Ergonomics
Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/terg20

Anaesthesia monitor alarms: a theory-driven approach

Karen E. Raymer [a b], Johan Bergström [b] & James M. Nyce [b c]

[a] Department of Anaesthesia, Faculty of Health Sciences, McMaster University, Hamilton, ON, Canada

[b] Centre for Risk Assessment and Management, Lund University, Lund, Sweden

[c] Department of Anthropology, Ball State University, Muncie, IN, USA

Version of record first published: 25 Sep 2012.

To cite this article: Karen E. Raymer , Johan Bergström & James M. Nyce (2012): Anaesthesia monitor alarms: a theory-driven approach, Ergonomics, 55:12, 1487-1501

To link to this article: http://dx.doi.org/10.1080/00140139.2012.722695

Ergonomics
Vol. 55, No. 12, December 2012, 1487–1501

Anaesthesia monitor alarms: a theory-driven approach

Karen E. Raymer[a,b]*, Johan Bergström[b] and James M. Nyce[b,c]

[a]*Department of Anaesthesia, Faculty of Health Sciences, McMaster University, Hamilton, ON, Canada;* [b]*Centre for Risk Assessment and Management, Lund University, Lund, Sweden;* [c]*Department of Anthropology, Ball State University, Muncie, IN, USA*

(*Received 15 January 2012; final version received 15 August 2012*)

The development of physiologic monitors has contributed to the decline in morbidity and mortality in patients undergoing anaesthesia. Diverse factors (physiologic, technical, historical and medico-legal) create challenges for monitor alarm designers. Indeed, a growing body of literature suggests that alarms function sub-optimally in supporting the human operator. Despite existing technology that could allow more appropriate design, most anaesthesia alarms still operate on simple, pre-set thresholds. Arguing that more alarms do not necessarily make for safer alarms is difficult in a litigious medico-legal environment and a competitive marketplace. The resultant commitment to the status quo exposes the risks that a lack of an evidence-based theoretical framework for anaesthesia alarm design presents. In this review, two specific theoretical foundations with relevance to anaesthesia alarms are summarised. The potential significance that signal detection theory and cognitive systems engineering could have in improving anaesthesia alarm design is outlined and future research directions are suggested.

Practitioner Summary: The development of physiologic monitors has increased safety for patients undergoing anaesthesia. Evidence suggests that the full potential of the alarms embedded within those monitors is not being realised. In this review article, the authors propose a theoretical framework that could lead to the development of more ergonomic anaesthesia alarms.

Keywords: alarms and warnings; advanced human–machine interfaces; anaesthesia alarms; anaesthesia equipment; equipment design; monitoring; patient safety; socio-technical systems

1. Introduction

The invasive surgical procedures which are very much a part of modern healthcare could not take place without the benefit of anaesthesia. Though preventable mishaps attributable to anaesthesia are rare, their importance is magnified by the sheer ubiquity of anaesthesia. For example, some 28 million anaesthetics are delivered in the United States each year (Stoelting and Miller 2006). Physiologic monitors are essential for the safe practice of anaesthesia, but play a disproportionate role in adverse events (Cooper *et al*. 2002). A review of the literature on anaesthesia monitor alarm design may interest readers of this journal.

Monitoring of patients under anaesthesia involves the interpretation of data related to both the performance of the anaesthetic machine and the physiologic condition of the patient. Equipment-related measurements directly capture variables of interest which include flows, circuit pressures and partial pressures of gases. By contrast, physiologic measurements present unique challenges. Monitored variables[1] are often surrogates for less easily-measured variables that more closely reflect the physiologic process of interest. Furthermore, measurements of those surrogate variables must often be done indirectly and are particularly susceptible to artefactual perturbation (Takla *et al*. 2006). Finally, accepted normal ranges (and therefore, alarm thresholds) are difficult to establish due to both a lack of clinical evidence and the degree of context-sensitivity that applies to the dynamic environment of surgery and anaesthesia (Table 1). These pragmatic realities are expressed theoretically in terms of imprecise mapping: of the system state to a symptom; of the symptom to an alarm and of the alarm to a response (Niwa and Hollnagel 2001).

Built on such a tenuous foundation, it is not surprising that dynamic fault management presents a major problem in anaesthesiology today (Seagull and Sanderson 2001, Edworthy and Hellier 2006, Hagenouw 2007, Schmid *et al*. 2011). Growing awareness of the multifaceted impact of alarms on patient safety has galvanised the efforts of several large professional organisations. The Emergency Care Research Institute (ECRI) recently identified alarms as the 'number one health technology hazard' for 2012 (ECRI 2011). The Association for the

*Corresponding author. Email: karenraymer@cogeco.ca

ISSN 0014-0139 print/ISSN 1366-5847 online

http://dx.doi.org/10.1080/00140139.2012.722695
http://www.tandfonline.com

Table 1. Challenges of physiologic monitoring in anaesthesia.

