

LANCASTER
UNIVERSITY

Computing
Department



Human Factors in Requirements Engineering

A Survey of Human Sciences Literature Relevant to the Improvement of Dependable Systems Development Processes

Stephen Viller, John Bowers and Tom Rodden

Cooperative Systems Engineering Group

Technical Report Ref: CSEG/8/1997

http://www.comp.lancs.ac.uk/computing/research/cseg/97_rep.html

To appear in the "Interacting with Computers", Special Issue on Human Error and Systems Development

ABSTRACT

Requirements Engineering (RE) is an inherently social process, involving the contribution of individuals working in an organizational context. Furthermore, failures in the RE process will potentially lead to systematic failures in the products that are produced as a result. Consequently, the RE process for dependable systems development should itself be considered as a dependable process, and therefore subject to greater scrutiny for vulnerabilities to error. Research on human error has typically focused on the work of individual actors from a cognitive perspective. This paper presents a survey which broadens the view on what contributes to human error by also examining work from the social and organizational literature. This review was conducted to inform efforts to improve the systems development process for dependable systems, and in particular their requirements engineering process.

CSEG, Computing Department, Lancaster University, Lancaster, LA1 4YR, UK

Tel: +44-1524-65201 Ext 93799; Fax: +44-1524-593608; E-Mail: julie@comp.lancs.ac.uk

<http://www.comp.lancs.ac.uk/computing/research/cseg/>

Human Factors in Requirements Engineering

A survey of human sciences literature relevant to the improvement of dependable systems development processes

Stephen Viller¹,
viller@comp.lancs.ac.uk,

John Bowers²,
bowers@hera.psy.man.ac.uk,

Tom Rodden³
tom@comp.lancs.ac.uk

^{1,3} *Computing Department
Lancaster University
Lancaster LA1 4YR
Fax: +44 1524 593608*

² *Department of Psychology
University of Manchester
Manchester M13 9PL
Fax: +44 161 275 2588*

Requirements Engineering (RE) is an inherently social process, involving the contribution of individuals working in an organizational context. Furthermore, failures in the RE process will potentially lead to systematic failures in the products that are produced as a result. Consequently, the RE process for dependable systems development should itself be considered as a dependable process, and therefore subject to greater scrutiny for vulnerabilities to error. Research on human error has typically focused on the work of individual actors from a cognitive perspective. This paper presents a survey which broadens the view on what contributes to human error by also examining work from the social and organizational literature. This review was conducted to inform efforts to improve the systems development process for dependable systems, and in particular their requirements engineering process

Keywords: human error, group performance, process losses, organizational failure, process improvement, dependable systems development.

Introduction

Work in the field of human error has typically focused on operators of safety-critical equipment, such as nuclear power plant controllers, and of the design of the human-machine interfaces in such settings. Limited consideration has been given to wider system development issues. Similarly, researchers and practitioners in the field of Dependable Systems are concerned with the design of computer-based systems which are intended to be operated in situations where the consequences of failure are potentially catastrophic. For example, the failure of a safety-critical system may cause great harm to people, property, or the environment. The work reported on in this paper is motivated by the need to ‘push back’ these concerns with the operation and design of dependable systems to the process by which they are developed.

It is widely recognised that the requirements phase of systems design is particularly problematic (Brooks Jr., 1987; Christel and Kang, 1992; Davis, 1993; Sommerville and Sawyer, 1997). Requirements may be expressed in a number of different ways, may come from a variety of sources, and the likelihood of conflicts is high. They also are liable to change throughout the development process, sometimes as a side effect of the process itself as the stakeholders’ needs and desires shift when exposed to early prototypes, and so on. This means that the requirements as specified in the project documentation may differ markedly from what is actually needed or expected.

Compared with other phases of systems development, there also exists a much stronger need to get requirements right. Errors in requirements lead directly to systems which fail, either because they do not function as desired, or because they do not match the actual needs of users and other stakeholders. Errors in requirements can remain latent until very late in the design process, and the longer that errors remain uncovered, the more they cost to rectify. When the systems being developed are to be used in any safety-critical setting, then the costs and consequences of failure are even greater, as is the motivation for any efforts to ensure that the requirements are error-free.

Whilst efforts to detect and rectify errors in Requirements Engineering (RE) and the whole of the development process are a necessity, the nature and cost of errors in requirements makes a strategy of avoidance rather than detection a more attractive prospect. The benefits of such an approach are primarily that the amount of rework can be reduced to a minimum, along with related savings in cost and time to completion of the system.

There is also a broadening consensus regarding the nature of RE as a social, as well as technical, process involving a variety of stakeholders engaging in diverse activities throughout (Bowers and Pycock, 1994; Goguen, 1993; Jirotko and Goguen, 1994; Quintas, 1993). Many of the specific details of the process followed for a given product will depend upon the nature of the product itself, the application domain, similarities and differences to existing products developed by the organization, and so on. When these variations are combined with an intense production pressure to release products on time, the importance of human skill and judgement in managing the contingencies, and human flexibility and artfulness in making RE processes work (sometimes in spite of the methods followed (Anderson *et al.*, 1993; Rodden *et al.*, 1994)) becomes readily apparent.

The RE process for any non-trivial system, therefore, will necessarily involve a number of individuals working both together and in isolation on a variety of manual and computer supported processes. Rarely (and almost never in a commercial context) is RE engaged in solely by an isolated individual working alone. Overwhelmingly, requirements are elicited, analysed and documented by teams of engineers working as part of a larger development team. Accordingly, the contribution of each individual needs to be coordinated with other team members as the RE process unfolds. Even when individuals do work alone (e.g. when one person assumes sole responsibility for drafting a requirements document), their contributions are oriented towards the contributions of others and specifically designed to mesh with them. Furthermore, in commercial contexts, relations that a supplier organization has with clients will always be borne in mind in some way while, for example, requirements documents are being developed for them.

These are intended as simple, non-controversial observations but it is worth emphasising them: RE is a form of cooperative work, involving the coordination of *individual work* within a *team or group* of engineers working in an *organizational context*. Further, the human-intensive nature of the RE process means that many of the errors which are attributed to this phase of systems development are of a human nature. If improvements to the RE process are to be proposed, therefore, it is important to first of all develop a good understanding of the human activities which are inherent in it.

Following the above observations, this paper approaches the understanding of human activity in RE from three perspectives, all of which are relevant at different points in the process. The first of these is from an individual perspective, concerned with the actions and tasks of requirements engineers working in isolation, or on individual tasks which contribute to a larger team or organizational goal. The source of research in this field is cognitive psychology, and more specifically *human error*, which has primarily approached problems of errors in high risk

activities from a cognitive standpoint. The first section covers human error, and the second section examines the related phenomena known as *violations*. The third section broadens this individual perspective to the consideration of people working not in isolation, but together in groups or teams. Research into how people perform in groups has a long tradition in psychology, and findings from this perspective come primarily from social psychology, but also from sociology. Having broadened the perspective once, it would be wrong to stop there when the enquiry can take in a further, wider reaching concern. The fourth section considers the way in which individuals and groups, teams, or other units function within organizations and the wider society at large. Here, the relevant research originates in sociology and in organizational and political studies. The purpose of studying the literature in these three areas is to inform, and provide the basis for, a human-centred approach to process improvement for RE. The final section briefly describes the way in which the content of this survey has been integrated into a method for the process improvement method called PERE.

Errors in individual work

The largest body of research on ‘human error’ (Reason, 1990; 1994) has its roots in cognitive psychology and cognitive understandings of peoples’ interaction with technology. Work in this area has typically focused on workplace settings such as nuclear power plant control rooms and on operational risks and operator errors in such environments. Rather less work on human error specifically concerns the use of computer-based systems, and there is even less devoted to the process of their development.

A major distinction to arise from this work is between different ‘levels’ of cognitive activity: e.g. skill-based, rule-based, and knowledge-based in Rasmussen’s (1983) formulation. These in turn lead to a number of error classes: skill-based slips and lapses, rule-based and knowledge-based mistakes. Skill-based slips and lapses happen during routine, familiar work, which requires little attention in order to be achieved. In RE, this would be typified by mundane activities, involving everyday skills (e.g. typing, reading, filing, etc.). Rule-based mistakes are related to errors in the plan of action when working in previously encountered situations. They can result from the application of ‘bad’ rules, or the misapplication of ‘good’ rules. In RE, the application of generic solutions can be prone to this type of error. Knowledge-based mistakes arise when working in novel situations, where no existing rule or plan can be applied and attempts are made to apply analogous rules which have worked in similar situations. This describes a great deal of RE work where either there is no previous system which is relevant to the current development, or where the personnel involved are inexperienced in the domain of application.

In human factors work on human error an important distinction is usually made between slips and lapses on the one hand and mistakes on the other. The distinction hinges on recognising two different ways in which planned action can fail:

- The plan may be adequate but the actions associated with realising or executing the plan do not go as intended. A *slip* occurs when an action is incorrectly performed or when some similar or related action is performed instead. A *lapse* occurs when an action which should be performed is omitted.
- The actions go entirely as planned but the plan itself is faulty in that it does not achieve its desired outcome. These failures are *mistakes*.

Mistakes, then, are typically errors of plan formulation, while slips and lapses are typically errors of plan execution. This distinction is often held to be important because slips and lapses are believed to have different origins and be influenced by different factors than mistakes. In the

terms of Rasmussen (1983), slips and lapses are to be understood in terms of the exercise of pre-existing *skills*, while mistakes relate to the inappropriate application of *rules* or prior *knowledge*. Accordingly, strategies which might be undertaken to remedy mistakes may not be effective in alleviating the risk of slips and lapses. On this argument, recognising different types of error is important to formulating appropriate strategies for anticipating and preventing error (Reason, 1990). The following sections briefly review Reason's work on Human Error in terms of this classification.

Slips and Lapses

Slips and lapses are errors which result from some failure in the execution and/or storage stage of an action sequence, regardless of whether or not the plan which guided them was adequate to achieve its object. (Reason, 1990: p.9)

Slips and lapses often occur during performance of routine familiar tasks in their usual environments. This means that the more skill people develop in performing a particular task, the more vulnerable they become to making slips and lapses during execution of that task. Slips and lapses are very typically associated with some form of attentional distraction: while performing a routine task, one is distracted by something else and slips up. Alternatively, they can be provoked by unexpected change in the otherwise familiar environment. Reason (1994) categorises slips and lapses as follows.

Recognition failures.

