Signals Passed at Danger: A Case Study in the Application of Visualisation Techniques

Fraser Speirs and Chris Johnson

Department of Computer Science
University of Glasgow
17 Lillybank Gardens
Glasgow
G12 8QQ

Tel: *+44 0141 339 8855 x0917*
Fax: *+44 0141 330 4913*
EMail: {*speirsfr, johnson*}*@dcs.gla.ac.uk*
URL: *http://www.dcs.gla.ac.uk/˜*{*speirsfr, johnson*}

Incident reports are increasingly being used to learn from near-misses in many different industries. These systems are important to safety-critical HCI for two reasons. Firstly, they provide information about near-misses caused by HCI failures. Secondly, HCI techniques are required to help analysts gain an overview of the mass of information they contain. The US Aviation Safety Reporting System has over 800,000 records. This paper presents the reporting of Signals Passed at Danger (SPAD) - the term used to describe an incident when a train passes a signal displaying a danger aspect without authorisation. This provides case studies for the application of visualisation techniques to safety-critical systems. This example was chosen because failure to react to data about previous SPAD incidents has been cited as a primary cause in several recent rail accidents. The paper surveys a number of approaches that might be applied to better support users tasks and presents a number of prototype displays that are currently being developed to improve the presentation of this data. The closing sections review the validation issues that arise from this work.

1 Introduction

The train crashes at Southall and Ladbroke grove occurred as a result of SPADs. Given that SPAD records have been kept for many years, this raises a question as to why the issues were not tacked sooner. The reasons why SPADs occur are complex and not entirely understood as yet. We argue that HCI can help by allowing analysts to get a clearer view of the collected data through advanced visualisation techniques. There are certainly organisational and regulatory reasons why the issue of SPADs has not been eliminated, but we believe that HCI has a contribution to make.

2 Visualisation Techniques

Information visualisation has been given one definition as ”The use of computer-supported interactive, visual representations of abstract data to amplify cognition” (Card et al., 1999). A number of highly influential techniques for information visualisation have been developed in recent years. Techniques such as starfield displays (Ahlberg & Shneiderman, 1994), self-organizing maps (Lagus et al., 1996) and Bead (Chalmers & Chitson, 1992) have all found application in various domains. This paper examines the potential uses of such techniques to support the reporting of safety-critical incidents.

In particular, we focus on the reporting of Signals Passed at Danger (SPAD), the term used to refer to an incident in which a train has passed a signal displaying a danger aspect without authorisation. In the UK rail industry, the primary responsibility for safety lies with Her Majesty's Railway Inspectorate (HMRI), a division of the Health and Safety Executive (HSE). HMRI monitors SPAD incidents and produces a report on a monthly basis that is circulated in the industry and available on the HMRI website (http://www.hse.gov.uk/railway/spad/index.htm).

Many visualisations have been developed for the purpose of browsing large collections of documents, or large data sets, such as FilmFinder (Ahlberg & Shneiderman, 1994), shown in figure 1. These are directly applicable to the problem of visualising incident reports. Incident reporting, however, has its own requirements which are peculiar to the application domain, these must be fulfilled if the technique is to support the users' tasks. The data reported in the HMRI SPAD reports summarise important questions for incident investigators including:

1. Determining trends over time. Much of the SPAD report is taken up with discussing changes in the rate of occurrence of SPADs over the reporting period. Figure 2 shows the current approach to displaying trend information, a statistical diagram.

2. Detecting 'hot-spots' where there may be an inherent safety problem. In SPAD investigation, such hot-spots manifest themselves as 'multiple-SPAD' signals.