Measured variable (surrogate)	Variable of true or greater interest	Measurement technique of surrogate variable: direct or indirect?	Artefactual influences	Certainty of 'Normal Range'
Blood pressure	Cardiac output (flow)	(1) Indirect (blood pressure cuff) (2) Direct (arterial cannula placed invasively into a peripheral artery)	Numerous: (1) incorrect cuff size (2) movement (3) incorrect transducer level	Highly uncertain:Dependent on patient and surgical characteristics.Significant lack of guiding evidence
Alveolar carbon dioxide tension	Arterial blood CO2 levels. To measure CO2 levels in the blood would require an arterial puncture. Alveolar levels can be measured non-invasively through the expired breath but there is an unpredictable gradient between alveolar and arterial carbon-dioxide levels.	Indirect: unable to sample directly from alveoli; the end of an expiration (end-tidal) best approximates the gas that was in the alveoli.	Numerous: the measured value can be underestimated if true alveolar sample is not achieved, which may occur in rapid and/or shallow breathing pattern. Even if a true alveolar sample is achieved, it may not reflect arterial CO2 levels in conditions of decreased blood flow (shock, cardiac arrest).	Well-defined , however there is a high degree of context sensitivity depending on both patient (chronic lung disease) and surgical (laparoscopic) factors
Pulse oximetry	Arterial oxygen saturation; arterial oxygen content	Indirect: compares absorption of two different wavelengths of light across superficial skin capillaries as a reflection of the percentage of oxygenated and de-oxygenated haemoglobin.	Numerous: low temperature, low blood flow, skin pigmentation or nail polish	Well-defined and not context-specific
Bispectral Index (BIS)	Depth of anaesthesia; degree of unconsciousness	Indirect: a composite of electro-encephalographic (EEG) signals measured through the scalp: actual formula is proprietary, a 'black box' for clinicians	Numerous and poorly-understood	Highly uncertain: very low sensitivity and specificity for predictive value of BIS-derived numeric. No evidence for improved outcome when using BIS monitor to measure depth of anaesthesia.

Advancement of Medical Instrumentation held a summit, involving a diverse array of experts, which resulted in a series of practical recommendations (AAMI 2011, Solet and Barach 2012).

The technical challenges of monitoring are compounded by historical and medico-legal factors. The field of anaesthesiology saw its most profound improvement in safety result from the development of two physiological monitors (the oxygen saturation monitor and the end-tidal carbon dioxide monitor) which seems to have fostered an unquestioning belief in the safety benefits of alarms (Imhoff and Kuhls 2006). Secondly, the use of alarms is mandated by law. Thirdly, anaesthesia and critical care monitors developed in parallel, lacking site-specific alarm algorithms despite distinct differences in the dynamics and context of each environment. Finally, and perhaps most importantly, alarms algorithms and limits are set by manufacturers who are motivated to ensure a zero false negative rate, often at the expense of an intolerably high false positive rate (Imhoff *et al.* 2009). One result is that despite a net safety effect that may be adverse, (e.g. Bliss and Dunn 2000, Meyer 2001, Dixon *et al.* 2007) risk may be tolerated because liability is downloaded from manufacturer to the practitioner when he or she either disables the alarm audio or becomes immune to the auditory alert (Woods 1985, Imhoff and Kuhls 2006, Weil 2009).

The advances of the past 20 years in microprocessing, artificial intelligence and human factors engineering have been applied to the 'alarm problem' in anaesthesia. These advances have seen limited penetration into the clinical setting, in large part due to the reluctance of manufacturers to change technologies in the absence of strong evidence supporting improved safety of those newer technologies (Imhoff and Kuhls 2006, Kiefer and Hoeft 2010, Solet and Barach 2012). In this review, the authors will:

- Define the nature of the 'alarm problem' in anaesthesia.
- Summarise current strategies addressing the problem.
- Review the theoretical domains that are relevant to the goal of creating anaesthesia alarms that take into account the complexity of interaction in a machine-alerted human monitored system.
- Suggest future research directions.

2. The scope of the problem: anaesthesia alarms in the clinical setting

2.1. False alarms, disabled alarms and other challenges to utility

Hagenouw (2007) traces the evolution of anaesthesia alarms from the 1960s, when the anaesthesiologist,[2] finger on the patient's pulse, *was* the monitor. A vast array of devices has been developed since then. However, any gains made by 'miniaturization and combination of parameters into one unit [have been] more than offset by the continual introduction of more monitors and devices' (Hagenouw 2007).

False alarms can be technical, when the monitor generates an inaccurate measurement, or clinical, whereby the measurement is accurate but the alarm thresholds are not appropriate for that particular patient's situation. In the critical care environment, O'Carroll (1986) showed that only 8 of 1455 alarm-soundings represented an actual critical situation. More recently, Imhoff *et al.* (2009) showed that up to 90% of all alarms in the critical care environment were false positives, while Siebig *et al.* (2010) indicated that only 15% of alarms were clinically relevant. These results have been replicated in the anaesthesia environment. Kestin *et al.* (1988) showed that only 3% of alarms indicated actual patient risk, with 75% of them being patently spurious. Schmid *et al.* (2011) found that 80% of alarms occurring during anaesthesia for cardiac surgery resulted in no clinical action. The irony of alarms designed to 'never miss' is the paradoxical effect they can have on the user: popular media takes note when a patient dies quietly adjacent to a disabled alarm (Kowalcyzk 2011). Indeed, the Food and Drug Administration (FDA) received 566 reports of deaths related to monitor alarms between 2005 and 2008 in various critical care settings (Weil 2009). Most incidences were associated with the disabling or silencing of alarm function.

