

The Elephant in the Room: Normal Performance and Accident Analysis

Paul M. Salmon¹(✉), Natassia Goode¹, Erin Stevens¹, Guy Walker²,
and Neville A. Stanton³

¹ Centre for Human Factors and Sociotechnical Systems,
University of the Sunshine Coast, Locked Bag 4,
Maroochydore DC, QLD 4558, Australia
psalmon@usc.edu.au

² School of the Built Environment, Heriot-Watt University,
Edinburgh EH14 4AS, UK

³ Transportation Research Group, Faculty of Engineering and the Environment,
University of Southampton, Highfield, Southampton SO17 1BJ, UK

Abstract. Accidents, accident causation, and accident prevention remain key themes within human factors and ergonomics research efforts worldwide. Accordingly, there are a range of well-developed models of accident causation and various methodologies to support accident analysis efforts. State of the art models propose a number of features of accident causation that go beyond operator errors and failed defenses. Once such feature now widely accepted is the notion that ‘normal performance’ plays a role in accidents; that is everyday behaviors not deemed to be errors or failures at the time of occurrence, are implicated in causal networks. Despite this, it is questionable whether our accident analysis methodologies are equipped to identify normal performance and its role in accidents. This paper examines this, reviewing current state of the art accident analysis methods along with their previous applications. It is concluded that, of the three methods reviewed, only one (Accimap) is currently capable of considering normal performance (at least without reclassifying it as a failure or error of some sort). The implications for accident analysis methodologies and practice are discussed and future methodological requirements are articulated.

1 Introduction

The human factors problem space is characterized by complex accidents, how and why they occur, and how organizations can prevent similar occurrences moving forward. Accordingly, there is a long history of accident causation models and accident analysis methods, with each new paradigm providing more explanatory power than its predecessor. So-called systems models of accident causation are now widely accepted [1–3], and there are a range of methods that enable accidents to be analyzed from this perspective [2, 4].

There is no doubt that significant progress has been made in understanding accidents, although it is acknowledged that we still do not fully understand them [5].

Further, it has been suggested that little headway is being made in reducing accidents in many industries [6]. Accident analysis methods have been implicated in this lack of progress. Leveson [6], for example, suggests one explanation is that our methods do not fully uncover the underlying causes of accidents.

What we do know about the accidents that occur in complex sociotechnical systems is that they are systems phenomena. The so-called systems approach has a long legacy in safety science, from the foundational work of Heinrich [7] through to the evolution of a number of more recent accident causation models and analysis methods [2, 3, 8]. Accidents are now known to be emergent properties of complex sociotechnical systems, just as safety is. Both safety and accidents are emergent properties arising from non-linear interactions between multiple components distributed across complex sociotechnical systems [2, 3]. As it stands currently, the most up-to-date methods and approaches rely on a form of systems thinking closely linked to current debates in the science of complexity. It is precisely this form of thinking, and the evolution of it, that brings the methods we use into question.

One of the advances provided by state of the art models [1–3] centers around the idea that the behaviors underpinning accidents do not necessarily have to be errors, failures or violations. As Dekker [1] points out:

“Systems thinking is about accidents that are more than the sum of the broken parts. It is about understanding how accidents can happen when no parts are broken, or no parts are seen as broken” (p. 35).

This provides an advance over popular models [e.g. 8], which tend to subscribe to the idea that failure leads to failure [1], and provides two key tenets. First, ‘normal’ performance plays a role in accident causation, and second, accidents arise from the very same behaviors and processes that create safety [1, 5]. What are apparently normal day-to-day behaviors can interact in a way that leads to adverse events. In fact, it is only with hindsight and an investigation methodology that behaviors are treated as failures and so on. These normal behaviors include workarounds, improvisations, and adaptations [1], but may also just be normal work behaviors consistently undertaken to get the job done.

Theoretical advances such as this have important implications for the methodologies applied to understand accidents. It is imperative that our analyses of accidents are akin to the models that describe the mechanisms involved in accident causation. That is, understanding accidents requires appropriate methodologies that reflect how contemporary models think about accident causation. The tenets described above provide an interesting shift in the requirements of accident analysis methodologies, which in turn raises questions regarding the capabilities of current approaches. Dekker [15] describes how organizations should not look for known holes or the problems that appear in incident reporting data or safety management systems. Instead he argues that they should look in the places where there are no holes – in other words, normal work. This raises critical questions – do our accident analysis methodologies have the capability to incorporate normal performance into their descriptions of accidents? Do we currently incorporate normal work into accident analyses? And if we do, are we misclassifying it as errors, failures, inadequacies? Furthermore, do our methods push analysts toward the classification of errors, failures, and violations?