   The following SPAD incident occurred at a signal that had a previous history of SPAD incidents, eight in total:

   > "The SPAD at Narroways Hill Junction, Bristol (signals B248 and B250 on 1 November 2000) gave rise to a rear end collision between an empty Royal Mail train and a loaded coal train

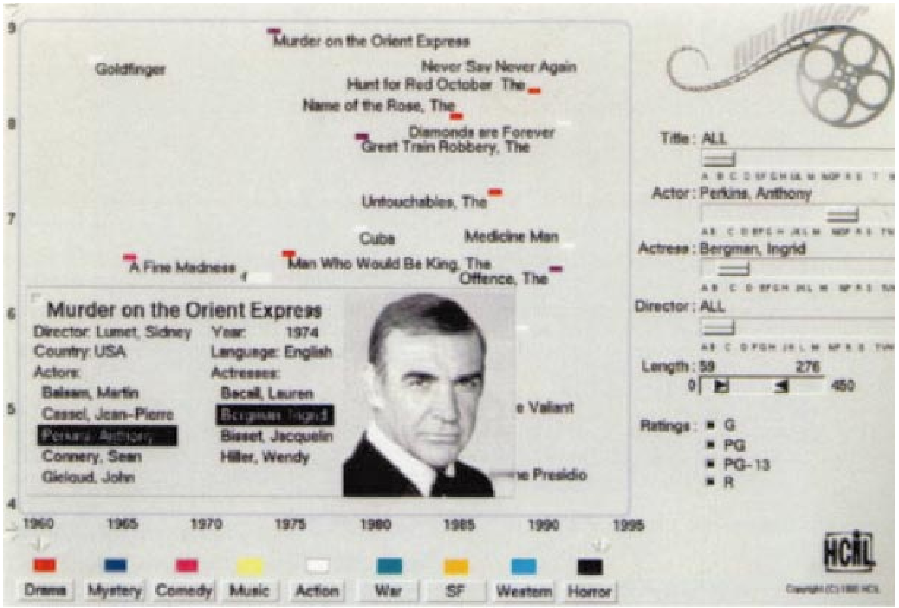

Figure 1: Filmfinder, a dynamically queried visualisation for film information

> proceeding at slow speed ahead. Both trains were operated by EWS. The Mail train had a locomotive at both ends of the train due to service disruptions. There had been problems in releasing the train brakes after the train was remarshalled at Bristol Parkway, and the train eventually left and descended Filton Bank (a gradient of 1 in 75) towards Bristol. The train was unable to brake in order to stop at the two signals at Danger protecting the coal train ahead and collided with it ." (Her Majesty's Railway Inspectorate, 2000b)

3. Determining causal factors, particularly factors common across many incidents. Rail adhesion is a particularly strong theme running through many SPAD incidents:

> "The SPAD at Three Bridges (signal T626 on 1 November 2000) involved a stoneblower machine (a self-propelled machine used in track maintenance) operated by AMEC Rail. After passing the signal at Danger, the machine ran over a level crossing with the barriers raised and open to road traffic. Fortunately, as the incident occurred in the early hours, there was no road traffic about. HMRI's investigation found that the SPAD occurred because of poor rail adhesion causing the wheels to slide when the brakes were applied. The area was not known for poor adhesion but has since been added to the Sandite programme (Sandite is a gell/sand mix laid on the