Though one might wonder why a healthcare provider would take the risk of disabling an alarm, this behaviour has long been recognised. In 1985, the following passage appeared in *Ergonomics*:

> [If alarms] become a 'normal occurrence', then the alarm will become part of the routine flow of action, and it will not function to break attention away from other mental and physical activity. In aggregate, trying to make all alarms unavoidable redirectors of attention overwhelms the cognitive processes involved in control of attention and exacerbates the alarm problem. One kind of operational response to this should not really be surprising—practitioners ignore or turn off the alarms. (Woods 1985)

Phipps *et al.* (2008) identified the conditions that provoke anaesthesiologists to violate the full range of 'rules' that would seem to govern their practice. The authors draw from the theory of behavioural economics, where both rule-following and rule-breaking are seen to incur costs as well as reap rewards:

> Violations occur when a person's attempts to optimize behavioural resources are not supported by the existing rules and breaching the rules incurs a lower overall cost to these resources. Where the perceived costs of violating-and the perceived benefits of compliance- are more salient, violations are less likely to occur. (Phipps *et al.* 2008)

Adverse events where the anaesthesiologist's actions were negatively influenced by the behaviour of (activated) alarms, either directly or indirectly are much more elusive than the ones that make the headlines. Hagenouw (2007) outlines the negative effects of false alarms on the anaesthesiologist's performance which include distraction, auditory fatigue and increased workload, each of which may impede attention to 'real' clinical problems. The sheer ubiquity of alarms limits their utility: Momtahan *et al.* (1993) showed that anaesthesiologists were unable to reliably identify the source or acuity of alarms from their daily work environment (cited in Hagenouw 2007), a result which substantiates Mondor and Finley's (2003) earlier finding that the perceived urgency of auditory alarms did not match their clinical relevance. Healey *et al.* (2006) showed that the operating room environment can have many distractions and interruptions, to which false alarms contribute, and through which relevant alarms must be discerned. As disturbing as the FDA's report is, it presents just one perspective on the role of monitor alarms in patient outcomes. So the question becomes not 'How can we design alarms that anaesthesiologists *cannot* disable?', but rather, 'How can we design alarms that anaesthesiologists will not *want* to disable?'

2.2. *Human response to false alarms in the experimental setting*

Warning alarms have an intuitive appeal, as the human operator is not believed to be well-suited to sustained visual scanning (Baig *et al.* 2011). In a simulated cockpit, even an imprecise automated warning system improved pilots' performance by promoting a closer inspection of the raw data (Xu *et al.* 2007). However, this experimental warning system operated with a 17% error rate (misses and false alarms, in equal measure), providing relevance that medical device alarms do not come close to approaching.

The psychology literature has shown us that individuals adjust their responsiveness to alarms sharply according the perceived reliability of the alarms (Bliss *et al.* 1995, Parasuraman *et al.* 1997). Bliss and his colleagues applied the term 'cry wolf effect' to this phenomenon, one that persistently finds its field manifestation categorised as human error[3] (e.g. Whittingham 2004, Nyssen and Blavier 2006, Weil 2009, Ferris and Sarter 2011, Kowalcyzk 2011).

The global performance impact of alarms has been studied in a laboratory where volunteers were asked to perform a task (tracking a ball on a screen) while simultaneously monitoring a gauge denoting the status of a parallel system. The gauge had an alarm, of variable reliability, to indicate when the system was failing. The authors found that *overall system performance* was more negatively impacted by a false-alarm prone gauge than a miss-prone gauge (Dixon *et al.* 2007).

Bliss and Dunn (2000) examined the effect of increased alarm and primary task workload on alarm mistrust. Their results supported earlier work showing that individuals adjusted their alarm-responsiveness according to the alarm system reliability, in a probabilistic fashion. They also found that the increased workload degraded alarm response performance.

The parallels to the healthcare setting in many of the experimental studies are limited by both the setting (non-naturalistic) and the design, which has participants performing tasks that are unrelated to the alarm task. In an experimental simulation study inspired by the observed responses of health-care professionals to the (mostly false) alarms in a medical ICU, Bitan and Meyer (2007) underscored the complex role that warning signals can play in guiding the operator's actions. Operators' responses (compliance, reliance and overall frequency of interventions) depended on both the characteristics of the warning system as well as how frequently the system required intervention. Even in an artificial environment, they were able to show that a warning signal was integrated with other data from the external environment and the operators' own internal resources.

2.3. *Anaesthesiologist and monitor alarms: an 'uneasy relationship'*

Though the operating room makes for a disobliging laboratory, several observational trials have added to our understanding of the 'uneasy relationship' (Hagenouw 2007) that anaesthesiologists have with alarms. Seagull and Sanderson (2001) found that the response to alarms varied significantly depending on the type of surgery and phase of anaesthesia, though the most common response across all phases was 'ignoring'. Alarms were most frequently ignored during anaesthetic induction and emergence, the most dynamic and treacherous phases of care, but also the phases that engage the anaesthesiologist most directly with the patient. The authors identified 'the inability to filter out irrelevant alarm information in a useful way' as a major annoyance for practitioners, as silencing measures had to be non-selectively applied to all alarms. Although their data argued for improving context-sensitivity of alarms, the authors highlighted the major limitation of machines whose algorithms aim to respond to context: those machines will predictably fail as conditions (patient or surgical) arise that cause the machine's understanding of the context to be incorrect (Seagull and Sanderson 2001).