These lead to slips which involve the incorrect identification or non-identification of details important to the plan. *Misidentification* is typically caused by confusion between correct and incorrect objects arising through their similarity. *Non-detection* (or *omissions* or *false negatives*) can be due to many reasons. For example: operator fatigue, interruption and misleading expectancies can all lead to failures due to non-detection. *False positives* (wrongly identifying problems which are not actually present) can also lead to action slips. Often safety-critical processes are designed so as to be relatively tolerant of false positives but this is not always the case, especially where corrective actions undertaken on the basis of a false positive identification have considerable costs in their own right.

Attentional failures.

Slips can arise both through insufficiently attending to how planned action is progressing and through over attending. Failures due to inattention constitute the most common source of error in this category but quite familiar classes of error can arise when one attends to something at moments where it is better to just execute a highly practised sequence (e.g. concentrating on one's precise leg movements is not advised when walking down stairs!).

- **Inattention slips.** While these can take many forms, they are almost all due to attention being captured by some other detail of the situation or some change in it. This kind of distraction or pre-occupation can commonly manifest itself as: *Branching slips*, *Overshoots and undershoots*, *Omissions following interruptions*, and *Unawareness that the plan is inappropriate*.
- **Slips through over-attention.** Although these errors are perhaps less familiar, people are still prone to make them. Interestingly, they can often occur in an attempt to compensate for an error (or a 'near-miss') due to earlier inattention. Two sub-classes can be identified: *Mistimed checks* and *Disrupting well practised actions*.

Memory failures

Memory failures typically lead to lapses—the omission of some action important to a planned action sequence. It is often difficult to distinguish between different explanations of a lapse. If something is forgotten, it is often equally explicable as a retrieval or as an encoding failure. Some commonly occurring errors due to memory failures include: *forgetting intentions*, *forgetting or misremembering preceding actions*, *encoding failures*, *retrieval failures*, and *reconstructive memory errors*.

Selection failures

Even if opportunities for action have been correctly identified, if the actor is attending appropriately to unfolding events, if plans and prior actions are appropriately remembered, it is still possible to select the incorrect action out of the range of alternatives that might be available. These failures often occur when an actor is having to engage in several different planned sequences simultaneously. The following sub-classes can be identified: *Multiple side-steps*, *Misordering*, *Blending actions from two current plans*, *Carry-overs*, and *Reversals*

What distinguishes slips and lapses from other forms of human error is the nature of the task being performed. The more routine and familiar the activity—the more we operate in *skill-based* behaviour—the more prone we are to making errors of these types. This has implications for the design and improvement of processes which have or may have a routine or repetitive element.

When conditions are less familiar or routine, skill-based behaviour becomes less of a factor, and we must turn to *rule-* or *knowledge-based* behaviour in order to complete tasks of increasing degrees of novelty. Human errors in these situations cease to be labelled slips or lapses, but are termed *mistakes*. The following section presents the types of mistake encountered in rule-based or knowledge-based activity.

Mistakes

Mistakes may be defined as deficiencies or failures in the judgmental and/or inferential processes involved in the selection of an objective or in the specification of the means to achieve it, irrespective of whether or not the actions directed by this decision-scheme run according to plan. (Reason, 1990: p.9)

In contrast to slips and lapses, mistakes occur in the formulation and construction of plans, rather than in their execution. It is commonplace (cf. Rasmussen, 1983; Reason, 1990) to distinguish two kinds of mistake: rule-based and knowledge-based mistakes. Rule-based mistakes involve the application of ‘pre-packaged solutions’ or acting in accordance with some rule of practice which is inappropriate to the current situation. In contrast, knowledge-based mistakes occur when there is no ‘rule’ or existing solution to apply (or misapply). Rather, knowledge-based mistakes occur when an actor’s general knowledge is called upon in the formulation of new plans or action sequences and various characteristics of how humans use their general knowledge lead to mistakes of various sorts. The occurrence of knowledge-based mistakes is notoriously hard to predict.

In this section, then, mistakes are analysed into these two classes and then further subclasses are identified.

Rule-based mistakes.

Rule-based working is best characterised as a process of tackling familiar problems where a person must first decide on a classification for the problem they aim to solve, followed by the

selection of a solution. This process is vulnerable at two main points: on classification of the situation, and on selecting the solution. Failure at these two points gives rise to two sub-classes of error: *misapplication of good rules*, and the *application of bad rules*.

Rule-based mistakes, therefore, are particularly likely to occur when someone is engaged in problem-solving activity in a relatively familiar domain, where previous strategies are applied in order to achieve their objectives. When previous solutions are not applicable, the individual must turn to knowledge-based procedures, and then becomes vulnerable to knowledge-based mistakes.

Knowledge-based mistakes.

When one has to construct new solutions to problems and formulate wholly new plans for action without recourse to existing rule-based solutions, one is subject to various biases which have been documented extensively in psychological literature over the last 25 years (Kahneman *et al.*, 1982). These include: *availability biases, frequency and similarity, confirmation biases, over-confidence, inappropriate exploration of the problem space, attending and forgetting in complex problem spaces, bounded rationality and satisficing, problem simplification through halo effects, control illusions and attribution errors, and hindsight biases and the 'i-knew-it-all-along-effect'*.

This concludes the brief review of human errors due to cognitive factors, i.e. relating to the work of individuals. This paper's consideration of individual human activity has not, however, quite concluded yet. The literature also identifies a further class of 'errors' where actions do not follow the specified plan or procedure. These are distinguished from what has gone before by the complicity of the actor(s) concerned. Whereas slips, lapses, and mistakes are generally taken to be inadvertent, *violations* are usually deliberate deviations from the plan, and are the subject of the following section.

Violations

The distinction between violation and error has been debated, but it hinges on the intentional disobedience of a rule or plan. Many such actions are violations in name only, because people will often disobey a bad rule in order to fix it. They can be classified in a similar manner to errors, according to whether they take place at a skill-, rule-, or knowledge-based level. Violations frequently occur in RE as short-cuts are taken in order to meet deadlines or engineers artfully present their work in project reviews or reckon with other constraints and contingencies (Anderson *et al.*, 1993; Rodden *et al.*, 1994).

Violations are deviations from safe operating processes, practices, procedures, standards or rules. Deviations can be deliberate (breaching rules for safe practice when knowing that such rules exist) or erroneous (acting against the recommendations of a rule without being aware of the existence of such a rule). Of these two classes of violation, deliberate violations have been most studied by psychologists and human factors researchers. However, the research on violations is still small in comparison with what is known about slips, lapses and rule and knowledge-based mistakes.

According to Reason (1990), deliberate violations differ from the errors covered so far in a number of respects. These are summarised in table 1.

There can be many different motivations for deliberate violation. Violations are not necessarily due simply to the wilful negligence of operators, though this sometimes can be the case.

Violations also relate to organizational issues such as: the nature of the workplace; the quality of tools and equipment; whether supervisors and managers turn a ‘blind eye’ to violations to get the job done; the quality of the processes, rules, regulations and operating procedures; and the organization’s overall safety culture.

<i>Errors</i>	<i>Violations</i>
Mainly <i>informational</i> in origin (incorrect or incomplete information leads to error).	Mainly <i>motivational</i> in origin (certain attitudes, social norms or an organizational culture encourages violation).
Errors are <i>unintended</i> .	Violations are typically <i>deliberate</i> .
They can be explained in terms of <i>individual</i> information processing characteristics.	They have to be understood in relation to the <i>social</i> context.
Errors can often be remedied by <i>improving the relevant information</i> .	Violations can only be remedied by <i>changing</i> attitudes, social norms or organizational culture.

Table 1—Comparison of errors versus violations

Classifying Violations

Reason (1990) suggests on the basis of a number of studies that violations can be classified using the same 3-level framework as applied to errors. Skill-based violations occur where some aspect of safe operating practice is violated by the skilled, routine performance of workers. That is, their routinely used skills are violational. Such *routine violations* include corner-cutting and making-do. Within the class of skill-based violations are *optimising violations*. These occur in the performance of some routine task where the actor will optimise how the task is done in non-functional ways. Rule-based violations can be termed *situational violations* as they typically involve breaking restrictive procedures in the light of particular situational exigencies. Situational violations tend to be deliberate acts carried out in the belief that they will not result in bad consequences. Knowledge-based violations are typically engaged in when circumstances are exceptional and unfamiliar for the actor. Thus, violations at this level can be classified as *exceptional violations*. Table 2 summarises this approach:

Performance Level	Error Type	Violation Type
Skill-based	Slips and lapses	Routine and optimising violations
Rule-based	Rule-based mistakes	Situational violations and ‘misventions’
Knowledge-based	Knowledge-based mistakes	Exceptional violations

Table 2—Skill-, rule-, knowledge-based violations

This concludes the consideration of violations by individuals of plans and procedures. It should be noted that violations may well become more likely as the maturity of a process increases and more restrictions are imposed on how the process should be followed as a consequence (Reason, 1995). It is also interesting to note in a similar vein that violations can often take place in order to make the process work, and can be seen to be in the spirit in which the process is intended to be carried out. Violations, therefore, are not necessarily as problematic in themselves as their categorisation in the literature suggests, but may well be indicative of problems in the specification of the process which they are circumventing.

This section has considered the role of the individual in originating errors, which have been classified in terms of skill-based slips and lapses, and rule-based and knowledge-based mistakes. These have been examined principally in terms of their relation to a plan (or process) and their originating factors have been briefly considered. Violations—deliberate deviations from the process—have also been examined. These occur for a host of reasons, many of which are social in origin. Any RE process improvement technique, and especially one explicitly focusing on the human factors component of processes, must incorporate facilities to minimise individual errors. However, consideration needs to move beyond that of errors in individual activity in order to examine the social process which is central to the development of any real system. The following section, therefore, considers group processes, and their potential for error.

Group performance failures and process losses

Many work processes are group activities. That is, they involve the participation of several individuals acting and interacting together. The management and use of safety-critical systems is almost invariably something done by a group or team of individuals. Even under those circumstances where individuals have quite clear roles and responsibilities, the actions they perform in those roles have to be coordinated with the actions performed by others in other roles for the whole team's work to be an effective and dependable process. These remarks are generally true of most advanced industrial production processes—few are the responsibility from start to finish of only one individual—and certainly true of working with dependable systems. Furthermore, the design and development of dependable systems, including the engineering of their requirements, are practically never undertaken by single individuals.