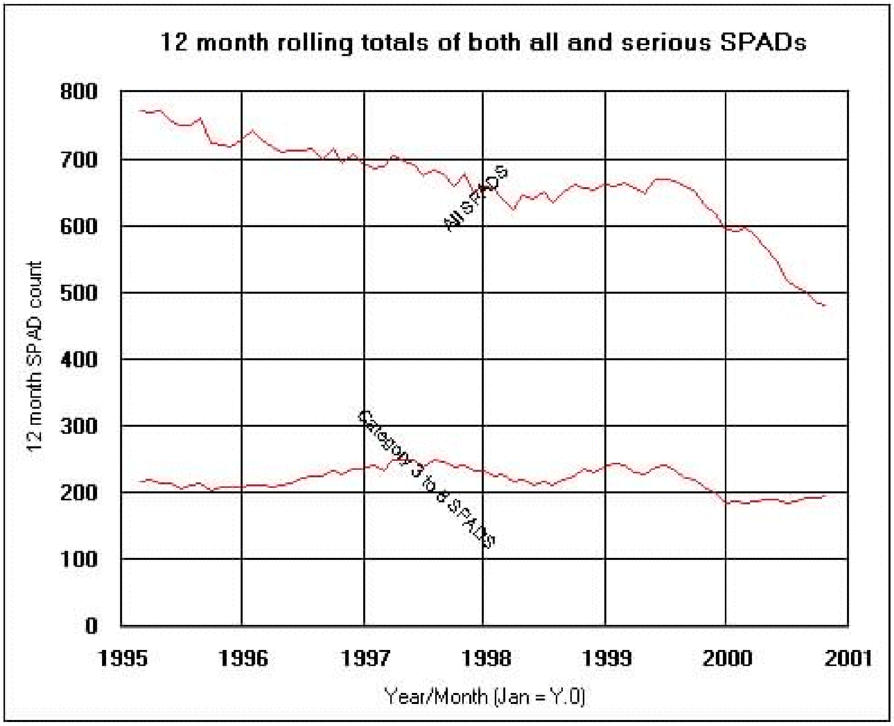

Figure 2: Rolling 12-month total of SPAD numbers

> rail head to improve adhesion during the leaf fall season)." (Her Majesty's Railway Inspectorate, 2000b)

3 Incident Reporting Schemes

Incident data are typically collected through a confidential incident reporting scheme, such as Reason discusses in the context of organisational failure (Reason, 1997). Such schemes anonymise the data, aggregate it and present it to the wider community.

3.1 The HMRI SPAD Reporting System

When a SPAD occurs, it is normally immediately apparent to the signaller that it has happened because the signalling equipment gives a real-time display of train locations on the network. In the UK, it is the signaller's duty to report the incident to Railtrack, the rail infrastructure company, who then conduct an investigation into the incident along with the train operating company involved. The SPAD reporting system is neither anonymous nor confidential, unlike many schemes operating in other areas of the rail system. In these reports, SPADs are categorised according to severity, from 1 to 8, according to the following scale:

1. Overrun 0 - 25 yards, no damage.

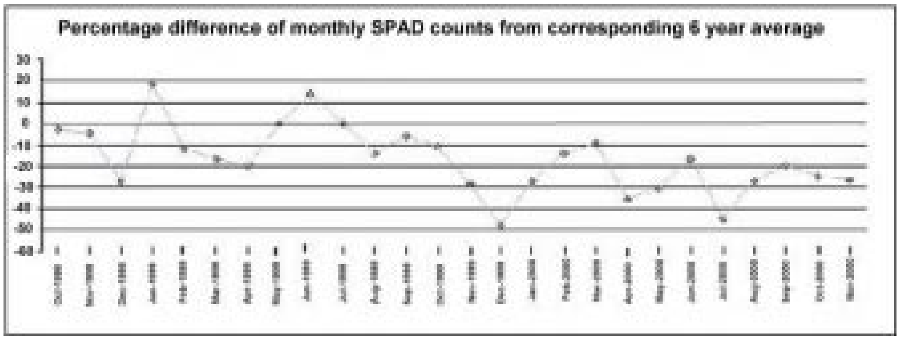

Figure 3: Percentage difference of monthly SPAD numbers from six-year average

2. Overrun 26 - 200 yards, no damage.
3. Overrun greater than 200 yards, no damage.
4. Track damage, no casualties.
5. Derailment, no collision, no casualties.
6. Collision (with or without derailment), no casualties.
7. Injuries to staff or passengers with no fatalities (including driver absence consequent upon trauma, without actual physical injury).
8. Fatalities to staff or passengers.

Severity is a particularly important attribute of a SPAD report, as analysing it with respect to other factors is likely to yield vital insight. For example, examining severity and location together would indicate places where the infrastructure is particularly poor from a human factors point of view.

3.2 The HMRI SPAD Report

Every month HMRI publishes a document that summarises SPAD incident data collected in that month. The examples here come from the November 2000 issue (Her Majesty's Railway Inspectorate, 2000b). The SPAD report is published as both a statistical record of SPAD incidence and as a digest of notable SPAD incidents from the set of reports for the month. The preceding quotations from the SPAD incidents at Narroways Hill and Three Bridges are examples of SPADs selected for inclusion in the monthly digest. The report gives aggregate figures for SPAD incidents in a given month, and makes comparisons based on 6- and 12-monthly reporting periods, as well as six-year averages. Figure 2 shows a 12-month rolling total of SPAD numbers; figure 3 shows a percentage difference from the six-year average. These statistical diagrams represent the extent to which visualisation techniques are currently used in incident reporting.