Smith *et al.* (2003) used a variant of ethnography to explore how anaesthesiologists enact sense-making in response to alarms. They found that anaesthesiologists seamlessly integrate information from monitors and the physical examination of the patient, drawing information from one or both sources according to the circumstances, the context and the anaesthesiologist's own expectations. The attitude towards alarms highlighted the novice-expert divide, with novices responding reflexively to outlying parameters while experts used the electronic data along with their own internal resources to form an image of the patient's current condition and where that condition might be headed. In experts, the authors observed scepticism towards out-of-range monitor data. Even when not patently false, alarms were often redundant, the anaesthesiologist having already detected the changing clinical state. The authors concluded that while a monitor may meet its manufacturer's specifications, it requires significant input from the practitioner to enable it to 'work' within the human–machine system. Corroborating findings were published by Wright *et al.* (2011) who examined the perceived value of alarms in alerting the anaesthesiologist (or nurse-anaesthetist) to the onset of a crisis. Among nine critical incidences discussed, none of the six anaesthesia providers was able to recall being alerted to the problem by an alarm. The importance of waveform interpretation and the use of multiple data points in recognising the onset of a crisis were recurring themes among the interviewees. Importantly, the confusing nature of the many false alarms that occurred, even in a real crisis, was another recurrent theme. Nyssen and Blavier (2006) examined the nature of error-detection in the anaesthesia setting and found that routine monitoring of the environment was the most frequent method that resulted in detection of errors. Different detection methods were used for different types of errors, with alarms being associated most frequently with the detection of technical and procedural errors. Echoing the work of Smith *et al.* (2003), this study showed that experts had a richer repertoire of error detection methods compared with novices.

Clearly, though anaesthesiologists are mandated by law to use alarms, they cannot be mandated to find them useful. The technological advances of the previous two decades have yielded scant onsite progress in the ability of alarms to support the joint human–machine cognitive system. What stands in the way of change are not technological issues. Rather it is the risk-reduction imperative of equipment manufacturers results in a commitment to both the status quo and a zero- false-negative rate. Accordingly, anaesthesiologists today are using systems that would be easily recognised by the pioneers of the specialty who practiced when those physiological monitors were first developed and put to use (Kiefer and Hoeft 2010).

2.4. Current strategies for improving alarm performance

Imhoff and Kuhls (2006) reduced the problem to two issues: the accurate identification of conditions that require the operator's attention and the effective communication of those conditions. Existing strategies aimed at improving alarm performance have been outlined in detail (Imhoff and Kuhls 2006) and can be categorised within that framework:

Detection and decision approaches (to reduce technical and clinical false positives):

(1) Technical measures that improve device hardware.
(2) Simple statistical methods, such as linear filtering or weighted averaging of data.
(3) Sophisticated approaches that utilise artificial intelligence and fuzzy logic.
(4) Methods that integrate device data. For example, an ECG asystole alarm could be overridden if the oximeter waveform showed pulsatility.
(5) Methods that constrain alarm responsiveness. Watson and Sanderson (1998) modelled that eliminating low-priority alarms during induction and emergence would reduce the alarm rate by 70 and 90% respectively (cited in Seagull and Sanderson 2001). The effectiveness of this approach is limited by the degree to which pre-programmed alarm priorities could successfully reflect the moment-to-moment clinical importance of those alarms through a range of clinical situations. Graham and Cvach (2010) described a quality improvement project undertaken in an ICU setting in order to address perceived alarm desensitisation on their unit. The intervention, whereby nurses were trained to individualise patients' alarm parameter limits, resulted in a 43% reduction in critical alarm incidence. Gorges *et al.* (2009) showed that delaying alarms by 19 seconds reduced the incidence of critical alarms by 67%.

User interface and human factors approaches:

(1) Alarm salience: the International Electrotechnical Commission (IEC) has attempted to increase the salience of alarms by stipulating the standardisation, priority-coding and validation of the melody attached to each individual alarm. Unfortunately, very low recognition rates have been reported, even after training (Sanderson *et al*. 2006, Wee and Sanderson 2008). Speech alarms are highly salient but limited by their potential to cause anxiety in conscious patients (Edworthy and Hellier 2006).
(2) Interface: efforts have centred on attempts to reduce demands on auditory perception and visual scanning while taking into account the requirements for multi-tasking in the anaesthesiologists' workflow. Head-up multisensory displays have been modelled after those used by fighter-pilots (Kiefer and Hoeft 2010). Other efforts have focussed on hybridisation of monitors and alarms, where physiological data is continuously broadcast, through sound (sonification) or vibrotactile display, allowing the clinician to decide the boundaries of safety in any given situation (Ferris and Sarter 2011).

3. Theoretical foundations

3.1. Signal detection

3.1.1. Single-stage signal detection theory

The ideal anaesthesia monitor alarm would detect all abnormal events (100% sensitivity) without ever sounding an inappropriate alert (100% specificity). Unfortunately, that 'all reasoning and decision-making takes place in the presence of some uncertainty' (Heeger 1997) is a truth that can be applied to machines as well as to the humans who design them. During the design of any alarm system, choices are made that result in a trade-off between sensitivity and specificity. Signal detection theory helps us to understand the factors that both inform those choices and determine their impact on performance.

Most physiologic processes can be represented as occurring within a normal range. Alarms are triggered when the reading lies outside that preset range. Each machine response can be categorised in one of four possible ways (Figure 1).

High and low alarm limits present separate signal detection challenges, creating additional complexity when compared to an alarm, such as a smoke alarm, that alerts abnormality in a single direction. Heart rate (HR) is of the most fundamental variables monitored under anaesthesia. The principles outlined in Heeger's (1997) review of signal detection theory can be exemplified using a HR monitor alarm.