Errors which arise through the ineffective or inappropriate coordination of individuals in groups will be termed *coordination failures*. Related to coordination failures, the work which has been conducted by social psychologists on *group process losses* will be analysed. Interacting in a group can sometimes lead to individuals making fewer contributions to the solution of some problem than they may have made working individually. Such losses may, under certain circumstances, lead to coordination failures, where, for example, group losses entail the omission of some critical action or contribution to the group's work. These remarks mean that how individuals coordinate their actions in group and team activities are crucial to understanding the origins and development of certain kinds of errors and failures in design and use. Just as for errors in individual activity, the RE processes for dependable system development must be designed so as to protect against failures of this nature.

In the remainder of this section, these issues of group performance are analysed under the following headings:

- the facilitation or inhibition of individual task performance by the presence of others (supervisors or an audience);
- the performance of interacting social groups and how this relates to individual performance;
- how factors such as leadership, status, and expertise can affect group interaction and influence group decision making.

The purpose of proceeding through the subject matter in this manner is to build up further classifications of vulnerabilities to error in human processes in much the same way as in the previous section which focused on individual activity. The main difference in this section is that the sources of the research and findings which arise out of it are much more diverse. The

outcome for this paper, however, is much the same, and takes the form of a classification of error and failure types which is built up section-by-section.

Each of the following sections visits a particular focus of research in social psychological studies of social settings and group work in particular. These are: social facilitation and inhibition; performance in interacting groups; group leadership; conformity and consensus; minority influence; and group decision making. In each case, one or more candidates are suggested for addition to a classification of potential problems with group work as was the case for individual work in the previous section.

Social facilitation and inhibition

The mere presence of others can often affect performance even on individual tasks. That is, an individual's performance can be affected even if there is no interaction taking place between the task-performing individual and others present. Paradoxically, these effects can be either facilitatory (presence of others leads to better performance) or inhibitory (presence of others leads to worse performance). Travis (1925) found that, on a simple manual task, people improved their performance in the presence of an audience. On the other hand, Pessin (1933) found that it took people longer to learn a list of items when facing an audience than when practising alone. It is likely that the presence of others improves performance on easy, well-learned tasks (that is, those which are conducted at the skill level) whereas social inhibition occurs when subjects are engaged in difficult or unfamiliar tasks which are not (yet) well learned (that is, those tasks which are more likely to be conducted at the rule and knowledge-based levels)¹.

Any full theoretical explanation of social facilitation and inhibition effects would have to integrate both arousal/motivational effects and information processing or problem solving effects of the presence of others. For details of such a theory, see (Paulus, 1989).

Findings of the sort documented here should lead, for example, to consideration of questions surrounding when and where the direct supervision of workers is appropriate. Direct supervision, even if it is founded in the aim of monitoring and checking for errors in performance, can sometimes lead to performance losses if the task in question implicates the use of one's general knowledge in problem solving activity or if the task is being conducted using pre-packaged proceduralised rules. On the other hand, the performance of a routine, skilled task may even be improved by the presence of a supervisor, though whether the dedication of a supervisor to monitoring an individual performing a routine task is worthwhile use of another individual's time may turn out to be debatable.

Performance in interacting groups

How does the performance of groups of people jointly engaged in the performance of some task relate to the performance of the individuals comprising them? Will a group outperform the best individual within that group (in which case conducting the task in a group situation has benefits) or will a group lead to performance worse than this (in which case the task would be better performed by either the most able individual alone or by a 'nominal group' of people working independently from whom the best solution is drawn? Again, at first glance, the research presents

¹ See (Zajonc, 1965) for further details and (Manstead and Semin, 1980) for the connection of these phenomena to distinctions between automatic (skill-based) and consciously controlled (rule and knowledge-based) performance.

a contradictory picture. There are circumstances where groups outperform the level the best individual within it is capable of performing alone. On the other hand, there are circumstances where individuals working alone, when their performance levels are appropriately summed, outperform groups.

Steiner (1972; 1976) convincingly argues that the relationship between group and individual performance, and hence whether there are significant group productivity losses, depends upon the kind of task being performed. To give a simple example, a team of people building a house will obviously complete it faster than an individual working alone. Furthermore, a building team where the jobs that each does are matched to their abilities will complete the job faster than individuals of identical abilities. On the other hand, it is likely that the fastest relay race team will be made up of the four fastest runners. A single slow runner will impact directly on the whole team's performance. Thus, in many activities, the performance of the least able member is critical to group success and cannot be compensated for by the good performance of more able members. Table 3 summarises Steiner's classification of tasks:

Question	Answer	Task type	Examples
Can the task be broken down into sub-components, or is division of the task inappropriate?	Sub-tasks can be identified	Divisible	Playing a football game, building a house, preparing a six-course meal
	No sub-tasks exist	Unitary	Pulling on a rope, reading a book, solving a mathematics problem
Which is more important: quantity produced or quality of performance?	Quantity	Maximising	Generating many ideas, lifting the greatest weight, scoring the most runs
	Quality	Optimising	Generating the best idea, getting the right answer, solving a mathematics problem
How are individual inputs related to the group's product?	Individual inputs are added together	Additive	Pulling a rope, stuffing envelopes, shovelling snow
	Group product is average of individual judgements	Compensatory	Averaging individuals' estimates of the number of beans in a jar, weight of an object, room temperature
	Group selects product from pool of individual members' judgements	Disjunctive	Questions involving 'yes-no, either-or' answers, such as mathematics problems, puzzles, and choices between options
	All group members contribute to the product	Conjunctive	Climbing a mountain, eating a meal, relay races, soldiers marching in file
	Group can decide how individual inputs relate to group product	Discretionary	Deciding to shovel snow together, opting to vote on the best answer to a mathematics problem, letting leader answer question

Table 3—A summary of Steiner's typology of tasks (1972; Steiner, 1976). Reproduced from (Wilke and Knippenberg, 1988) in (Hewstone *et al.*, 1988: p325)

Steiner and others who have followed him (e.g. Wilke and Knippenberg, 1988) have shown that by classifying tasks into the different ways in which individual inputs are combined in group performance, one can make predictions about the nature and extent of group performance losses. These results are summarised below:

Additive tasks

For these, a group will always outperform any single individual comprising the group because, by definition, group performance equals the sum of individual performance levels. However, it was discovered over one hundred years ago by Ringelmann (see Wilke and Knippenberg, 1988) that, while this is true, the contribution that individuals make as part of a team is often less than the contribution they would make if acting alone—a phenomenon known as 'social loafing' (for a substantial recent review, see Karau and Williams, 1993). Indeed, individual contributions often go down as a function of increasing group size. Thus, although the actual productivity of

the group exceeds the productivity of the best member, the potential productivity is higher still. These losses from potential productivity are of two main sorts:

- **Motivational losses.** As part of a team, individuals need not be motivated to perform as well as they would when performing individually. One can be a 'free-rider'.
- **Coordination losses.** Some of each individual's activity has to be devoted to coordinating their efforts with others rather than the direct performance of the task itself.

Compensatory tasks

Consider a task in which a group of individuals are each making estimates of the number of beans in a jar or of the weight of an object. Shaw (1981) argues that the bulk of evidence for tasks of this sort indicates that the statistical average of a group of people making such estimates is more reliable than the judgements of most of the individual making up the group. That is, the overestimates of some cancel out the underestimates of others. Steiner (1972; 1976), however, suggest that this conclusion should be taken with some care as it is not always possible in daily life to statistically average in simple ways. Equally, one is not always acquainted with the skills and biases of the individuals making up the group so one cannot be sure whether the variation in their biases will be likely to lead to overestimates cancelling out underestimates.

Disjunctive tasks

These are tasks where the group's solution or overall outcome will be a selection from the individually proposed solutions or individually contributed performances. Many problem solving tasks are of this sort, where a single option (or a range of options which is less in number than all those proposed by the group) must ultimately be selected. Very often, not surprisingly, group performance in disjunctive tasks will equal the best performance of the individuals who make up the group. Shaw (1932) explains this result by noting that groups have the opportunity to correct the errors and reject the incorrect suggestions made by individuals. However, further work has made the picture more complex for, on the one hand, it is not always the case that a group member does propose the best solution to a problem and, on the other hand, it is not always the case that a groups happen to adopt the best solution even if it is proposed by one of its members. Much depends on whether the best solution is *recognised as such* by the group members. This may only be possible for certain kinds of problem or task. Tasks where optimal solutions are easily recognised as such are known as *eureka tasks*. In contrast, for non-eureka tasks, it is quite possible that a correct solution will not be proposed or an incorrect solution will amass support from the group members. The critical aspects for group success in disjunctive tasks appear to be (Steiner, 1972; Thomas and Fink, 1961):

- **Potential performance and member expertise.** Do group members possess the right expertise for solving the problem?
 - if there is not at least one competent member the group is unlikely to succeed.
- **Motivation.** Do group members, possessing the correct solution, actually propose it?
 - if a low status member happens upon the solution, they may feel unable to express it.
- **Coordination.** Do correct solutions elicit more support than incorrect ones so that they emerge as the group's overall solution?
 - if a low status member or one held to be inexperienced does express the solution, it might be resisted by other group members or ignored;
 - in some problem domains, the overhead of the task of convincing others that a solution is correct may be too high for correct yet non-obvious solutions to be adopted.

Conjunctive tasks

In these it is necessary that every member contributes lest the group fails. For runners in a relay race team, it is only necessary for one of them to drop the baton or run out of their lane for the whole team to be disqualified. Conjunctive tasks, where group success depends upon the least proficient member, are more likely to fail with increasing group size, as the probability that the group will contain at least one member who does not contribute increases the more members there are. This is prominently true for unitary group tasks. However, many tasks in everyday life are divisible. That is, different members can adopt different sub-tasks. If the sub-tasks have a relative degree of independence in their execution then the effects of failure of one sub-task (allocated to one individual or a subgroup) may not be catastrophic for the whole enterprise. Additionally, in divisible conjunctive tasks, if the competencies of the individuals match the sub-tasks they are engaged in, then the potential productivity of the group can rise above the productivity of the least able member.

Discretionary tasks

A discretionary task is one in which the members of the group themselves decide upon what kind of task it is, how the task is to be performed, and how individual inputs are to be coordinated into a group outcome. Discretionary tasks, if this can be put so paradoxically, are by definition ill-defined. They are likely to resolve into one of the forms of task already discussed in which case one can expect the corresponding performance levels to be achieved and losses (if any) encountered. However, this very process of resolving the task and its conduct is itself an overhead to the performance of the task in which coordination losses are likely to become critical.

The results and analysis above are summarised in Table 4.