SPADs categorised as levels 1 and 2 are not reported individually. The following is an example of a summary of a level 3 SPAD:

> *"The SPAD at Sandy (signal P10 on 15 August 2000) occurred when the train service was disrupted due to the failure of a WAGN train in*

the Up (London bound) platform. So that another WAGN Up train could call at Sandy, the signaller routed this train through the station on the Up fast line and then, following the train's reversal, into the Down platform. The train then departed from the station under the authority of a position light signal, signal P14, and passed the next signal, signal P10, also a position light signal, at Danger. Fortunately, the route ahead of signal P10 had cleared for the movement so there was little chance of adverse consequences occurring. HMRI has written to Railtrack requiring a review of the current signalling at Sandy and has advised that signal P10 should be able to clear before signal P14 can clear. This would prevent a further SPAD incident of the type occurring at signal P10." (Her Majesty's Railway Inspectorate, 2000b)

This excerpt is typical of the type of information that is available about any given SPAD incident. Such incidents are typically so numerous that sufficient resources do not exist to investigate them all in depth.

4 A case study of incident visualisations

A database of incident reports, such as the SPAD report, can show which activities have regular safety issues. These regular issues may include multiple-spad signals or common root causes of SPAD incidents such as poor rail adhesion, which is now a well known cause of SPAD incidents. Databases can also help an investigator to discover why an accident has not occurred (Reason, 1997), or in Reason's terms, what defences were not breached and thus worked to prevent the accident. This is significant from an HCI perspective because human intervention is a key factor in mitigating potential consequences of a SPAD. Such intervention is increasingly being mediated through computer-based signalling systems. It is therefore *critical* that usability concerns be explicitly considered when analysing SPAD reports.

Current techniques for communicating this information are quite rudimentary in that they allow he reader to check certain statistics, but not to truly explore the data. The actual statistical techniques used are quite sophisticated, but the presentation of the results is static and inflexible. The current methods of presentation do not take advantage of the possibilities offered by recent advances in information visualisation. The remainder of this paper looks at a number of visualisation systems that have been developed and attempts to apply their ideas to the SPAD incident data.

4.1 Computerising incident data

Many organisations that run incident reporting schemes have developed databases of their reports. Given the volume of reports that some schemes receive, this is a necessity. A database, coupled to reporting tools and data-mining visualisation techniques is a very appropriate means of dealing with a volume of incident reports.

Incident reporting data is typically a large set of data; long-running schemes such as the NASA-run Aviation Safety Reporting System (http://asrs.arc.nasa.gov) may have as many as 800,000 separate incident reports. The SPAD report presents selections from around 600 incidents per month. The challenges involved in learning from such a knowledge-base are quite different from those posed by individual, large documents such as an accident report.

The entire field of information retrieval developed around such problems, but the newer field of information visualisation looks specifically at the interaction issues surrounding the need to learn from large data sets. Shneiderman (Shneiderman, 1996) gives a "Task by Data Type Taxonomy for Information Visualization", which discusses tasks and supporting strategies for visualising large high-dimensional data sets.

The next three sections take each of the three key tasks - monitoring trends, finding hot-spots and determining causality - and suggests possible visualisation approaches to the problem.

4.2 Location Identification

A primary role for incident reporting is to identify locations that may be particularly notable for the number of incidents occurring there. With SPAD incidents, a multiple-SPAD signal may be sited poorly, badly aligned, obscured by undergrowth or affected by the sun at particular times of day. A visualisation system for SPAD reports would need to include some means of indicating problem areas.

Geographic Information Systems, or GIS, are probably the best known geographically-based visualisation tools today. They are used in many organisations for planning and analysis of information such as census data, urban planning and conservation projects. Other systems have been developed that display traffic on data networks with displays showing geographical location (Becker et al., 1995).