It is therefore important to question whether the accident analysis methods currently employed by researchers and practitioners recognize the two key tenets described above. The aim of this paper is to investigate this by evaluating three accident analysis methods that are considered to be state of the art and are the most commonly applied within the academic literature: Accimap [3], STAMP [2] and HFACS [9]. The investigation involved reviewing each methodology, its theoretical underpinning, and applications of the method presented in the academic literature.

2 Accident Analysis Methods

2.1 Accimap

The Accimap method accompanies Rasmussen's now extremely popular risk management framework [3]. Accimap provides a methodological framework for describing accidents in terms of contributory factors and the relationships between them, enabling a representation of the full network of contributory factors involved. It does this by decomposing systems into six levels across which analysts place the decisions and actions that enabled the accident in question to occur (although the method is flexible in that the number of levels can be adjusted based on the system in question). Cause-effect relationships between the decisions and actions are subsequently mapped onto the diagram to show the relationships between contributory factors within and across the six levels. The output is an exhaustive description of the network of contributory factors involved in the incident under analysis. A notable feature of Accimap is that it does not provide analysts with taxonomies of failure modes; rather, analysts have the freedom to incorporate any behavior deemed to have played a role in the accident in question.

2.2 STAMP

The Systems Theoretic Accident Model and Process method (STAMP) views accidents as resulting from the inadequate control of safety-related constraints [2], arguing that they occur when component failures, external disturbances, and/or inappropriate interactions between systems components are not controlled [2, 11]. Leveson [2] describes various forms of control, including managerial, organizational, physical, operational and manufacturing-based controls. STAMP uses a 'control structure' modelling technique to describe complex systems and the control relationships that exist between components at the different levels. A taxonomy of control failures is then used to classify the control failures that played a role in the incident under analysis. An additional component of STAMP involves using systems dynamics modeling to analyze the behavior of the system over time. This enables the interaction of control failures to be demonstrated along with their effects on performance.

2.3 HFACS

Although not based on contemporary models of accident causation, the Human Factors Analysis and Classification System (HFACS) [9] remains highly popular [e.g. 10, 11]. HFACS is a taxonomy-based approach to accident analysis that provides analysts with taxonomies of error and failure modes across four system levels based on Reason's Swiss cheese model of organizational accidents: unsafe acts, preconditions for unsafe acts, unsafe supervision, and organizational influences. Although developed originally for use in analyzing aviation incidents, the method has subsequently been redeveloped for use in other areas including mining, maritime, road safety, and healthcare. Later versions of the method have extended the levels to incorporate an 'external influences' level which considers failures outside of organizations such as legislation gaps, design flaws, and administration oversights [e.g. 12].

3 Evaluation of Methods

The three methods were subjected to review by the authors with regard to the extent to which they are capable of identifying and representing normal performance in accident analyses. The review was performed based on the co-authors' knowledge of accident causation models, their experiences in applying the methods and an examination of the previous applications published in the academic literature.

3.1 Theoretical Underpinning and Methodological Constraints

Table 1 provides an overview of the extent to which each methodology and its underpinning model of accident causation consider the role of normal performance in accidents, followed by a judgment as to whether the method could be extended to do so.

In the case of the models underpinning each method, all three do acknowledge the role of normal performance in accidents, albeit to differing degrees. Rasmussen's risk management framework acknowledges the role of normal performance in accidents (e.g. the role of financial and production pressures). Leveson's control theory approach embodies a focus on the interactions between "perfectly functioning components" [6, p. 56]. Finally, although Reason's Swiss cheese model does have a strong focus on unsafe acts, failed defenses, and latent failures, his description of latent conditions does incorporate conditions "not necessarily the products of bad decisions" [8, p. 11].

The second column in Table 1 asked whether the methods themselves provide the capability to consider normal performance. In these authors' opinion only Accimap fully achieves this, as it does not provide taxonomies of failure modes, allowing analysts to include anything that they consider contributed to the accident in question. Both HFACS and STAMP provide taxonomies that lead analysts to explicitly consider failures. In the case of HFACS this includes error and failure modes, whereas STAMP's taxonomy includes control failures. It is acknowledged that aspects of normal behavior are incorporated in the HFACS taxonomy, such as weather; however, the predominant focus is on errors and failures.

