) \quad [15]$$

The semantics of the Kn operator is introduced in Appendix A. In contrast the remainder of this section goes on to apply this notation to represent and reason about the causes of the communications failures that led to the near miss at Heathrow Airport. This relatively simple notational extension opens up a range of user modelling techniques. For instance, it is also possible to reason about an individual's knowledge of the information that is available to their colleagues. In our case study, it can be argued that the Departure Officer did not know whether or not the Arrivals Officer knew about BAW 818:

$$\text{not Kn}_{\text{atc_departures}}(\text{Kn}_{\text{atc_arrivals}}(\text{conflicting_traffic}(\text{sab603}, 271, \text{baw818}))) \wedge \text{not Kn}_{\text{atc_departures}}(\text{Kn}_{\text{atc_arrivals}}(\text{not conflicting_traffic}(\text{sab603}, 271, \text{baw818})))) \quad [16]$$

The previous clause illustrates the use of epistemics to reason about an individual's view of their colleague's knowledge. It is also possible to reason more generally at the level of "common knowledge". These shared assumptions often lie at the heart of major accidents. They, typically, cause problems for operators in one of two ways. Firstly, the knowledge may not be available to all of the members in a group. Secondly, the shared information may be incorrect. This is particularly dangerous because users seldom challenge common knowledge when they rely upon the understanding of their colleagues. In order to represent this form of knowledge we must first define the group which shares the knowledge:

$$\text{atc_officers} = \{\text{atc_arrivals}, \text{atc_departures}, \text{departures_mentor}\} \quad [17]$$

An epistemic operator E_{group} can then be used to state that everyone in a group knows that a certain proposition is true. As before the semantics for this and the other epistemic operators are provided in Appendix A. This operator is particularly useful for our case study because the Departure Officer's decision not to communicate with the Arrival Officer about BAW 818 was partly justified by a belief that SAB 603 would be turned to the left. "the Mentor stated that he had never previously witnessed a missed approach from Runway 27l being turned right. The assumption was therefore made that SAB 603 would be turned left away from the departure runway and Air Departures, acting on this expectation failed to inform Air Arrivals of the second departure i.e., BAW 818 since she perceived the main requirement at the time as dealing with AFR 813":

$$E_{\text{atc_officers}}(\text{alter_course}(\text{atc_arrivals}, \text{sab_603}, \text{left})) \quad [18]$$

The previous clause states that all ATC officers know that SAB 603 is requested to alter course to the left. However, the previous quotations indicate an even stronger assumption in which not

only did the officers know that SAB 603 would be turned to the left but that they assumed that the others in the group knew this. The stronger C_{group} operator can be used to capture this assumption:

$$C_{atc_officers}(\text{alter_course}(\text{atc_arrivals}, \text{sab_603}, \text{left})) \quad [19]$$

Unfortunately, neither clause [18] nor clause [19] capture the true state of operator knowledge in the lead up to the incident. The Arrivals Officer did not share the same set of assumptions as either the Departure Officer or their mentor. They chose to turn SAB 603 to the right rather than to the left. The following section looks at the reasons for this decision.

4 Epistemic Logics and the Causes of Operator Error

The near miss stemmed from the Arrival Officer's decision to alter the course of SAB 603 to the right. This decision went against normal procedure but was based upon a desire to increase the separation between SAB 603 and AFR 813. "Air Arrivals decided to turn SAB603 right because he considered that a left turn would have been a possible confliction to the 'Midhurst' departure (AFR 813) and a right turn would cause less disruption as he was not aware of the 'Brookmans Park' departure" (AAIB, 1998):

$$\text{Kn}_{atc_arrivals}(\text{conflicting_traffic}(\text{sab603}, 271, \text{afr813})) \wedge \\ \text{not Kn}_{atc_arrivals}(\text{conflicting_traffic}(\text{sab603}, 271, \text{baw818})) \Rightarrow \\ \text{alter_course}(\text{atc_arrivals}, \text{sab_603}, \text{right}) \quad [20]$$

The previous clause illustrates how the decision to turn the aircraft to the right was based both on the controller's knowledge of AFR813 and their ignorance of BAW818. This analysis was confirmed, with hindsight, by the subsequent investigation. "Air Arrivals stated, however, that had he known about BA 818, he would have turned the missed approach aircraft left in the manner anticipated by his colleagues":

$$\text{Kn}_{atc_arrivals}(\text{conflicting_traffic}(\text{sab603}, 271, \text{baw818})) \Rightarrow \\ \text{alter_course}(\text{atc_arrivals}, \text{sab_603}, \text{left}) \quad [21]$$

None of this, of course, explains why the Departure's Controller decided not to inform their colleague about BAW818. "The assumption was therefore made that SAB 603 would be turned left away from the departure runway and Air Departures, acting on this expectation failed to inform Air Arrivals of the second departure i.e., BAW 818":

$$\text{alter_course}(\text{atc_arrivals}, \text{sab_603}, \text{left}) \Rightarrow \\ \text{Kn}_{atc_departures}(\text{not conflicting_traffic}(\text{sab603}, 271, \text{baw818})) \quad [22]$$

Clauses [21] and [22] illustrate how epistemic logic helps to focus on the communication problems that lies at the heart of our case study. If the Arrivals Officer had known about the

conflict with BAW 818 then they would have routed SAB 603 to the left. The Departures Officer did not inform them of the potential conflict because if SAB 603 had been routed to the left then there would have been no conflict.