For SPAD reporting, such displays would aid an investigator's task of identifying 'hot-spots' on the railway network. When combined with interactive display techniques such as starfield displays (Ahlberg & Shneiderman, 1994) and dynamic queries (Shneiderman, 1994) this can provide an overview of the data by filtering out unnecessary details, such as exact signal numbers, which can be viewed in the full-text report.

We have developed an early prototype that combines these techniques into a tool for browsing SPAD data. The interface comprises of a geographical display of the rail network, over which is laid 'markers' that identify SPAD locations. The user is provided with paired sliders that are used to define ranges for numerical attributes of the data, and pop-up menus to select certain predefined sub-sets of the data such as incidents caused by drivers of a particular company.

Figure 4 shows our prototype displaying a section of the Midlands rail network. It should be noted that hypothetical data has been used for the purposes of illustration. Each small square represents one SPAD incident, the stacked squares indicate a multiple-SPAD location. As the sliders at the right are moved to dynamically change the query parameters, the indicators are hidden if the SPAD they represent is not in the query set. Similarly, other SPAD indicators appear when they become part of the query set.

4.3 Trend Identification

An important part of HMRI's purpose in preparing the SPAD Report is to identify whether the trends in SPAD incidents are rising or falling. This is in line with their

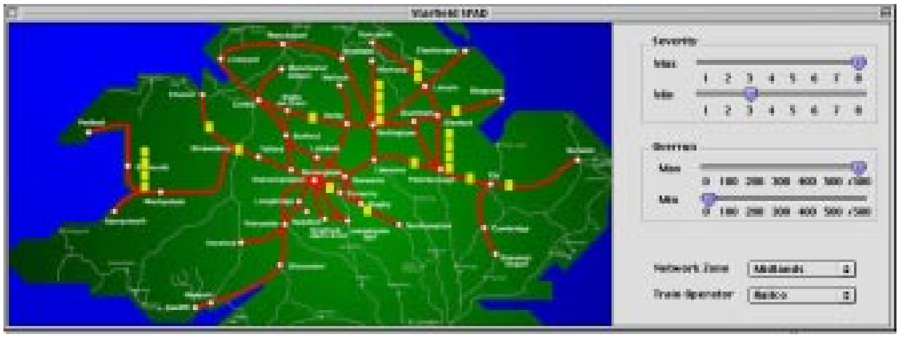

Figure 4: An early prototype of a SPAD incident browser

role in ensuring that train operating companies run their services safely, and take steps to address problems.

HMRI has used traditional statisical graphing to report numerical analysis of SPAD incidents ever since the SPAD report published in December 1999, but this is itself not without problems:

> ”This month, HSE statisticians have advised that the method of presenting SPAD information should be revised to avoid potentially misleading interpretation of the standard analysis. Previously, as in [figure 2 & 3], the standard presentations each month have been represented using the ratio of that month’s SPAD count with the average of the corresponding months of the six preceding years. These data were then plotted out month by month, together with a ”trend” line fitted to these points (technically a linear least-squares squares regression line). The main visual message of this representation is the trend line, which could lead to readers naturally assuming that this represents the trend in SPAD numbers (i.e. if the slope is up, SPADs are increasing, and vice versa). However this natural assumption is wrong. The slope of the trendline indicates whether the change in SPAD numbers (change being measured over the past six years) is getting bigger or smaller month by month, regardless of whether the change itself is upwards or downwards. A flat trend indicates a steady increase or decrease: it does not discriminate between the two. The previous presentation thus shifts attention from the issue of primary interest (are SPADs increasing or decreasing?), to a secondary issue (is the rate of change in SPAD numbers increasing or decreasing?).” (Her Majesty’s Railway Inspectorate, 2000b)

This illustrates a key point, that standard statistical visualisations can be misleading and that usability issues are critical for any visualisation technique. While such graphs perform well in some circumstances, they do not support the interactive discovery of hidden trends within the data set. Shneiderman, in particular,