Even when the HR is normal, an outlying measurement may occur to trigger an alarm. This would be classified as *noise*. Figure 2 shows hypothetical probability of occurrence curves for a machine that is monitoring HR in a group of normal patients and a group of tachycardic[4] patients. If the 'noise' and 'signal plus noise' distributions shared no common territory then the monitor would reliably discriminate, providing only 'hits' and 'correct rejections'.

The discriminability index (d' = separation/spread) represents how effectively the monitor can distinguish signal from noise. Improving the strength of the information will increase the separation between curves.[5] Reducing noise will narrow the spread of each curve. Either manoeuvre will serve to reduce the overlap between the two curves,

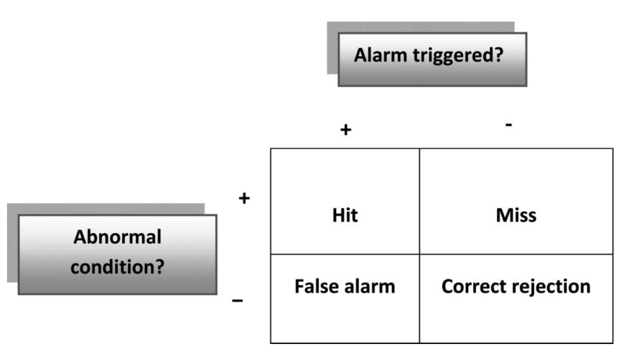

Figure 1. Categories of alarm performance.

increasing d'. It is such an overlap that forces compromises to be made, compromises that are defined by the 'criterion'.

The criterion is the strength of response required to trigger an alarm. The alarm-designer sets the 'criterion' according to the perceived importance of avoiding a miss.[6] A very low criterion will ensure that few or no abnormalities are missed, but will also result in many false positives. Figure 3 shows us the impact of shifting criteria, while highlighting the reality that as long as there is overlap between curves, no criterion can produce zero misses and zero false alarms.

3.1.2. *Signal detection in a machine-alerted human monitoring system*

With devices like anaesthesia monitor alarms, the machine is just one part of a human machine-alerted monitoring system. Sorkin and Woods (1985) expand signal detection theory to model a two-staged monitoring system in an effort to understand how the characteristics of both human and machine subsystems impact on the performance of the combined human–machine system. They developed a series of modified receiver operator curves to define the properties of a two-staged system and characterise theoretical interactions between the automated and human monitors within that system. By necessity, assumptions were made. Each of the following findings should be viewed as reflecting a specific set of assumptions:

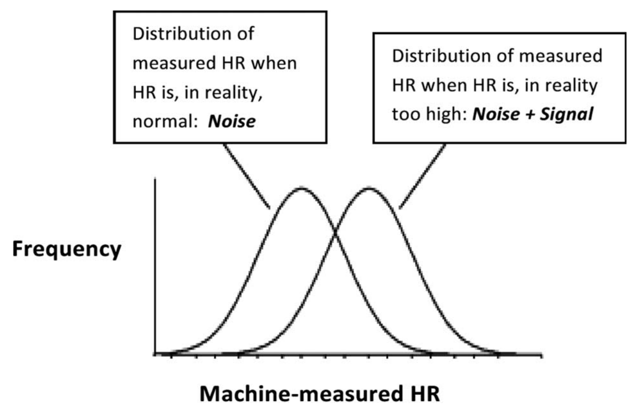

Figure 2. Probability of occurrence curves for normal and tachycardic patients.

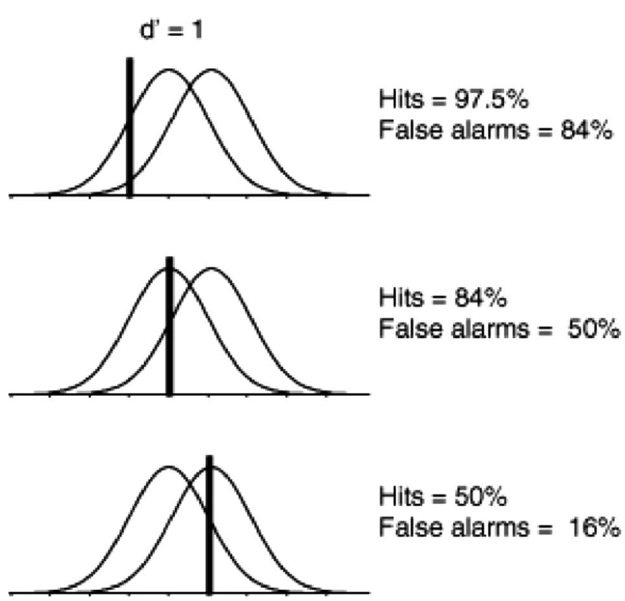

Figure 3. Effect of shifting the criterion. Source: Heeger (2007).

(1) Assuming equivalent sub-system discriminability (d'), overall system performance is 1.414d'. This 41% advantage over the performance of a single-staged system falls short of the sum of the performance of the two parts.[7]
(2) If the human operator increased his criterion in response to a high false positive rate (low criterion) in the machine, the dual system would cease to be effective except within a very limited range of low output rates from the machine.
(3) When the operator is overloaded (by frequent alerts in the face of other concurrent roles), he or she might use strategies to task-shed, either by attending only a subset of machine responses or by observing all alerts with reduced detectability (d').