Table 1. Extent to which methods and underpinning models do and could consider normal performance.

Method	Does underpinning model consider Normal Performance (NP)?	Does method consider NP?	Have previous applications considered NP? Has NP been classified as errors, failures etc.?	Could the method be extended to consider NP
Accimap	Yes – Rasmussen’s risk management framework acknowledges the role of normal performance in accidents.	Yes (method does not adopt a taxonomic approach and does not constrain analysts to failures only).	Yes. Yes	Yes – as the method is generic and does not use failure mode taxonomies normal performance can be included.
HFACS	Yes – Reason’s Swiss Cheese model focuses on unsafe acts and failed defenses but also acknowledges that latent conditions may not arise from failures, errors etc.	No – HFACS uses a taxonomy of errors and failure modes.	No. Applications have classified the errors and failures involved in accidents. Yes	Yes – although this would require a development of a taxonomy of normal behaviors
STAMP	Yes – Leveson’s model considers normal performance and its role in control failures	No – the method uses a taxonomy of control failures.	No. Applications have classified control structures and the control failures across them Yes	Yes – the systems dynamics component of STAMP can incorporate normal performance and show the links to control failures

The next question looked at the extent to which previous applications of the methods, or at least those published in the academic literature, included normal performance in their analyses. In addition, we asked whether previous applications show

evidence that normal performance has been considered but classified as errors and/or failures. In response to the former, only Accimap achieved a yes. Accimap applications have incorporated normal performance in accident descriptions, or at least have incorporated behaviors that have not been explicitly defined as failures or errors. Salmon et al. [13], for example, incorporated aspects of normal system performance such as financial and production pressures. Further, the decisions and actions included in Accimap are not always classified as failures or errors, and normal performance resulting from emergence is often included. HFACS applications have focused almost exclusively on errors and failure modes, whereas STAMP has focused on the control failures involved in accidents. In response to the latter, previous applications suggest it is highly likely that normal performance has been incorporated in analyses but classified as errors, failures etc.

Finally, the extent to which the methods could be extended to more explicitly focus on normal performance was examined. All three could be extended, however different levels of adaptation would be required. Accimap does not actually require any development, since it is generic and does not use a taxonomy of any sort. Rather, only explicit guidance for analysts to consider non-failures is required. HFACS would require development of a taxonomy of normal behaviors, which in turn would require substantial investigation in terms of which normal behaviors are involved in accidents. Finally, STAMP would require development of a taxonomy of control behaviors.

3.2 Accimap Revisited

The examination of the three methods leads to the conclusion that Accimap provides an analysis methodology that is the most consistent with contemporary models of accident causation. The next part of this case study involved examining a previous application of Accimap in order to exemplify the extent to which normal performance has and could be considered. For this purpose, a recent Accimap analyses undertaken by the lead author was examined: the Mangatepopo gorge walking incident [13]. The analysis was reviewed to determine if the contributory factors included were representative of normal performance rather than errors and failures. Following this, the analysis was revised so that each contributory factor was expressed as normal behavior, rather than a failure of some sort.

The Mangatepopo gorge walking tragedy occurred on the 15th April 2008 when a number of college students and their teacher drowned during a gorge walking activity in the Tongariro National Park, New Zealand. The group and their instructor experienced a flash flood whilst in the gorge and became trapped on a small ledge above the water. Fearing the group would be washed off the ledge, the instructor attempted to evacuate them from the gorge by first entering the river herself and then extracting the students in pairs as they followed her into the water in 5-min intervals. Only the instructor and two students managed to get out of the river as intended, with the remaining eight students and teacher being swept downstream and then over a spillway. Six students and their teacher eventually drowned, with only 2 of those swept over the spillway surviving.

Salmon et al. [13] used Brookes et al.'s [14] investigation of the Mangatepopo incident to conduct a comparison of Accimap, HFACS and STAMP. In Fig. 1 a revised version of the original Accimap diagram is presented. Those factors in the Accimap that were not prefixed with descriptors such as 'failed to', 'failure', 'inadequate' etc. are shown via bold outlines. In Fig. 2, a revised version of the Accimap is presented in which the contributory factors originally classified as failures and could conceivably be characteristic of normal performance, are highlighted. In addition, each description has been revised to remove items such as 'failed to', 'failure', 'inadequate' etc. For example, the contributory factor originally described as 'Poor adventure program design' has now been described as 'Adventure program'.