Task	Group Productivity	Description
Additive	Better than best	Group out-performs the best individual member
Compensatory	Better than most	Group out-performs a substantial number of group members
Disjunctive (eureka)	Equal to the best	Group performance matches the performance of the best member
Disjunctive (non-eureka)	Less than best	Group performance can match that of the best member, but often falls short
Conjunctive (unitary)	Equal to the worst	Group performance matches the performance of the worst member
Conjunctive (divisible with matching)	Better than the worst	If sub-tasks are properly matched to ability of members, group performance can reach high levels

Table 4—Group performance of groups working on various types of tasks (Forsyth, 1983; Steiner, 1972; 1976). Reproduced from (Hewstone *et al.*, 1988: p332)

Summary of group performance failures

Just as Steiner’s system can be used to classify tasks, this scheme as presented in table 4 can be used to classify group performance failures and errors in such tasks. Thus, one can refer to additive task errors, compensatory task errors and so forth as errors made in additive and compensatory tasks. In addition, this task-based classification scheme can be complemented with

one based on some of the social psychological phenomena noted. This gives rise to further additions to the classification of social and group performance failures.

Group leadership

There is a considerable social psychological literature on the topic of leadership in small groups. This literature covers many different aspects of the issue but perhaps the most relevant for our purposes concerns the leadership factors which are likely to promote the productivity or success of a group in the performance of the group's task.

One approach has been to study experimental groups who do not have a leader initially to see under what circumstances a leader might emerge and what characteristics that leader may have. In an early series of observational studies of groups in interaction, Bales (1955) suggest that a specialisation often occurs in groups between 'socio-emotional specialists' who are oriented towards the solidarity of the group and who resolve tensions in the group and so forth, and 'task specialists' who are more concerned with the execution of the task itself than the group's internal social dynamics. Bales (1955) argue that group leaders rarely embody both aspects and that effective groups often have two 'leaders'—one concerned with the socio-emotional aspects of the group, one concerned with its effective task performance. This conclusion was borne out by Likert (1967) in studies which suggested that effective leaders manifested (and rather rarely both) employee centred or production centred behaviour. Steiner (1976) argues that employee centred (or socio-emotional) behaviour is necessary to ensure that the unrealised productivity of the group is kept to a minimum (e.g. by encouraging all group members to participate) and that production (or task) centred behaviour is necessary to ensure that the group's potential productivity is as high as it can be. Thus, according to this literature, there seem to exist two leadership styles, each with its own impact on group performance and effective groups often need an appropriate balance between the two.

Of course, it is not necessary for groups to have a leader. Indeed, some groups may perform better without an explicit leader and may not have the need for a leader to emerge. There seem to be a number of critical factors of relevance to whether groups have a need for a leader.

- **size**—Hemphill (1961) argues that effective performance in a large group is often dependent on the group having a leader to coordinate various specialised subgroups and facilitate overall decisions.
- **availability**—the group should have someone at their disposal who has had relevant leadership experience.
- **the value of success**—success must be important to the group for it to seem worthwhile to install a leader.

Finally, Rutte (1984) suggest that group leadership may be important in resolving the 'free-rider problem' (that is, the possibility that some group members may contribute unequally and 'ride' on the contributions of others) and *if task success is in danger*.

This analysis leads to a number of potential sources of group failure and error if leaders are inappropriately or ineffectively installed into groups. For example:

- Errors or failures due to inappropriate leadership style (or balance between styles)
 - e.g. a group lacks a 'task sensitive' leader
- Errors or failures due to inappropriate leadership skills
 - e.g. the person appointed leader does not have appropriate experience

- Errors or failures due to the excessive influence of the leader
 - e.g. a high status leader who does not encourage contrary opinions to emerge.

Conformity and consensus: normative and informational influence

Sooner or later in working together as a group, group members will become aware of the opinions or contributions of others. Specifically, they may often become aware of whether they are in the majority or minority of group opinions. Indeed, periodically, many groups explicitly assess the level of agreement within the group through formal voting, a 'show of hands' or other means. When group members become aware of the overall position of the group, how does this influence their views? Do individuals within groups come to realign their views in accordance with the majority view ('majority influence') or do they retain their private views in spite of the majority view? Equally, under what circumstances can minority opinion come to influence the views of the majority?

Certainly, the existence of a majority position can influence the views of minority individuals. Classic social psychological experiments by Sherif (1953) and Asch (1951) are often claimed to demonstrate just that. In Asch's work, for example, a series of lines of varying lengths are shown to a group of people. One pair of the lines are identical in length and the group members' task is to say which two are identical. Asch arranged the experiment so that the majority of the group members are 'confederates' of the experimenter who are instructed to give a consistent, yet incorrect response. Asch found that as many as 37% of people, when confronted with a clearly incorrect majority opinion, nevertheless fell in line with the majority view.

To further analyse why people conform in such settings, social psychologists often make a distinction between informational and normative influence. People may be influenced in their opinions by the information provided by others, what their opinions are and the reasons they give for them. In contrast, people may also be influenced by normative reasons to conform with the views of others. For example, an individual may wish to avoid being disliked and so agree with a majority view to promote the chances that they will be popular within the group. Alternatively, an individual of low status within a group may change their view to match that of the majority if that majority contains high status members. Both of these are examples of normative influence in action.

Thus, consensus in a group can come about through the combination of two factors: the normative and informational influence that different group members have on one another. A number of studies have been conducted to try and tease apart the effects of normative and informational influence to gain an impression of when each factor is most potent. This work can only be crudely summarised here (for more details, see Van Avermaet, 1988):

- Normative influence is heightened by, for example:
 - rewarding conformity itself
 - increasing the interdependence of group members on each other
 - insisting that opinions and contributions to the group are made public by being spoken aloud rather than written down privately or anonymously
 - informing the group that it will be compared with other groups.
- Informational influence is greater, for example:
 - for group members who are perceived as being competent in the task domain
 - as sources of information become more reliable (e.g. improvements in viewing conditions)

- as the majority increases (but only if the majority are seen to be acting independently and are not merely repeating the same reasons for the majority opinion)
- when the range of opinion within the group increases (that is, if the majority is not unanimous).

The distinction between normative and informational influence again suggests two classes of process error or failure which may occur as groups interact in the pursuance of some task:

- Errors or failures due to conformity arising from inappropriate normative influence
 - e.g. when an incorrect judgement of a high status member commands influence because others respect that status
- Errors or failures due to conformity arising from inappropriate informational influence
 - e.g. when the judgement of one member is based on false evidence or is misunderstood by another group member (at least some of these errors may arise due to slips, lapses or mistakes being made within the group).

Innovation: minority influence

Of course, the existence of a majority opinion or a subsequent overall group consensus is no guarantee in itself of the correctness or worth of the opinion. Indeed, the main problem with conforming to majority opinion is that important minority views may be ignored. As one can imagine, there is much that a group leader or facilitator can do to prevent minority opinions being passed over. A study by Maier (1952) suggests that leadership style is an important determinant here of whether minority opinions can come to have influence. A group leader who merely monitored the procedures and agenda followed by the group did not help minority opinions find their voice while a leadership style which encouraged a more even group participation did allow effective minority opinions to emerge. Under these circumstances, encouraging uniform participation went some way to ensuring that *all opinions* and not just the majority one were given equivalent discussion. Note that this strategy might also have assisted in ameliorating process losses due to the free-rider problem (see above).

In principle, of course, it *must* be possible for a minority to influence majority opinion or otherwise change and innovation would be impossible. However, it is equally clear that, due to both normative and informational factors, a majority is hard to displace. Moscovici (1976) argues that minority opinion can alter majority views provided that the minority adhere to a *behavioural style* in which they propose a clear position and hold firmly to it. Of particular importance to this behavioural style is the *consistency* with which the minority defend and advocate their position. This consistency is made up of two components: intra-individual consistency over time (individuals will not waver within themselves in their views) and inter-individual consistency over time (individuals will not waver between themselves in their views). Note it is important for minorities to sustain their positions consistently *over the long term*. This contrasts with the effects of majority influence which can be immediate.

Provided the conditions noted above are held, there is a chance for the minority to influence the majority. However, it has been shown if these very same strategies are used in turn by the majority against a consistent minority, the effect of the minority can swiftly disappear (Doms and Van Avermaet, 1985).

Interpretations of exactly how minority influence takes place vary (see Maass and Clark, 1984). However, it is clear according to this literature that, if an adequate range of opinions are to be considered within a group, strategies must be found for permitting minority opinions to emerge. Without this, a further class of group process error or failure may emerge:

- Errors or failures due to the exclusion of minority opinion.

Group decision making: the risky shift, group polarisation and groupthink

Consider a decision making task in which a group has to resolve on a course of action. Consider further that each course of action has a set of risks and probabilities attached to it. Naturally, this is a very common situation in the design of dependable systems or in the engineering of requirements for them. Under such circumstances, what are the relations between the views of individual group members (the course of action they as individuals would decide upon) and an overall group decision. In a famous experimental study, Kogan (1964) showed that groups seem to be more tolerant of risks than the individuals composing them. That is, the course of action resolved upon by the group was more risky in general than the decisions that the individuals would have tended to make in isolation. This phenomenon is often known as *the risky shift*.

However, since this early work, it has been shown that groups are not always 'riskier' than the individuals comprising them. Quite often groups can be more cautious than individuals and the risky shift is not the general phenomenon it might have at first appeared to be. What seems to be important is the initial level of opinion within a group. If the individuals initially favour moderately risky strategies, then the group will adopt yet more risky options. However, if the individuals who comprise the group initially favour moderately safe options, then the group decision is likely to be even less risky. That is, the group 'shifts' further in the direction already favoured. Myers (e.g. Myers, 1982) terms this phenomenon *group polarisation*. Group polarisation seems to occur in a wide variety of contexts (see Lamm and Myers, 1978).

The previous discussion would suggest two reasons why group polarisation can occur. First, group polarisation may occur for normative reasons as each group member compares their views with other members' positions. Members—on realising the overall group norm—may come to adopt more extreme positions to align themselves more fully with the direction of the group's thinking. Alternatively, group polarisation may occur due to processes of informational influence. Group interaction will yield a number of arguments, most of which are in support of the position already favoured by the group. Group discussion therefore will tend to *increase* the amount of support that the overall group position will have. In the light of this, members may take even more extreme positions as they will be encountering arguments for their view which they had not heard before. Group polarisation, on this view, becomes a matter of mutual persuasion. Clearly, both informational and normative influence can operate in explaining group polarisation and Isenberg (1986) explicitly argues that any plausible theory of group polarisation should combine these two factors.