The authors postulate that in the first and third situations, the human's performance (criterion choice and discriminability) are *dependent on* the machine criterion. In this setting, the system would become constrained, optimally performing within very limited boundaries.

This highly theoretical research leads to a most practical conclusion: the criterion position that optimises performance of a two-stage detection system may be different than that which optimises performance of a single-staged system (Sorkin and Woods 1985). Accordingly, in a system where a machine is alerting a human supervisor, the machine's criterion should be set to optimise overall system performance rather than the performance of the machine alone. Sorkin and Woods' (1985) work was referenced in a later study by Bustamante *et al.* (2007), who used signal detection theory to examine the effects of varied alarm thresholds and workloads on human performance of a complex task. As expected, their results confirmed that humans respond more quickly to alarms when using a system with the lowest rate of false alarms. Unexpectedly, and contradicting the earlier work of Bliss and Dunn (2000), they found that the benefit on overall performance was lost, especially under high workload conditions, where the operators tended to miss conditions that were not alarmed. The authors postulated that when the system was more reliable (low false alarm rate), operators tended to rely on it more, neglecting to account for the inherent corollary, that the system would miss more problems (Bustamante *et al.* 2007). It is important to note that the system with the highest rate of false alarms in this study had a 65% false alarm rate, still much lower than those reported in health care. We don't know what impact a 90% false alarm system would have had on performance in the high workload condition in this experimental setting, or how these factors would manifest their impact in a naturalistic setting.

The two-staged signal-detection model, with machine and human operating in series, must be applied cautiously to the anaesthesia environment. This is because an anaesthesiologist often acts in *parallel* to the machine, monitoring the machine's responses as well as the noisy raw input data. The anaesthesiologist's interpretation of that mutually-available data is modified by his or her own biases, past experiences, and any number of other human factors. Importantly, the anaesthesiologist's interpretation of data is also informed by his or her access to other monitor information and trends.

The discussion of two-staged signal detection hints at the complexity of interplay between alarm, operator, task and workload. To better understand how the machine might influence and improve the performance of the human operator, we now turn to the cognitive systems literature.

3.2. Cognitive systems engineering

Early machines were designed to extend the (human) operator's *physical* capabilities. The impact of the physical relationship between human and machine was addressed through the field of ergonomics. With the advent of microprocessors, machines became an extension of the operator's *mental* function. The architecture of the interaction between human and machine was such that the human operator no longer controlled the machine, as with a machine that manufactured widgets, but instead 'now had to control a process-or to monitor a self-controlling process' (Hollnagel and Woods 1983).

With the introduction of cognitive systems engineering (CSE), Hollnagel and Woods (1983) addressed the cognitive impact of the machine on the human–machine system. Together, the human and machine comprise a joint cognitive system (JCS)[8]: it receives data from the environment then makes a plan to achieve a goal. Its arrived-upon plan is both data-driven and concept-driven, the latter based on its own knowledge of the world apart from the specific input data at any given moment. Hollnagel and Woods (1983) outline principles of JCS design that remain relevant today:

- Human beings are naturalistic (not 'rational') decision makers.
- Human fault-finding performance varies with the degree to which the operator is engaged within the system control-loop: *whether system- engagement of the operator improves or degrades performance depends on workload.*

- The total performance of the JCS cannot be explained by the performance of the individual sub-systems; accordingly, tasks should be assigned (to human or machine) with the function of the whole system in mind.
- Human 'errors' are usually a symptom of a flawed system. Merely automating what humans are 'not good at' can merely shift or expose other system vulnerabilities.

Concept-driven behaviour hinges on one's internal representation of the environment. That the human has a 'model' of the system in which he or she operates (that may be less than complete or accurate) seems intuitive. What this model is and the role it plays in human–machine interaction has been the issue on which CSE research has traditionally focused. Hollnagel and Woods (1983) have extended this logic in positing that the machine itself has its own 'image' of the human operator built into the very design of the machine. Unfortunately, 'the system's image of the user ... is virtually never explicitly designed to enhance the joint function of [human] and machine', and thus creates 'mismatches between [hu]man and machine' (Hollnagel and Woods 1983). The authors stress that machine-designers must first understand human cognitive function in order to design machines that can be compatible with and support human performance and ultimately, that of the JCS. Without that key step, the human will be left to adapt to the machine, with predictable results. Applying these concepts to the field of anaesthesia monitors, Kiefer and Hoeft (2010) state that anaesthesia-related adverse events are not random accidents, but instead are the direct result of the systemic constraints imposed on the interaction between human and machine. The authors go on to emphasise that machine design must accommodate human cognition, rather than the other way around.

3.3. Data overload

Nearly 30 years after Hollnagel and Woods' work, designing machines that adapt to human cognition remains an elusive goal. One pervasive example of machine–human mismatch is data overload. Woods *et al.* (2002) analyze data overload within the cognitive systems framework.