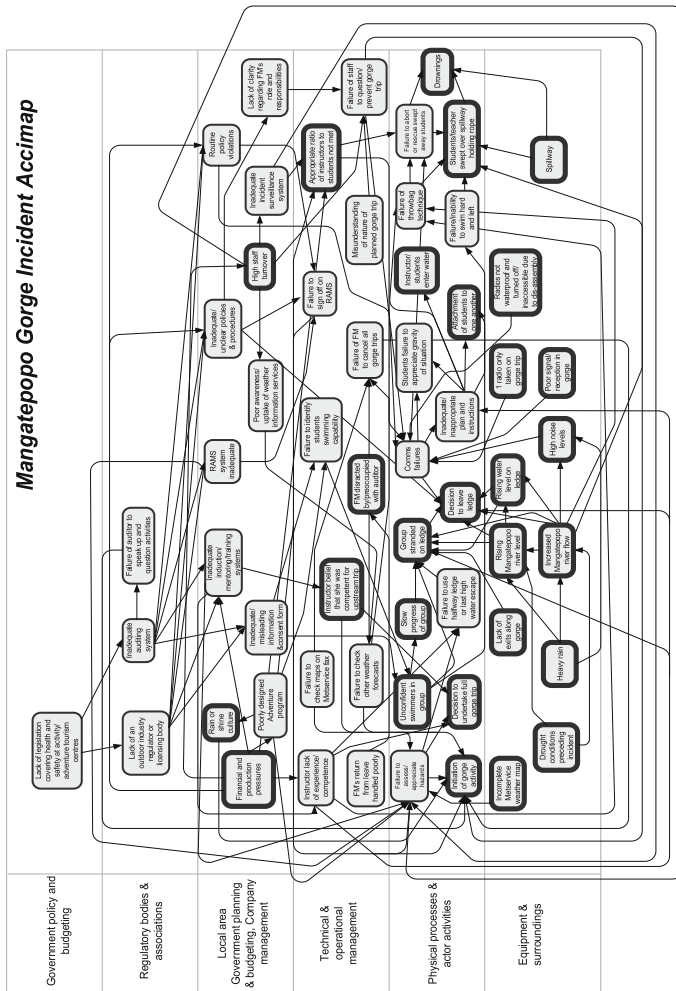


Fig. 1. Original Accimap from Salmon et al. [13]. Factors outlined in bold are those that were not classified as failures or errors.

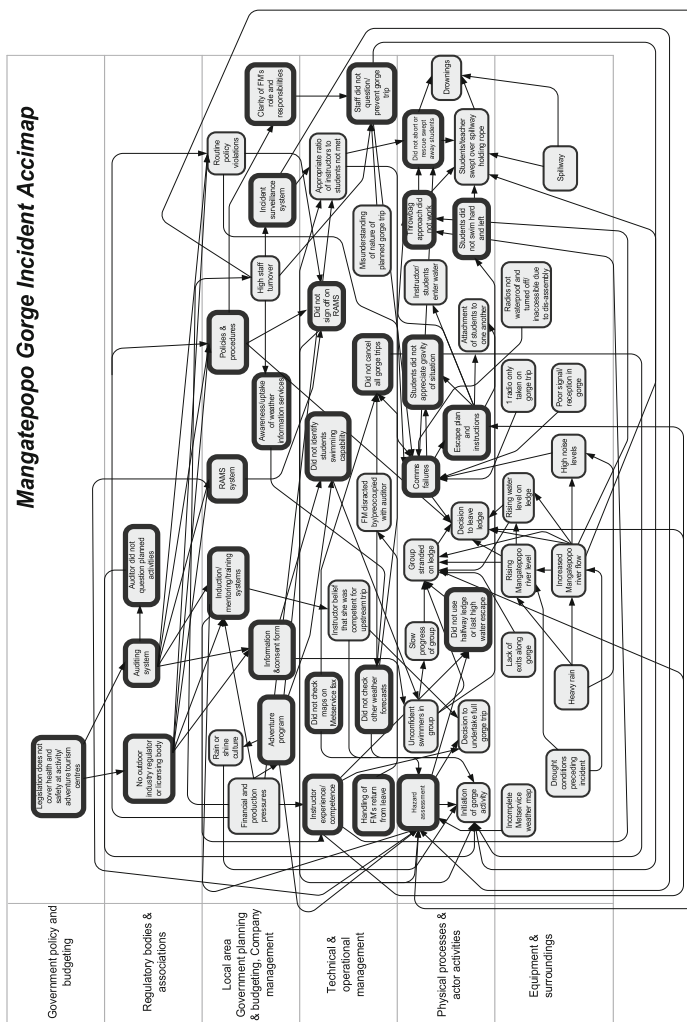


Fig. 2. Revised Accimap with failure descriptors removed. Factors outlined in bold are those that were modified to represent normal performance.