Our approach is limited in the sense that it focuses on knowledge requirements during interaction with safety-critical interfaces. We have not attempted to represent the physiological factors that contribute to many major accidents. Elsewhere we have described how logic can be used to reason about the physical layout of complex working environments (Johnson, 1996). Future work intends to build on this to represent the ways in which operators must integrate computer-based tasks with more general, physical activities in safety-critical systems. This can be achieved by using epistemics to represent and reason about components of more general user models such as EPIC and ICS that both account for physiological aspects of interaction. For instance, epistemic logics can express the propositional and implicational information that is represented in the ICS model (Busse and Johnson, 1998). The intention would be to provide a concrete syntax and vocabulary for analysing the sorts of information and knowledge that might be held within these subsystems during particular traces of interaction.

5 Conclusion

This paper has argued that operators often make incorrect assumptions about the information that is known to their colleagues. This, in turn, can prevent them from correctly anticipating the future behaviour of other workers. Although this is a serious problem, it does not always threaten safety. Users can initiate repair activities by asking colleagues about their knowledge and intentions at regular intervals during interaction. Safety can, however, be jeopardised if operators forget or consider it unnecessary to verify their understanding of their colleagues' view of the situation.

This argument has been illustrated by a case study in which the Air Traffic Controllers at Heathrow Airport failed to ensure adequate separation between two aircraft. The case study illustrated the complexity that arises when using natural language to discuss the recursive nature of shared knowledge. We quickly begin to talk about an individual's knowledge of another individual's knowledge. Epistemic logics provide a means of addressing this complexity because they provide a clear and coherent language in which to reason about both individual and shared knowledge. Using this notation, it was possible to identify the mutual dependency that led to the incident in the first place. This dependency prevented each of the operators from gaining an accurate impression of the information that was known to the other. This can be summarised as follows: if the Arrivals Officer had known about the conflict with BAW 818 then they would have routed SAB 603 to the left. The Departures Officer did not inform the Arrivals Officer of the potential conflict because if SAB 603 had been routed to the left then there would have been no conflict.

This work represents a first step towards a more detailed understanding of the errors that arise when groups of individuals must share their knowledge to improve a collective "situation awareness". Much remains to be done. We have already argued that epistemic logics must be integrated with more general user modelling techniques. The great weakness, and also the great strength, of this notation is that it focuses upon a specific aspect of interaction. This makes it

appropriate for accidents that stem from communications problems, such as those mentioned above. It is far less appropriate for accidents that involve perceptual or physiological problems. However, the importance of this focus should not be underestimated. The Air Accident Investigation Branch (1998) concluded “An important element of Air Traffic Control is communication. It is not altogether surprising, therefore, that communication breakdown occurs or that it should figure so prominently in the occurrence of incidents. The issue is to understand why communication fails and what can be done to prevent such failure in the future or, at least, to mitigate its effects”. User modelling provides tools to improve our understanding of these failures. Without such tools, many incidents will continue to be classified as simple instances of operator ‘error’.

Acknowledgements

Thanks go to the members of the Glasgow Interactive Systems Group (GIST) and to the Formal Methods and Theory Group in Glasgow University. This work is supported by the UK Engineering and Physical Sciences Research Council, grants GR/JO7686, GR/K69148 and GR/K55040.

6 Appendix A: Semantics of the Epistemic Logic

In formal terms, the Kn operator relies upon a Kripke Semantics which is represented by the following tuple: $(S, \pi, K_1, \dots, K_n)$. S is a set of states or possible worlds. π is an interpretation function that associates the values true or false with primitive propositions in each state. Therefore, we can apply $\pi(s)(p)$ to determine whether or not the proposition p is true in state s . In contrast, K_i is a binary relation between states. It is used to denote whether an agent i believes a state to be possible from the information that they have in the current state. We can define the truth of a proposition inductively as follows:

$$\begin{aligned} (M, s) \models p & \text{ (for a primitive proposition } p) \text{ iff } \pi(s)(p) = \text{true} \\ (M, s) \models p \wedge q & \text{ iff } (M, s) \models p \text{ and } (M, s) \models q \\ (M, s) \models \text{not } p & \text{ iff } (M, s) \not\models p \\ (M, s) \models K_n p & \text{ iff } (M, t) \models p \text{ for all } t \text{ such that } (s, t) \in K_n \end{aligned}$$

The final case illustrates how an agent i knows a proposition p to be true in a state s if p is true in all states that i considers to be possible from s . In passing it should be mentioned that K_n is an equivalence relation and is based on the initial epistemic model introduced by Fagin, Halpern, Moses and Vardi (1995).