While today we can collect ever increasing volumes of data, technology has been less effective in helping the human unearth meaning from the mound of data to find the item that might advance or hinder the operator's goal(s). Despite considerable effort, further technological advances seem to have only deepened this 'data-availability paradox' (Woods *et al.* 2002). In the operating room environment, the development of more (and more sophisticated) physiologic monitors creates more displayed data and ultimately, more triggered alarms. Unless those monitors are developed and applied with attention to human–machine interaction, they risk contributing to the widening gap between the quantity and meaningfulness of data (Kiefer and Hoeft 2010). The characterisation of data-overload as a problem of 'finding the significance of data when it is not known a priori what data from a large data field will be informative' (Woods *et al.* 2002) has particular relevance to this discussion as monitor alarms themselves can be seen as a data-overload condition, where the alarms themselves are the data items, and the operator has to extract meaning from noise.

One approach to the data-overload problem has been to harness the power of computers, which, unlike humans, are able to siphon large volumes of diverse data streams both quickly and accurately (Baig *et al.* 2011). Though appealing, this approach can overlook the true challenge of data overload which lies in context sensitivity: data finds its meaning in its relationship to other data, to goals and to expectations. Rather than seeing humans as a weak link, Woods *et al.* (2002) assert that people are uniquely capable of directing their attention towards the data within their perceptual field that is most likely to be relevant, even when what is relevant reflects changing contexts and goals. Woods *et al.* (2002) examine how people can focus on what is important to better understand how machines get in the way of those natural abilities. To this end, the authors highlight three fundamental skills unique to human cognition:

- Perceptual organisation (lumping the data into meaningful groups).
- Attentional control (knowing where to focus; keeping focus simultaneously granular and global).
- Anomaly-based processing (noting departures from typicality as well as relative changes or trends).

Woods *et al.* (2002) conclude that machines must be designed to *support* (rather than replace) the human's innate context sensitivity and only then will the problem of data overload recede. They outline necessary design constraints that emphasise the organisation and conceptualisation (but not the minimisation) of data. Maintaining that humans alone are able to achieve context-sensitivity, they identify attempts to offload context sensitivity to machines as misguided, yielding brittle and limited results. Attempts to 'finesse' context sensitivity are categorised as follows:

- Reducing the amount of data available to the operator.
- Static prioritisation – deciding in advance which classes of data will take priority.
- Syntactic approaches where keywords within the content trigger alarm priority.
- Attempting to build in intelligence into the machine so that it can present only the 'important' data.

The latter technique, artificial intelligence, has been pursued in medical decision-making. One example is 'fuzzy logic', an approach whose appeal lies in its representation of data as a shade of grey in between 0 and 1, congruent to that of the clinician, who rarely views a parameter as strictly 'normal' or 'abnormal'. In fuzzy logic, the range of possible measurements is divided into classes. Any given numeric value could have membership, of different degrees, in more than one class creating a fuzzy set for that measured variable. Expert knowledge is used to define rules that allow the eventual classification of the possible combinations of fuzzy sets in the binary fashion that an alarm requires.

There is circularity to Woods' argument, in that an alarm is itself, by definition, a machine-generated attempt at context sensitivity, *the machine trying to tell the operator what is interesting*. It is not surprising then, that methods to improve context-sensitivity, simplistic or sophisticated, are frequently applied to the alarm problem in anaesthesia. In order to better understand how to manage the problems associated with data overload in the joint cognitive activities of anaesthetic care we need to turn to the literature on the co-ordination among participants in joint cognitive activity.

3.4. Co-ordination in joint cognitive activity

Co-ordination has been defined as 'the intent to work together to align goals and to invest effort to sustain common interests' (Klein *et al.* 2005). Participants may have to relax -or even sacrifice- their own local goals while acting in support of the global goals of the group. The authors define three requirements for effective co-ordination in joint activity:

- *Interpredictability*: one must be able to predict the actions of the other in any given situation. Interpredictability results from parties having a 'shared script' and is contingent on each party being able to view the situation from the other party's perspective.
- *Common ground*: the parties must negotiate a shared set of knowledge, beliefs and assumptions which will allow the parties to act independently in the pursuit of their joint activity. In order to achieve common ground, each party must have a baseline understanding of the others' roles, skills, competencies and goals. Furthermore, as the activity unfolds, the stance of each participant (e.g. his or her perception of time pressure, level of fatigue and competing priorities) must be observable to the others.
- *Directability*: one party must be able to influence the actions of other parties despite shifting goals and conditions.

During joint activity, communication between participants can occur on two levels: there is the 'task-work' that relates directly to the joint activity and the 'teamwork' which serves to keep the choreography of the activity on track. Co-ordination is particularly important as the activity transitions through phases. Participants must have ways of signalling to each other and judging when the other participant is at an 'interruptible' phase of his or her own task(s). The toll exacted, in time and effort, by these efforts to maintain choreography has been termed 'coordination cost'.

The authors identify the loss of common ground as the most frequent cause of joint activity breakdown. Some of the most salient causes of loss of common ground are:

- Discrepant access to data.
- Lack of clarity regarding the joint goal.
- Lack of awareness of differing stances between workers (workload, competing priorities).
- Confusion over who knows what.

Klein *et al.* (2005) highlight six unique challenges[9] that must be overcome to allow successful JCS activity. For example, the machine that is designed to be responsive to context is simultaneously less predictable to the human operator. Another challenge lies in the area of attention management: how can the machine communicate status changes to its human partner while still allowing him or her to work without unnecessary interruption?