As suggested in Table 1, it is clear from Figs. 1 and 2 that Accimap enables analysts to represent the role of normal performance in accidents. This is demonstrated through the inclusion of both failures and normal performance in Fig. 1. Figures 1 and 2 have some important implications. First, from Fig. 1, it is clear that the absence of a taxonomy of failure modes is enabling analysts to incorporate normal performance into Accimap analyses. This is seen in Fig. 1 through the inclusion of contributory factors that are characteristic of normal everyday performance in the led outdoor activity context. These include factors such as ‘financial and production pressures’, ‘rain or shine culture’, ‘high staff turnover’, and ‘initiation of gorge activity’. All of which seemingly represent normal behaviors that had been ongoing and present in the system for sometime.

Second, from Fig. 2 it is clear that a high proportion of the contributory factors identified were deemed to be failures (and this is characteristic of Accimap analyses generally). In this case half of the contributory factors identified were classified as a failure of some sort. It is not possible in this case to reinvestigate whether these factors actually represent failures or normal performance; however, it is questionable whether some of the factors reported in the original Accimap actually do represent failures. A clear example is the contributory factor originally labelled 'Staff failed to question/prevent gorge trip'. This refers to the fact that staff at the activity center did not step in to question or prevent the planned gorge walking activity despite the adverse weather conditions. The activity center had a strong 'rain or shine' culture (i.e. they prided themselves on conducting activities come rain or shine) and the group in question was designated for a water-based activity on the day in question. Normal behavior in such circumstances was probably not to question planned activities. In fact questioning the activity could conceivably have been seen as abnormal behavior. Rather, in line with the rain or shine culture the expectation was that activities would go ahead as planned even during extreme weather.

An important question to ask then is whether the appropriate number of failures are being classified, or whether analysts have a tendency to classify normal performance as failures? Dekker [1] talks at length about the inappropriate focus on identifying broken components. It may be that Accimap users are classifying decisions, actions, and behaviors as broken when in fact they are far from it. Although an analysis of all Accimap analyses was not undertaken for this paper, various authors have suggested that the latter may be the case in accident analysis generally [1].

4 Discussion

Modern day accident analysis efforts have to consider normal performance [1–3, 15]. Not doing so potentially provides one explanation for why satisfactory progress is not being made in preventing accidents across many safety critical domains. The aim of this paper was to examine whether contemporary accident analysis methods are consistent with contemporary accident causation models in that they can be used to represent the role of normal performance in accidents (as opposed to merely errors and failures). Based on an examination of Accimap, STAMP and HFACS, it is clear that the answer for the latter two methods (at least in their current format) is no, and even for Accimap more work may be required with regard to how analysts apply the method. Both HFACS and STAMP provide taxonomies of error and failure modes that are used to classify the errors and failures involved in accident scenarios, which in turn means that there is little scope for analysts to include behaviors other than those deemed to have been failures of some sort. Aside from factors such as the weather, there is no opportunity for analysts to incorporate normal behaviors in their descriptions of accidents – they have to force fit events into one of the error or failure modes provided. The output is a judgment on what errors or failures combined to create the accident under analysis. Whilst this is inappropriate given current knowledge on accident causation, a worrying consequence may be that the normal behaviors that contribute to accidents are not picked up by organizations during accident analysis efforts. This may impact

accident prevention activities by providing a false sense of security that nothing is failing and thus nothing needs fixing (apart from error producing human operators). A more sinister implication is that organizations who apply methods such as HFACS may not have sufficient understanding of accidents to prevent them. Certainly the idea that the aviation domain has no analysis of the role of normal performance in air crashes is a sobering one. Extending HFACS and STAMP to incorporate analyses of normal performance in accidents is therefore a pressing requirement.