So-called *groupthink* (Janis, 1972) is an extreme case of group polarisation. Janis described a number of cases of military and political decision making (most notably the decision making of the Kennedy administration leading to the 'Bay of Pigs' invasion in 1961) in which group polarisation takes extreme forms. Groupthink occurs when a group of already like-minded individuals form a highly cohesive group and mutually reinforce themselves in a course of action which may well turn out unwise in spite of the group's extreme conviction. Janis argues that the following antecedent conditions make groupthink possible:

- the decision making group is highly cohesive;
- the group is isolated from alternative sources of information;
- the group's leader clearly favours a particular option.

If these conditions are met, then groupthink will be characterised by discussions in which the group develops:

- an illusion of its own invulnerability;
- a tendency to mutually rationalise actions which are in line with the proposed option;
- while ignoring or discounting inconsistent evidence and arguments.

Groupthink and less extreme forms of group polarisation, then, constitute another possible source of error or failure in group or team work:

- Errors or failures due to group polarisation and groupthink.

This concludes our consideration of vulnerabilities to failure or error due to group processes. A number of categories of failures have been drawn out from the literature, each of which may be applicable to the activities of requirements engineers, or systems developers in general, when working as a group or team. This literature was turned to so as to widen the perspective on human error in systems development from the individual to the social. The next section broadens our enquiry once more, from the social to the organizational.

Organizational problems and failures

The third broad area of research which is turned to in this paper is that relating to work at an organizational level, and the errors and failures which organizations are vulnerable to. Previous sections have characterised the work of requirements engineers as a combination of individual and cooperative work. One justification for opening up the coverage to the social psychology of groups was that individual activity does not occur in isolation, and that for much of the RE process human activity is predominantly oriented towards the work of others. The same argument applies here, in that the various individuals and their groupings in project teams and so forth all exist within some organizational setting, and all their activity pertains to some organizational goal or other. As such, an understanding of organizations, how they are made up, and how they function, is extremely relevant to this paper.

The remainder of this section considers three perspectives on the ways in which organizations have been found to function in hazardous situations, and how they can contribute to failures, but also to their avoidance. First of all, organizational failures are viewed as ‘accidents waiting to happen’ in the work on latent organizational failures. Following this a classification scheme for organizations is presented in terms of the degree of interactive complexity and tightness of coupling between components. On the basis of this classification, it is suggested that, for some types of organization, accidents are inevitable and should therefore be considered normal. Finally, this school of thought is contrasted with work from a number of researchers who contest that organizations can be highly reliable in hazardous settings provided a number of recommendations are adhered to.

Latent organizational failures

In this section, the importance of organizational factors in understanding the origins of error and failure are turned to. An increasing amount of work in the field of accident and error analysis is concentrating on the factors that can be attributed to failures at an organizational level. Reason (1990; 1992), in particular, has coined the term *latent organizational failures* to describe the errors resulting from organizational factors which may remain dormant for some time before combining with one or more other factors in the cause of an accident. These latent failures frequently take the form of fallible decisions taken high up in the organization hierarchy, and are so named because of the likelihood that they will remain unnoticed for some time before being transmitted through the organization’s various levels to combine with a triggering event or active failure (unsafe act) to breach the system’s safety defences and cause an accident (see figure 1).

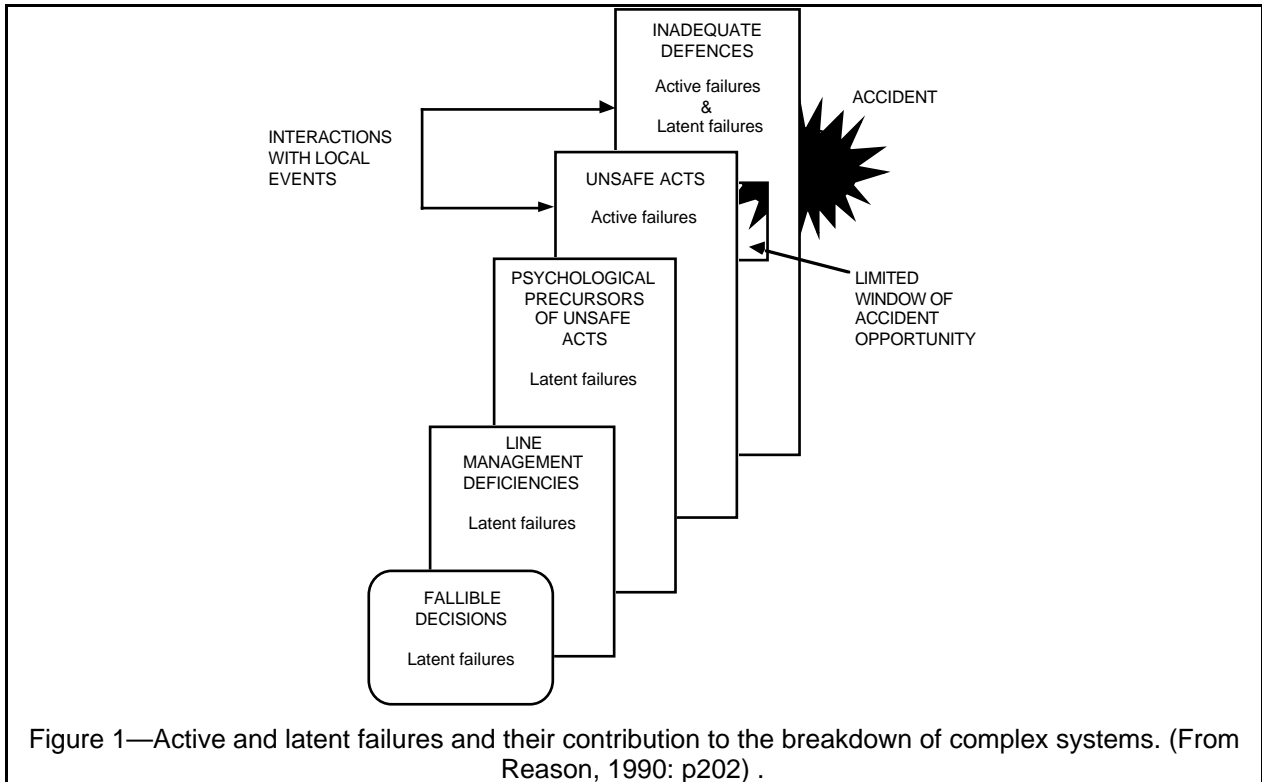


Figure 1—Active and latent failures and their contribution to the breakdown of complex systems. (From Reason, 1990: p202) .

Turner (1992) also proposes that system failures develop over a period of time, and are usually due to a number of factors, rather than a single catastrophic event. According to Turner, the development of a system failure is typified by the sequence in figure 2 below:

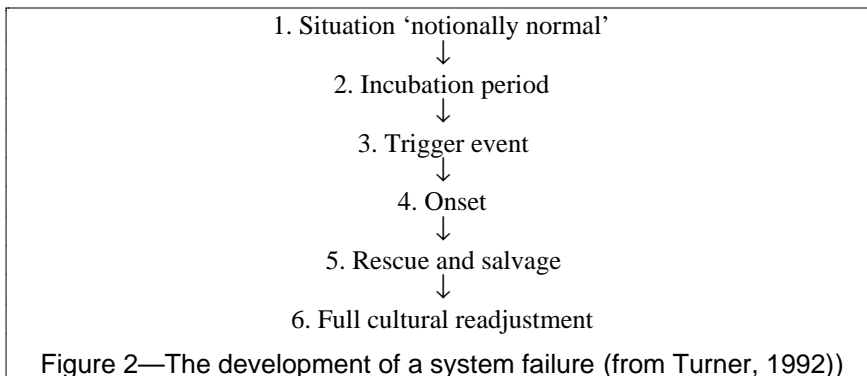


Figure 2—The development of a system failure (from Turner, 1992))

Turner (1992) sets out the predisposing features which will typically interact in the incubation period, during which time the system is a “disaster waiting to happen”, as follows:

- Organizational rigidities of perception and belief
- Decoy phenomena which distract attention from genuine hazards
- A range of many types of *information* and *communication difficulties* associated with the ill-structured problem which eventually generates the accident. Such ambiguities, noise and confusion are frequently complicated by unexpected elements injected into the situation by ‘strangers’ who are unfamiliar with the system, most frequently members of the public, and by additional surprises arising from unanticipated aspects of the ‘site’ or of the technical system involved

- Failure to comply with existing safety regulations
- A variety of modes of minimizing or disregarding emergent danger, especially in the final stages of the incubation period (Turner, 1992: p194, original emphasis)

Reason (1992) identifies ten organizational failure types, and three latent workplace factors which can combine in many ways in order to produce various situations that are the “early warning signs” of an accident. The workplace factors are:

- *violation producing conditions*, such as unfamiliarity with the task, poor human-system interface, and irreversibility of errors.
- *error-producing conditions*, including lack of organizational safety culture, management/staff conflict, and poor supervision and checking; and
- *inadequate defences*, where human and technical elements fail to deal with an accident in terms of protection, detection, warning, recovery, containment, or escape.

The organizational failure types that these factors can combine with are:

- incompatible goals
- organizational deficiencies
- inadequate communications
- poor planning and scheduling
- inadequate control and monitoring
- design failures
- unsuitable materials
- poor operating procedures
- inadequate maintenance, and
- poor training.

This work naturally leads to consideration of the types of organization that are error-prone as a means of identifying potential problem areas that need to be addressed.

Typologies of organizations

Reason (1992) proposes a seven-point organization rating scale, which can be used to classify the safety culture that exists within an organization. This scale, based on how organizations react to hazards, ranges from *pathological*, where safety practices are at the bare minimum, through to *generative-proactive*, where the organization is constantly striving to improve safety measures. Reason adopts medical terminology, and talks about organizational safety ‘health indicators’ and how safety research has so far concentrated on performance and warning indicators that are concerned with the immediate to medium term picture of ‘health’. It is Reason’s contention that progress in understanding what makes a safe organization will be made only by further study of high reliability organizations, and by using global ‘health’ indicators that give a more predictive and long term view of organizational safety.