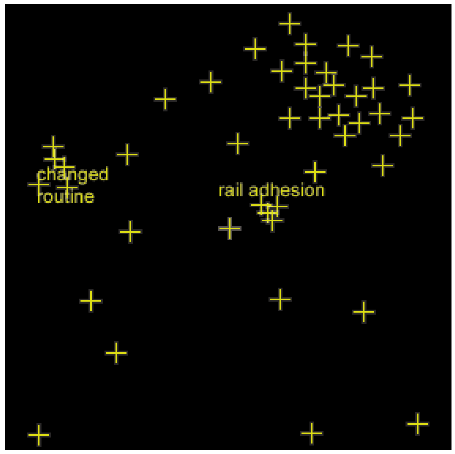

Figure 5: A prototype clustering representation of SPAD data

has pioneered the 'dynamic query' technique that allows a user to interactively specify parameters to a query and see the updated results in real-time (Shneiderman, 1994). Dynamic queries can be applied to most sets of data, but there are still questions of how one displays the set of results selected by the query. There are many possibilities. One example may be starfield displays (Ahlberg & Shneiderman, 1994), which lay out representations of data points in a two-dimensional display. This is illustrated by figure 4.

Dynamic Querying could be applied to a database such as the SPAD data with good effect, to support dynamic discovery of interesting or unusual reports in the data set. The same technique could be used in a geographical layout of the data to filter the information presented on the display, for example to interactively change the lower threshold of SPAD severity that is shown.

4.4 Determining Causality

The primary reason for investigating and cataloguing SPAD incidents is to understand why they happen and put in place appropriate defences against the risk that they pose. Presentations of SPAD data should support drawing connections between different incidents, perhaps automatically. If we assume that a report about a SPAD includes some indication of the causes of the incident, we can apply similar techniques to those applied to document collections. Several systems have been developed that attempt to visualise the relationship between documents in a collection. The Bead system (Chalmers & Chitson, 1992) uses a force-based clustering approach to showing the similarity between documents.

Figure 5 shows another early prototype showing a representation of SPAD data. Again, the actual data is hypothetical for the purposes of illustration. Each cross represents a single SPAD incident, and their distance from each other is a function of their similarity. This is similar to Bead and other force-based layouts. As figure 5 shows, certain close clusters have a summary showing - 'rail adhesion' and 'changed routine'. At the top right, there appears to be a possible weak similarity between a number of incidents, and along the bottom there are three incidents that would seem to be quite unique in their causes.

These document collection visualisation techniques do not, however, support the other tasks well. Bead is designed to show relationships through spatial metaphor, rather than trends or geographical distributions. Again, it will be necessary to combine multiple techniques into a tool that will support investigation tasks well.

5 Conclusion

We have argued that HCI visualisation techniques provide one means of helping incident investigators to detect incident 'hot spots', emmergent trends and common causes in SPAD reports. It is important to emphasise, however, that these visualisations do not provide a panacea to the problems posed by Signals Passed at Danger. The problems of reviewing and interpreting incident data form part of the wider problems that affect organisational learning from previous incidents. For example, Her Majesty's Railway Inspectorate enquiry into the investigation of incidents involving 'Signals Passed at Danger' (SPADs) found that "in some cases greater emphasis was placed on completing a multi-page form than getting to the root cause of the SPAD incident". (Her Majesty's Railway Inspectorate, 1999) The consequences of this were identified in a recent internal report that examine the failure of Her Majesty's Railway Inspectorate to respond adequately to previous problems at the signals which were involved in the Ladbroke Grove rail crash:

> "During the almost five years preceding the Ladbroke Grove accident, there had been at least three occasions when some form of risk assessment analysis on the signaling in the Ladbroke Grove area has been suggested or proposed. The requests were: the Head of Technical Division's letter of 11 November 1996 which requested a layout risk assessment of the re-signaling (paragraph 43); the Field Inspector's letter of 16 March 1998 to Railtrack (paragraph 64); and the Railtrack Formal Inquiry of 1 July 1998 (paragraph 66). In addition there was an earlier request for details of measures taken to reduce the level of SPADs in the area around SN109 recorded in the Head of Technical Division's letter of 1st March 1995 (paragraph 39). None of these requests appear to have been pursued effectively by HMRI." (Her Majesty's Railway Inspectorate, 2000a)