Accimap, on the other hand, does not use a taxonomy of failure or error modes and so enables analyst to incorporate normal performance and to show its relationship with other behaviors. There is freedom for analysts to include any form of behavior in the network of contributory factors. Despite this, however, Accimap descriptions tend to incorporate contributory factors prefixed with descriptors such as ‘failure to’, ‘lack of’ or ending with ‘error’. A pressing question here then is the extent to which the failures described in Accimap analyses actually represent failure or are in fact normal behaviors. A further important line of further inquiry is the extent to which researchers and practitioners understand the need to incorporate ‘normal’ behaviors in accident analyses. A downside of Accimap’s flexibility, however, is that there are no prompts for analysts to look beyond failures. A step-by-step procedure specifying this would be beneficial.

The conclusion then from this initial investigation is that there is room for methodological improvement in the area of accident analysis. Not all state-of-the-art methods are consistent with our current understanding of accident causation. Further, even for the methods that are, it is questionable whether they are being used in a manner consistent with contemporary models of accident causation. This paradox represents a key issue for researchers and practitioners and for safety science generally. On the one hand there is now a widespread understanding that the role of normal performance in accidents is apparent and needs to be understood [1–3]. On the other hand, however, accident analysis efforts, regardless of domain, do not seem to be dealing particularly well with this feature. At best this means our understanding of accidents may be incomplete. At worst, it may mean that the countermeasures proposed based on accident analyses are inappropriate and doomed to fail. Dekker [15] rightly points out that we need to look where there are no holes; equally, we need methods that do not dig holes or take us down them.

Acknowledgements. Paul Salmon’s contribution to this article was funded through his current Australian Research Council Future Fellowship (FT140100681).

References

1. Dekker, S.: *Drift into Failure: From Hunting Broken Components to Understanding Complex Systems*. Ashgate Publishing, Ltd., Farnham (2011)
2. Leveson, N.G.: A new accident model for engineering safer systems. *Saf. Sci.* **42**(4), 237–270 (2004)
3. Rasmussen, J.: Risk management in a dynamic society: a modelling problem. *Saf. Sci.* **27**(2/3), 183–213 (1997)

4. Svedung, I., Rasmussen, J.: Graphic representation of accident scenarios: mapping system structure and the causation of accidents. *Saf. Sci.* **40**, 397–417 (2002)
5. Hollnagel, E.: *Barriers and Accident Prevention*. Ashgate, Aldershot (2004)
6. Leveson, N.G.: Applying systems thinking to analyze and learn from events. *Saf. Sci.* **49**, 55–64 (2011)
7. Heinrich, H.W.: *Industrial Accident Prevention: A Scientific Approach*. McGraw-Hill, New York (1931)
8. Reason, J.: *Human Error*. Cambridge University Press, New York (1990)
9. Wiegmann, D.A., Shappell, S.A.: *A Human Error Approach to Aviation Accident Analysis. The Human Factors Analysis and Classification System*. Ashgate Publishing Ltd., Burlington (2003)
10. Daramola, A.Y.: An investigation of air accidents in Nigeria using the human factors analysis and classification system (HFACS) framework. *J. Air Transp. Manag.* **35**, 39–50 (2014)
11. Mosaly, P.R., Mazur, L., Miller, S.M., Eblan, M.J., Falchook, A., Goldin, G.H., Marks, L.B.: Assessing the applicability and reliability of the human factors analysis and classification system (HFACS) to the analysis of good catches in radiation oncology. *Int. J. Radiat. Oncol.* **90**(1), S750–S751 (2014)
12. Chen, S.T., Wall, A., Davies, P., Yang, Z., Wang, J., Chou, Y.H.: A human and organisational factors (HOFs) analysis method for marine casualties using HFACS-maritime accidents. *Saf. Sci.* **60**, 105–114 (2013)
13. Salmon, P.M., Cornelissen, M., Trotter, M.: Systems-based accident analysis methods: a comparison of Accimap, HFACS, and STAMP. *Saf. Sci.* **50**(4), 1158–1170 (2012)
14. Brookes, A., Smith, M., Corkill, B.: Report to the trustees of the Sir Edmund Hillary Outdoor Pursuit Centre of New Zealand: Mangatepopo gorge incident, 15 April 2008 (2009)
15. Dekker, S.: *The Field Guide to Understanding Human Error*, 3rd edn. Ashgate Publishing Ltd., Farnham (2014)