Perrow (1984) classifies organizations according to their *interactions*, which may be linear or complex and their *coupling*, either loose or tight. The classifications are summarised in tables 5 and 6 respectively:

Complex Systems	Linear Systems
Tight spacing of equipment	Equipment spread out
Proximate production steps	Segregated production steps
Many common-mode connections of components not in production sequence	Common-mode connections limited to power supply and environment
Limited isolation of failed components	Easy isolation of failed components
Personnel specialisation limits awareness of interdependencies	Less personnel specialisation
Limited substitution of supplies and materials	Extensive substitution of supplies and materials
Unfamiliar or unintended feedback loops	Few unfamiliar or unintended feedback loops
Many control parameters with potential interactions	Control parameters few, direct, and segregated
Indirect or inferential information sources	Direct, on-line information sources
Limited understanding of some processes (associated with transformation processes)	Extensive understanding of all processes (typically fabrication or assembly processes)
Summary Terms	
Complex Systems	Linear Systems
Proximity	Special segregation
Common-mode connections	Dedicated connections
Interconnected subsystems	Segregated subsystems
Limited substitutions	Easy substitutions
Feedback loops	Few feedback loops
Multiple and interacting controls	Single purpose, segregated controls
Indirect information	Direct information
Limited understanding	Extensive understanding

Table 5—Complex vs. Linear systems (from Perrow, 1984: p88)

Tight Coupling	Loose Coupling
Delays in processing not possible	Processing delays possible
Invariant sequences	Order of sequences may be changed
Only one method to achieve goal	Alternative methods available
Little slack possible in supplies, equipment, personnel	Slack in resources possible
Buffers and redundancies are designed-in, deliberate	Buffers and redundancies fortuitously available
Substitutions of supplies, equipment, personnel limited and designed-in	Substitutions fortuitously available

Table 6—Tight and loose coupling tendencies (from Perrow, 1984: p96)

Perrow then classifies organizations according to where they fall in the two-dimensional categorisation of linear-complex interactions versus tight-loose coupling. He uses this classification when considering whether authority in an organization should be centralised or decentralised in order to reduce the risk of accidents. It can be seen from this (see table 7) that tightly coupled, complex interactions produce incompatible demands on the organization. Tight coupling requires authority to be centralised, whilst complex interactions require decentralised authority. It is this class of organizations that Perrow believes to be especially vulnerable to

system accidents. In fact, he argues that accidents are inevitable in tightly coupled, interactively complex systems, and to this extent can be considered ‘normal’. Perrow’s argument is taken up by Mellor (1994) and is supported with a number of cases where Mellor argues that the use of computers in any system will increase both the interactive complexity and the degree of coupling, and therefore make the occurrence of normal accidents more likely.

INTERACTIONS	
Linear	Complex
<p>CENTRALISATION for tight coupling. CENTRALISATION compatible with linear interactions (expected, visible).</p> <p>e.g.. Dams, power grids, some continuous processing, rail and marine transport.</p>	<p>CENTRALISATION to cope with tight coupling (unquestioned obedience, immediate response). DECENTRALISATION to cope with unplanned interactions of failures (careful slow search by those closest to subsystems).</p> <p>Demands are incompatible.</p> <p>e.g.. Nuclear plants, weapons; DNA, chemical plants, aircraft, space missions.</p>
<p>CENTRALISATION or DECENTRALISATION possible. Few complex interactions; component failure accidents can be handled from above or below. Tastes of elites and tradition determine structure.</p> <p>e.g.. Most manufacturing, trade schools, single-goal agencies (motor vehicles, post office).</p>	<p>DECENTRALISATION for complex interactions desirable. DECENTRALISATION for loose coupling desirable (allows people to devise indigenous substitutions and alternative paths), since system accidents possible.</p> <p>e.g.. Mining, R&D firms, multi-goal agencies (welfare, DOE, OMB), universities.</p>

Table 7—Centralisation/Decentralisation of authority relevant to crises (from Perrow, 1984: p332)

Normal accidents vs. High reliability organizations

Sagan (1993) portrays the normal accident approach of Perrow and others as being a somewhat pessimistic view, and contrasts it with the more optimistic² work of a number of researchers whom he groups under the ‘High Reliability Theory’ school of thought. These researchers have examined systems such as U.S. Navy aircraft carriers, the American Federal Aviation Administration’s (FAA) air traffic control system, nuclear power plants, and the human body’s immune system (see, for example, La Porte and Consolini, 1991; Marone and Woodhouse, 1986; Roberts, 1989; Wildavsky, 1988) amongst others. The systems they have studied all display high levels of reliability, and the researchers believe that this can be explained by a number of common features which they have discovered in the organizations concerned. In particular, the studies have pointed to four critical causal factors that they believe if satisfied will lead directly

² Although some of the high reliability theorists object to this optimistic label, see (Sagan, 1993, p47)

to highly reliable operations, even for organizations working with hazardous technologies. These factors are expanded upon below³:

- **Organization leadership prioritises safety.** Two reasons are given for why it is important that the political elite and leaders of the organization should place a great emphasis on the importance of safety.
 - First, the needs for high levels of redundancy and constant operational training requires a great financial commitment. Therefore, if the political authorities and organizational leadership concerned are willing to devote considerable resources to safety, then accidents will be less likely.
 - Second, the leadership must see safety as a priority in order to be able to transmit this to the rest of the organization and in turn lead to the development of a strong organizational culture of safety.
- **High levels of redundancy exist in personnel and technology.** In the words of one of the high reliability theorists, “duplication is a substitute for perfect parts” (Bendor, 1985), and redundancy is seen as a must in the quest to build “reliable systems from unreliable parts”. There are two types of redundancy which can be employed
 - *duplication* Two (or more) different units are dedicated to performing the same function. The studies of aircraft carrier operations have highlighted the importance of both technical and personnel duplication.
 - *overlap* More than one unit has the same functional area in common. For example, different officers may be assigned the same duties, whilst their overall responsibilities may differ, thus allowing each one to cross-check the other’s work.
- **Decentralised authority, continuous training, and strong organizational culture of safety are encouraged.** These three factors are seen as relieving some of the pressure created by individual failures such that redundant systems are not over stressed.
 - *decentralisation* In order to allow for those closest to problems to respond rapidly and appropriately to any situation as it develops, a high degree of decentralised authority for decision-making is required.
 - *continuous operations and training* Organizations are more likely to relax vigilance and become complacent when the conditions of operation become stable and routine, leading in turn to carelessness and error.
 - *culture of reliability* In a stable operating environment, an organization can rely upon standard operational rules and procedures for maintaining reliability because the actions performed and decisions made by operational staff will fall within a predictable set. This is not usually the case for organizations that are working with hazardous technologies, where staff must react rapidly to an unpredictable environment in an appropriate manner. Developing a reliability or safety culture at all levels of the organization through recruitment, socialisation, and training of staff is seen as a way of achieving this degree of assurance that staff will respond to dangerous situations in the appropriate manner.
- **Organizational learning takes place through trial-and-error, simulation, and imagination.** It is of great importance that an organization is capable of learning over time, if it is to achieve a highly reliable status. This trial-and-error process must work

³ See chapter 1 of (Sagan, 1993) for a good review of high reliability theory, normal accident theory, and a comparison of the two. See also the *Journal of Contingencies and Crisis Management* 2 (4) for a special issue devoted to further debate between Normal Accident Theory and High Reliability Theory.

such that both safety- and danger-inducing activities are recognised as such, and that operating procedures are adjusted in order to increase the operational level of safety. High reliability theorists cite the changes in working practices in nuclear power plant control rooms after the Three Mile Island incident, and point to how many of the safety procedures in place on U.S. aircraft carriers were introduced following crashes or deck fires. Obviously, it is imperative that organizations learn from such serious incidents, and from lesser ones as well, but it would be most unwise for an organization—*especially* one working with hazardous technology—to court disaster for the benefit of a potential learning experience. For this reason two supplementary strategies for improving organizational learning are proposed:

- *simulations* Rather than waiting for an accident to happen, the organization can simulate a possible scenario, and use this both as a training exercise for the staff, as well as allowing procedures to be altered in the light of this experience. This is routine practice in both the nuclear and aerospace industries, where simulations are used to provide operators and pilots with the experience of trial-and-error learning, without the serious consequences.
- *imagination* In addition to simulating accident scenarios using dedicated simulators or operational equipment, it is also possible to envisage hazardous events and their consequences with pen and paper or more sophisticated tools. This is where risk or hazard analysis fits in, or any such method which is used to anticipate possible operator or design errors, with safety consultants analysing the potential for errors in existing procedures, and proposing solutions to these problems.

The high reliability theorists believe that organizations working with hazardous technologies can operate safely through good management and organizational design which apply the above factors. This is in contrast to the normal accident school of thought which states that accidents are inevitable in such organizations, which are by definition highly complex and tightly coupled. Sagan (1993) applies the two schools of thought to the problem of safety with nuclear weapons operations, possibly the most hazardous of hazardous technologies, in order to test the assumptions that the two theories are based on and how well their predictions fit with reality. Table 8 provides a summary of the contradictory assumptions, statements, and predictions of the two schools of thought.

High Reliability Theory	Normal Accidents Theory
Accidents can be prevented through good organizational design and management.	Accidents are inevitable in complex and tightly coupled systems.
Safety is the priority organizational objective.	Safety is one of a number of competing objectives.
Redundancy enhances safety: duplication and overlap can make “a reliable system out of unreliable parts.”	Redundancy often causes accidents: it increases interactive complexity and opaqueness and encourages risk-taking.
Decentralised decision-making is needed to permit prompt and flexible field-level responses to surprises.	Organizational contradiction: decentralisation is needed for complexity, but centralisation is needed for tightly coupled systems.
A “culture of reliability” will enhance safety by encouraging uniform and appropriate responses by field-level operators.	A military model of intense discipline, socialisation, and isolation is incompatible with democratic values.
Continuous operations, training, and simulations can create and maintain high reliability operations.	Organizations cannot train for unimagined, highly dangerous, or politically unpalatable operations.

Trial and error learning from accidents can be effective, and can be supplemented by anticipation and simulations.

Denial of responsibility, faulty reporting, and reconstruction of history cripples learning efforts.