These organisational factors illustrate some of the problems that arise when attempting to validate our work. It is perfectly possible for organisations to misinterpret or ignore the insights that can be provided by Starfield displays or dynamic query interfaces. This is a significant problem. Most HCI validations of these techniques have focussed on short-term comprehension and recall tasks rather than on longitudinal effects on organisational learning. One means of addressing these concerns is to introduce the visualisations at a local or team-based level before attempting to assess their impact on a national level. It is for this reason that we are currently exploring a number of small-scale field trials in which rail staff will be provided with access to some of the visualisations shown in previous pages. Contextual and workplace studies will then be conducted to determine whether or not these novel visualisations contribute to the operators' existing tasks.

Ulitmately, technology will be applied to resolve the hazards that are created by SPADs. This includes the Train Protection and Warning System at a cost of some £310 million and ultimately the Advanced Train Protection system, estimated to cost in the region of £2 billion. Research within the wider field of human factors has, however, shown that such automated systems do not always reduce the number of HCI related incidents (Reason, 1997). The net effect is often to change the nature or many interaction problems. At a more general level, therefore, HCI visualisation techniques are essential if investigators are to identify and respond to the changing nature of safety-critical interaction problems. Without this application of HCI techniques there is a danger that we may fail to learn the lessons of past failures.

6 Acknowledgments

Thanks are due to the members of Glasgow Accident Analysis Group. This work has been supported by EPSRC grant GR/M98302; Communicating Knowledge about Accidents from Synthesised Web Sites.

References

Ahlberg, C. & Shneiderman, B. (1994), Tight Coupling of Dynamic Query Filters with Starfield Displays, *in Proc. CHI'94, ACM Conference on Human Factors in Computing Systems*, pp.313–317.

Becker, R. A., Eick, S. G. & Wilks, A. R. (1995), "Visualizing Network Data", *IEEE Transactions on Visualization and Computer graphics* pp.16–28.

Card, S. K., Mackinlay, J. D. & Shneiderman, B. (eds.) (1999), *Readings in Information Visualization*, Morgan Kaufmann.

Chalmers, M. & Chitson, P. (1992), Bead: Explorations in Information Visualisation, *in Proc. ACM SIGIR'92*, ACM Press, Copenhagen, pp.pp. 330–337.

Her Majesty's Railway Inspectorate (1999), Report on the inspection carried out by HM Railway Inspectorate during 1998/99 of the management systems in the railway industry covering signa ls passed at danger, Technical Report, Health and Safety Executive, London, United Kingdom. http://www.hse.gov.uk/railway/spad-01.htm.

Her Majesty's Railway Inspectorate (2000a), Internal Inquiry Report: Events leading up to the Ladbroke Grove rail accident on 5 October 1999, Technical Report, Health and Safety Executive, London, United Kingdom. http://www.hse.gov.uk/railway/paddrail/inq-03.htm.

Her Majesty's Railway Inspectorate (2000b), Signals Passed at Danger (SPAD)s Report For November 2000, Technical Report, The Health and Safety Executive. http://www.hse.gov.uk/railway/spad/nov00.htm.

Lagus, K., Honkela, T., Kaski, S. & Kohonen, T. (1996), Self-organizing maps of document collections: A new approach to interactive exploration, *in* E. Simoudis, J. Han & U. Fayyad (eds.), *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, AAAI Press, Menlo Park, California, pp.238–243.

Reason, J. (1997), *Managing the Risks of Organizational Accidents*, Ashgate.

Shneiderman, B. (1994), “Dynamic Queries for Visual Information Seeking”, *IEEE Software* **11**(6), 70–77.

Shneiderman, B. (1996), “The eyes have it: A task by data type taxonomy for information visualizations”, *Proc. 1996 IEEE Conference on Visual Languages* .