The semantics of common knowledge can be introduced as follows for E_{group} (read as ‘everyone knows’), C_{group} (‘everyone knows that everyone knows...’) and D_{group} (‘information is shared amongst the group’):

$$\begin{aligned} (M, s) \models E_{group} p & \text{ iff } (M, s) \models K_n p \text{ for all } i \text{ in } group \\ (M, s) \models C_{group} p & \text{ iff } (M, s) \models E_{group} p \text{ for all } i \text{ in } group \\ (M, s) \models D_{group} p & \text{ iff } (M, s) \models p \text{ for all } t \text{ such that } (s, t) \in \bigcap_{i \text{ in } group} K_n \end{aligned}$$

The final definition illustrates how a proposition p is known amongst a group by removing all of the worlds that either agent considers to be impossible. What remains is the intersection of their knowledge. This models the situation where one or more agents could communicate to achieve a better understanding than would have otherwise been possible.

References

- Air Accidents Investigations Branch, Department of Transport. Report On The Accident To Boeing 737-400 G-OBME Near Kegworth, Leicestershire on 8th January 1989, number 4/90, Her Majesty's Stationery Office. London, United Kingdom, 1990
- Air Accident Investigation Branch, Report on the Incident near London Heathrow Airport on 27 August 1997, Aircraft Incident Report Number 5/98, Her Majesty's Stationery Office, London, 1998.
- P. Barnard, Interacting Cognitive Subsystems: A Psycholinguistic Approach to Short Term Memory. In A. Ellis (ed) Progress in the Psychology of Language, Lawrence Erlbaum Associates, Hove, United Kingdom, 197-258, 1985.
- J. Barwise and J. Perry. *Situations And Attitudes*. Bradford Books, Cambridge, United States of America 1983.
- D. Busse and C.W. Johnson, Modelling Human Error within a Cognitive Theoretical Framework. In F.E. Ritter and R.M. Young (eds.) The Second European Conference on Cognitive Modelling, Nottingham University Press, 90-97, 1998.
- R. Fagin, J. Halpern, Y. Moses, and M. Vardi. *Reasoning About Knowledge*. MIT Press, Boston, United States of America 1995.
- C.W. Johnson, *A Probabilistic Logic For The Development of Safety-Critical Interactive Systems*. International Journal of Man-Machine Systems, 39(2):333-351, 1993.
- C.W. Johnson, *The Formal Analysis Of Human-Computer Interaction During Accident Investigations*. In G. Cockton, S.W. Draper and G.R.S. Weir, editors, People And Computers IX, 285-300. Cambridge University Press, Cambridge, 1994.
- C.W. Johnson, *The Application of Petri Nets to Represent and Reason about Human Factors Problems During Accident Analyses* In P. Palanque and R. Bastide, editors, The Specification And Verification Of Interactive Systems, 93-112, Springer Verlag, Berlin, 1995.
- C.W. Johnson, *Decision theory and Safety-Critical Interfaces*. In K. Nordby, P.H. Helmersen, D. Gilmore and S. Arensen (eds.) Interact '95, 127-132, Chapman Hall, London, United Kingdom, 1995a.
- C.W. Johnson, Impact of Working Environment upon Human-Machine Dialogues: A Formal Logic for the Integrated Specification of Physical and Cognitive Ergonomic Constraints on User Interface Design. Ergonomics (39)3:512-530, 1996.
- D. Kieras, S.D. Wood, D.E. Meyer, Predictive Engineering Models based on the EPIC Architecture for Multimodal High Performance Human Computer Interaction Tasks. To appear in ACM Transactions on Human Computer interaction, 1998.
- W. Morehouse and M.A. Subramaniam, The Bhopal Tragedy. Technical Report. Council for International and Public Affairs, New York, United States of America, 1986.
- D. Norman, *The 'Problem' With Automation : Inappropriate Feedback And Interaction Not Over-automation*. In D.E. Broadbent and J. Reason and A. Baddeley (eds.), Human Factors In Hazardous Situations, 137-145, Clarendon Press, Oxford, United Kingdom, 1990.
- J. Reason, *Managing the Risks of Organisational Accidents*, Ashgate, Aldershot, United Kingdom, 1997.
- Sheen, Formal Investigation into the Sinking of the mv Herald of Free Enterprise, UK Department of Transport, Report of court 8074, Her Majesty's Stationery Office, 1987.