Table 8—Competing perspectives on safety with hazardous technologies (from Sagan, 1993: p46)

At the end of Sagan’s inquiry into the applicability of the two sets of theory, he answers the question that he asks himself at the beginning of the book: “Which theoretical perspective proved to be most helpful in understanding the history of nuclear weapons safety?” (Sagan, 1993: p252) . Whilst acknowledging the useful insights provided by the high reliability perspective, he found much stronger support for the pessimistic views of Perrow and others with the normal accidents approach. Not only this, but based on the historical data about nuclear weapons operations in the U.S.A., he extends Perrow’s pessimism with four further issues that contribute to the causes of accidents in high technology systems. In brief, these are:

- **The dark side of discipline** Both high reliability theorists and normal accidents theorists agree that a strong organizational culture—with high degrees of socialisation, discipline, and isolation from the rest of society—can lead to greater safety when working with hazardous technologies. Goffman (1961) refers to such organizations as “total institutions” and the aircraft carriers of the high reliability theorists’ case studies are fine examples of this. Perrow, however, questions whether it is either possible or desirable for civilian organizations to be run on such a strict military model. Sagan cites several examples that point to such a culture leading to “excessive loyalty and secrecy, disdain for outside expertise, and in some cases even cover-ups of safety problems, in order to protect the reputation of the institution” (p254)
- **Conflicting interests** Whilst organizational leaders may place a great priority on achieving safe operations, they will also have many other, potentially competing interests, some of which may take priority.
- **Constraints on learning** Organizational learning is constrained by political and social pressures to portray a certain image of the organization to the outside world. What shocked Sagan was to find that not only would this lead to false reporting to the press and so on, but that the invented or altered stories would come to be believed by their creators. What he found was “...not just a further piece of evidence showing how difficult it is for large organizations to learn from success. These cases show something more disturbing: the resourcefulness with which committed individuals and organizations can turn the experience of failure into the memory of success.” (pp257-258)
- **The measure of safety** Finally, Sagan warns against believing the story-teller, especially when the story is about the teller, and a little bit too good to be true. He criticises the high-reliability theorists for relying upon accounts of safety in U.S. Navy aircraft carrier operations which have been produced by the Navy themselves, and urges those studying organizations working with hazardous technologies to exercise scepticism when dealing with such accounts.

It should be clear from the foregoing discussion that the debates about human error, latent failure and reliability in an organizational context are highly complex and unresolved. For this reason, it would be misleading and, indeed, dangerous to simplify the literature by siding with any single position or organizational theory out of those reviewed. Rather, the organizational aspects of human factors need to be understood on a case-by-case basis. The high reliability theorists select cases which strongly suggest the credibility of their approach to organizational safety, at least *for those cases* (and settings very similar to the ones they have studied). On the other hand, the ‘normal accidents’ view gains credibility from its own, different, cases and offers a plausible account of how accidents can indeed become a ‘normal’ feature of *certain settings* and are likely

to remain so. Sagan's work indicates that *for the particular setting he studied* a normal accidents view may be the more plausible. This does not rule out the possibility that, for other kinds of settings, the recommendations of the high reliability theorists might be more useful.

At the organizational level, dependability can be regarded as an issue subject to critical *organizational dilemmas*. For example, to heighten the reliability of a process against failure, one may feel (as recommended by the high reliability theorists) tempted to introduce redundancy into the process. However, as observed from the normal accidents perspective, this might heighten organizational complexity. Redundancy, then, is an organizational dilemma and, in the abstract, one can argue the case for introducing redundancy both ways. While this may be true in the abstract, *in specific cases*, it may become quite clear whether introducing redundancy is an effective protection or merely a source of 'secondary vulnerability'. However, this issue can only be resolved in the light of specific knowledge of the application domain and through continual monitoring of the effectiveness of in-service changes to processes.

The complexity of issues at the organizational level and how problems of one sort (e.g. redundancy) often have implications for problems of another sort (e.g. organizational complexity) means that potential process improvements need to be studied in the light of the full range of organizational context considerations. At the organizational level, there are no easy answers.

Nevertheless, in keeping with the other main sections of this paper, we need to be able to draw from this work a classification of vulnerabilities to error that can be applied when considering how to improve processes at an organizational level. The review of organizational factors contributing to errors and failures gives rise to the following vulnerabilities:

- 'Single points of failure' exist where a mistake by an individual can lead directly to a failure or hazardous condition
- Errors and failures propagate through the process
- Wide fluctuations in workload
- Reporting procedures and hierarchy of decision-making authority prevents rapid response to problems as they arise
- Working practices allowed to 'slip' into unsafe modes
- Failure to comply with existing safety regulations or develop new safety procedures
- Recurrent failures of a similar nature
- Potential safety hazards are allowed to pass unrecorded
- Organisational rigidities of perception and belief
- Significance of vulnerability is minimised
- There exists tight coupling of processes within a complex production system
- Process varies from project to project (ad-hoc)

This concludes this paper's survey of human sciences literature. As stated in the introduction, the survey was carried out in order to inform a process improvement method aimed at the RE process for the development of dependable systems. The following section provides a brief overview of the method, and how the findings covered in this review have been incorporated into it.

Applying human factors research in process improvement

The preceding sections have reviewed research in a number of human sciences and built up a number of categories of vulnerabilities to errors, poor performance, and unsafe behaviour in human activity. The various concepts developed in this review form the basis of a process improvement method called Process Evaluation for Requirements Engineering (PERE) which was developed in the REAIMS⁴ project. PERE encapsulates the concepts outlined above and provides a set of mechanisms that allows them to be applied in practice. This section briefly describes how the results from the survey in this paper are utilised within PERE.

PERE incorporates the survey results into a semi-structured approach to examining processes for human vulnerabilities in a number of levels. First of all, the categories of vulnerability from the literature are presented as a checklist which is structured according to the type of activity concerned (individual, social, etc.). Each entry in the checklist refers to one type of vulnerability, grouped with possible defences which could be put in place against it, and a reference to the full review to aid with points of clarification. This is necessary because of the diversity of research traditions turned to in the review. It is possible that more than one of the entries in the checklist might apply to a particular situation. It is only through providing access to the detail of the findings behind the vulnerability which will allow analysts to determine the correct course of action, in the light of the particular context of the situation they are concerned with. Figure 3 provides an excerpt from the checklist as an illustrative example.

unique reference number		reference to the section number(s) in the review where the vulnerability is discussed	
5.1	Social facilitation and inhibition (i.e. the degree and direction in which individual performance of a task is affected by being observed)	<ul style="list-style-type: none"> • Consider whether the introduction of direct supervision of activity is appropriate, whether its gains outweigh any potential performance losses as might be the case for skill-based tasks • Tend not to employ direct supervision (and prefer more indirect, deferred error-checking) for more knowledge-based tasks. 	3.2.1
the vulnerability to error		possible defence(s) against the vulnerability	

Figure 3—Excerpt from the human factors checklist

In order to assist in the application of the checklist to a particular development process, PERE also provides a step-by-step guide to focus the analysis. This takes the form of a series of questions about the process component under consideration, which guide the analyst to the relevant sections in the checklist. The result of applying the human factors viewpoint to a process is a completed table which, for each component of the process, details what vulnerabilities exist,

⁴ Requirements Engineering Adaptation and IMprovement strategies for Safety and dependability. ESPRIT project 8649 <http://www.comp.lancs.ac.uk/computing/research/cseg/projects/reaims/>

their likelihood and consequences, what possible defences could be put in place, and finally what possible secondary vulnerabilities could be introduced by the defences. Addressing secondary vulnerabilities takes into account the possibility of conflicting entries in the checklist where the defence for one vulnerability may introduce vulnerabilities of its own. The most common of these is the introduction of further complexity to the process, which can lead to organizational issues according to Normal Accident Theory. Consideration of the likelihood and consequences allows the recommendations to be prioritised according to the gravity of their occurrence.

The PERE method as a whole combines this human factors perspective on the process being evaluated with a mechanistic viewpoint. The mechanistic analysis is inspired by the Hazops (Kletz, 1992) approach which has been widely used in the chemical process industry. Whilst the human factors viewpoint is concerned with components of the process that are performed by people, the mechanistic viewpoint examines the whole process ignoring whether activities are performed by people or machines. The mechanistic viewpoint is used to build up a model of the process, which the human factors viewpoint then uses to focus on the parts involving human activity.

PERE was developed and evaluated in cooperation with the industrial partners on the REAIMS project, being applied internally to their own processes and to customers through consultancy. PERE has been applied to an organizational memory process used to generate design requirements based on experience (Mårtens *et al.*, 1996), and to the process for development of standards for an international standards body (Bloomfield *et al.*, 1996). Current work on PERE includes the investigation of support for the mechanistic viewpoint through the use of process algebra and quantitative modelling (Emmet *et al.*, 1997), as well as the development of tool support for the method. Further details on PERE are contained in the specification of the PERE module (Viller and Bowers, 1996)

Summary and conclusions

This paper has presented a review of human sciences literature relevant to the study of errors made by humans when working as individuals, in groups, and within an organizational context. This review was necessarily wide-ranging, covering a varied and disparate collection of sources. The result of this review is a number of generic error types which are applicable to different kinds of work activity. The three perspectives of individual work, group performance, and organizational factors are independent in the sense that the research reviewed here has arisen from radically different traditions in the human sciences, yet they are also related due to their shared focus on human activity, and how it can deviate from what is desired.

The work on individual error is already established as the field of human error. This approaches the problem from a cognitive psychological standpoint, and seeks to explain the slips, lapses, and mistakes that people make in terms of the cognitive processes taking place in their heads. The result of this work is a hierarchical classification of errors according to the cognitive level at which people operate when performing various tasks. They are classified into skill-based slips and lapses, and rule-based and knowledge-based mistakes.

Violations have very similar effects to errors, and can be classified according to the same skill-, rule-, knowledge-based framework. They differ from errors in that they result from deliberate action, and also because counteracting them typically requires intervention at an organizational level.

To date, the vulnerabilities of people functioning in social settings have been largely neglected by the human error and dependable systems communities. There is, however, a long tradition of social psychology research which examines the nature of group work. This includes how it compares with individuals performing similar tasks, what are the costs, and what might go wrong when working as a group. This gives rise to a broad classification of social factors to consider when examining the working of project teams, etc. in RE.

The work on organizational vulnerabilities will be more familiar to the dependable systems community, as safety-critical systems and hazardous technologies are the domains in which much of this work has been developed. In many ways, the avoidance of errors at individual and social levels depends upon factors at an organizational level. Authors in this field have proposed guidelines for how to operate in a highly reliable manner, as well as combinations of factors to avoid in order to steer clear of inevitable accidents.

The motivation for performing this review was to inform the development of a process improvement method, initially targeted at the RE process for safety critical systems. Requirements engineering is an inherently social activity, performed by individuals who carry out their work in an organizational context. The widely recognised problems associated with errors in requirements are magnified when the domain of application is considered critical in some sense, and safety critical in particular. The potential for errors to be introduced into requirements due to problems in the RE process itself leads us to the conclusion that the RE process for critical systems must be considered critical itself. We have been concerned, as a consequence, with understanding how RE might be vulnerable to errors and failures when considered from a human science perspective.

The categories of errors uncovered in this work contribute to a technique for the analysis of processes that considers how a process under evaluation might be vulnerable to errors of a human origin. They have been combined into a checklist which considers vulnerabilities under the various categories presented here, and suggests defences which could be put in place against them. This checklist itself forms part of the technique known as PERE (Process Evaluation for Requirements Engineering), which has been used by industrial partners on the REAIMS project both on their own development processes and those of customers. PERE has been applied successfully a number of times by the REAIMS project industrial partners. There is ongoing work to develop tool support for the method, and more formal support for the mechanistic analysis.

Acknowledgements

We would like to thank all of our partners in the REAIMS project for their support in developing PERE. These are GEC-Alsthom Transport, Adelard, Aerospaciale Avions, Aerospaciale Protection Systems (APSYS), Digilog, and TÜVIt. Thanks are also due to the anonymous reviewers for their helpful comments.

References

- Anderson, R., Button, G. and Sharrock, W.** (1993) 'Supporting the design process within an organisational context' *Proceedings of ECSCW'93* Milan, Italy 47-59
- Asch, S.E.** (1951) 'Effects of group pressure upon the modification and distortion of judgement' in **M. Guetzkow (ed)** *Groups, Leadership, and Men* Carnegie Press

- Bales, R.F. and Slater, P.E.** (1955) 'Role differentiation in small decision-making groups' in **T. Parsons and R.F. Bales (eds)**, *Family, Socialization and Interaction Process* Free Press
- Bendor, J.B.** (1985) *Parallel Systems: Redundancy in Government* University of California Press
- Bloomfield, R., Bowers, J., Emmet, L. and Viller, S.** (1996) 'PERE: Evaluation and Improvement of Dependable Processes' *Safecomp 96—The 15th International Conference on Computer Safety, Reliability and Security* Vienna
- Bowers, J. and Pycock, J.** (1994) 'Talking through design: requirements and resistance in cooperative prototyping' *Proceedings of CHI'94* Boston, MA 299-305
- Brooks Jr., F.P.** (1987) 'No Silver Bullet: Essence and Accidents of Software Engineering', *IEEE Computer* **20**, 4, 10-19
- Christel, M.G. and Kang, K.C.** (1992) *Issues in Requirements Elicitation* Technical Report, CMU/SEI-92-TR-12, September 1992, Software Engineering Institute,
- Davis, A.M.** (1993) *Software Requirements: Objects, Functions and States* Prentice Hall International
- Doms, M. and Van Avermaet, E.** (1985) 'Social support and minority influence: the innovation effect reconsidered' in **S. Moscovici, G. Mugny and E. Van Avermaet (eds)**, *Perspectives on Minority Influence* Cambridge University Press
- Emmet, L., Bloomfield, R., Viller, S. and Bowers, J.** (1997) 'PERE: Process improvement through the integration of mechanistic and human factors analyses' *First International Workshop on Human Error and Systems Development* Glasgow, UK 101-103
- Forsyth, D.R.** (1983) *An Introduction to Group Dynamics* Brooks/Cole
- Goffman, E.** (1961) *Asylums: Essays on the Social Situation of Mental Patients and Other Inmates* Doubleday
- Goguen, J.A.** (1993) 'Social issues in requirements engineering' *Proceedings of the IEEE International Symposium on Requirements Engineering: RE'93* San Diego, CA 194-195
- Hemphill, J.K.** (1961) 'Why people attempt to lead' in **L. Petrullo and B.M. Bass (eds)**, *Leadership and Interpersonal Behaviour* Holt, Rinehart & Winston
- Hewstone, M., Stroebe, W., Codol, J.-P. and Stephenson, G.M. (eds)** (1988) *Introduction to Social Psychology* Basil Blackwell
- Isenberg, D.J.** (1986) 'Group polarization: a critical review and meta-analysis', *Journal of Personality and Social Psychology* **50**, 1141-1151
- Janis, I.L.** (1972) *Victims of Groupthink* Houghton Mifflin
- Jirotko, M. and Goguen, J.A. (eds)** (1994) *Requirements Engineering: Social and Technical Issues* Academic Press
- Kahneman, D., Slovic, P. and Tversky, A. (eds)** (1982) *Judgement Under Uncertainty: Heuristics and Biases* Cambridge University Press
- Karau, S.J. and Williams, K.D.** (1993) 'Social loafing: a meta-analytic review and theoretical integration', *Journal of Personality and Social Psychology* **65**, 4, 681-706
- Kletz, T.A.** (1992) *Hazop and Hazan—Identifying and Assessing Process Industry Hazards* Institution of Chemical Engineers
- Kogan, N. and Wallach, M.A.** (1964) *Risk Taking: A Study in Cognition and Personality* Holt, Rinehart & Winston
- La Porte, T.R. and Consolini, P.M.** (1991) 'Working in practice but not in theory: theoretical challenges of 'high reliability organizations'', *Journal of Public Administration Research and Theory* **1**, 1, 19-47
- Lamm, H. and Myers, D.G.** (1978) 'Group-induced polarization of attitudes and behaviour' in **L. Berkowitz (ed)** *Advances in Experimental Social Psychology* Academic Press
- Likert, R.** (1967) *The Human Organization* McGraw-Hill

- Maass, A. and Clark, R.D.** (1984) 'Hidden impact of minorities—15 years of minority influence research', *Psychological Bulletin* **95**, 3, 428-450
- Maier, N.R.F. and Solem, A.R.** (1952) 'The contribution of a discussion leader to the quality of group thinking: the effective use of minority opinions', *Human Relations* **5**, 277-288
- Manstead, A.S.R. and Semin, G.R.** (1980) 'Social facilitation effects: mere enhancement of dominant responses?', *British Journal of Social and Clinical Psychology* **19**, 119-136
- Marone, J.G. and Woodhouse, E.J.** (1986) *Averting Catastrophe: Strategies for Regulating Risky Technologies* University of California Press
- Märtins, F., Schippers, H., Viller, S. and Sawyer, P.** (1996) 'Process evaluation in requirements engineering' *Design for Protecting the User: 13th Annual Centre for Software Reliability Workshop* 11th-13th September, 1996, Bürgenstock, Switzerland
- Mellor, P.** (1994) 'CAD: Computer-aided disaster!', *High Integrity Systems Journal* **1**, 2, 101-156
- Moscovici, S.** (1976) *Social Influence and Social Change* Academic Press
- Myers, D.G.** (1982) 'Polarizing effects of social interaction' in **H. Brandstätter, J.H. Davis and G. Stocker-Kreichgauer (eds)**, *Group Decision Making* Academic Press
- Paulus, P.B. (ed)** (1989) *Psychology of Group Influence* Erlbaum
- Perrow, C.** (1984) *Normal Accidents* Basic Books
- Pessin, J.** (1933) 'The comparative effects of social and mechanical stimulation on memorizing', *American Journal of Psychology* **11**, 431-8
- Quintas, P. (ed)** (1993) *Social Dimensions of System Engineering: People, Processes, Policies and Software Development* Ellis Horwood
- Rasmussen, J.** (1983) 'Skills, rules, knowledge; signals, signs and symbols; and other distinctions in human performance models', *IEEE Transactions on Systems, Man and Cybernetics* **SMC-13**, 3, 257-266
- Reason, J.** (1990) *Human Error* Cambridge University Press
- Reason, J.** (1992) 'The identification of latent organizational failures in complex systems' in **J.A. Wise, V.D. Hopkin and P. Stager (eds)**, *Verification and Validation of Complex Systems: Human Factors Issues* Springer-Verlag 223-237
- Reason, J.** (1994) *Error management in aircraft engineering: a manager's guide* British Airways Engineering, Heathrow,
- Reason, J.** (1995) 'A systems approach to organizational error', *Ergonomics* **38**, 8, 1708-1721
- Roberts, K.H.** (1989) 'New challenges in organizational research: high reliability organizations', *Industrial Crisis Quarterly* **3**, 2, 111-125
- Rodden, T., King, V., Hughes, J. and Sommerville, I.** (1994) 'Process modelling and development practice' *Proceedings of the Third European Workshop on Software Process Technology, EWSPT'94* 59-64
- Rutte, C.G. and Wilke, H.A.M.** (1984) 'Social dilemmas and leadership', *European Journal of Social Psychology* **14**, 105-121
- Sagan, S.D.** (1993) *The Limits of Safety: Organizations, Accidents, and Nuclear Weapons* Princeton University Press
- Shaw, M.E.** (1932) 'A comparison of individuals and small groups in the rational solution of complex problems', *American Journal of Psychology* **44**, 491-504
- Shaw, M.E.** (1981) *Group Dynamics: The Social Psychology of Small Group Behaviour* McGraw-Hill
- Sherif, M. and Sherif, C.W.** (1953) *Groups in Harmony and Tension: An Integration of Studies on Intergroup Relations* Octagon
- Sommerville, I. and Sawyer, P.** (1997) *Requirements Engineering: A good practice guide* John Wiley
- Steiner, I.D.** (1972) *Group Processes and Productivity* Academic Press

- Steiner, I.D.** (1976) 'Task-performing groups' in **J.W. Thibaut, J.T. Spence and R.C. Carson (eds)**, *Contemporary Topics in Social Psychology* General Learning Press
- Thomas, E.J. and Fink, C.F.** (1961) 'Models of group problem solving', *Journal of Abnormal and Social Psychology* **63**, 53-63
- Travis, L.E.** (1925) 'The effect of a small audience upon eye-hand coordination', *Journal of Abnormal and Social Psychology* **20**, 142-146
- Turner, B.A.** (1992) 'The sociology of safety' in **D.I. Blockley (ed)** *Engineering Safety* McGraw-Hill 186-201
- Van Avermaet, E.** (1988) 'Social influence in small groups' in **M. Hewstone, W. Stroebe, J.-P. Codol and G.M. Stephenson (eds)**, *Introduction to Social Psychology* Basil Blackwell 350-380
- Viller, S. and Bowers, J. (eds.)** (1996) *Process Evaluation in Requirements Engineering: PERE REAIMS Deliverable, REAIMS/WP5.1/LU035, Issue 4.1, 20th March 1996, Lancaster University, Lancaster, UK*
- Wildavsky, A.** (1988) *Searching for Safety* Transaction Books
- Wilke, H. and Knippenberg, A.V.** (1988) 'Group performance' in **M. Hewstone, W. Stroebe, J.-P. Codol and G.M. Stephenson (eds)**, *Introduction to Social Psychology* Basil Blackwell 315-349
- Zajonc, R.B.** (1965) 'Social facilitation', *Science* **149**, 269-274