The analytic focus on co-ordination shifts the object of analysis from humans and technology as arbitrary self-contained entities of the joint operating system. Instead, humans as well as technological artefacts are all actors in the joint cognitive activity in which they engage. Approached this way, an alarm is one of many measures used by one actor to direct other actors. Examples of questions to pose when analyzing alarms from the perspective of a JCS therefore include:

- Is the sounding of the alarm based on a common ground?
- Does the alarm behaviour reflect goal conflicts between actors?
- Is key information made available to all actors to inform their behaviour?
- How predictable are the actions of one actor to other actors?
- How does mutual directability impact on an individual's workflow?
- How high are the co-ordination costs between human and machine?
- Is there a mechanism for actors to be aware of the stance (e.g. workload, fatigue, level of engagement with the monitored unit) of other actors?

4. Application of theoretical principles: examples from the literature

The challenges and importance of integrating the theoretical principles of signal detection theory and CSE into the design of JCSs have received significant attention in non-medical industry (e.g. Roth and Pew 2008). A recent review on managing alarm fatigue in cardiac care (Solet and Barach 2012) recommended that the body of knowledge regarding the impact of technology on human performance that has been developed over the past 50 years must now be incorporated into alarm design. Despite the recognition of the applicability of these principles, most of the published research remains theoretical without resulting in concrete design recommendations. Furthermore, such literature within the medical field is negligible. A selection which includes studies from non-medical industries and from the larger sphere of decision-support illustrates how the theoretical principles discussed in this review could inform anaesthesia alarm design on a practical level.

Observable themes span the range of abstraction hierarchy: refining criterion selection; using fuzzy logic to mimic context sensitivity; achieving a detailed understanding of user needs and work domain; designing machine function to reflect human cognition; refining attentional control and function allocation; creating machines that recognise and adapt to the human stance; adjusting operator engagement according to workload; presenting data in a simultaneously global and granular data fashion and creating observability and directability.

Ansermino *et al.* (2008) incorporated principles of signal detection theory in a study examining ventilator alarms. Recognising that criterion choice is a balance between false alarms and missed events, the authors suggested that criterion performance could be optimised by providing information from either a different parameter or a different time. They interviewed experts to gain consensus on how other monitored parameters or trend changes could help pinpoint the best criteria for a series of ventilatory events, using the heuristics or 'rules of thumb' that anaesthesiologists applied to those events.

Masalonis and Parasuraman (2003) applied fuzzy signal detection theory techniques to air traffic control and found that by assigning each event a fuzzy membership (between 0 and 1) they were able to reduce the computed false alarm rate in both human and machine systems. Ortero *et al.* (2009) applied this approach to 16 commonly-occurring alarms in the ICU setting, using expert opinion to define the classification of events. They were able to show an astounding 7% false positive rate over 912 triggering events although the 'miss rate' was not measured. Stevens *et al.* (2012) were able to reduce false alarms by 59% in ICU patients post coronary artery surgery by classifying events according to the fuzzy set generated through the analysis of four monitored variables. Baig *et al.* (2011) developed an elegant fuzzy logic approach to the early detection of hypovolemia in surgical patients.

Pott *et al.* (2005) underscored the importance of achieving a deep understanding of the human's cognitive work to inform machine design. They developed a detailed cognitive work analysis of the anaesthesiologist with emphasis on naturalistic decision-making and situation awareness. They identified two classes of tasks that the anaesthesiologist undertakes (maintenance and repair), with the repair task being subdivided into situation assessment, familiarity-searching, urgency-assessment and diagnosis. They suggested that the machine should be designed to understand the nature of the cognitive work that the anaesthesiologist is undertaking at any given moment, including recognising transitions from maintenance to repair. The performance of the machine, then, would adapt according to the human state: distractions would be reduced during high workload states, while communication would be increased during low workload states in order to enhance vigilance. Presentation and volume of information would be adjusted according to the information processing capacity of the user in that condition. Decision-support methods could be designed to accommodate the varying needs of novice versus expert anaesthesiologists, or those of the fatigued practitioner.

Dadashi *et al.* (2012) undertook a cognitive work analysis of rail operators which included a deconstruction of the cognitive process involved in interpreting an alarm. The resulting understanding yielded concrete design recommendations. For example, they noted that the duplicate methods of alarm notification were redundant as the operator only attended the notice in the 'banner'. They did note that operators spent time categorising and filtering the alarm type, and recommended that the alarm be presented in a way that streamlined this process. The fact that the operator sought out a wide range of data in order to accept or reject the alarm led to the recommendation that specific situational data be presented along with the alarm.

Dalal and Kasper (1994) examined the impact of cognitive coupling on the effectiveness of a decision-aid. They demonstrated that depending on the cognitive nature of the problem, the style of decision-aid (analytic or heuristic) could either enhance or detract from the overall performance of the human–computer system compared to the human working alone.

Niwa and Hollnagel (2001) deconstructed alarm response into evaluation and response phases and hypothesised that when the sum of times required for evaluation and decision exceeded the available time to act, the operator would lose control of the situation. They suggested that the machine could be designed to recognise and adapt to loss of operator control by taking over more of the operator's functions and presenting information in an 'overview' fashion. When re-establishment of control was detected, function allocation and data presentation would return to baseline.

Dorneich *et al.* (2012) described a method to incorporate 'etiquette' in the timing of interruptions of the human by an automated system. The machine modified its interactions with the human according to the classification of the human's cognitive state, based on brain and heart sensors. The timing of interruption, therefore, was based on an algorithm that considered message priority, user workload and system state.

McKenna and Elm (2006) describe the application of CSE within the intelligence community, including a case where a decision support algorithm was designed to be sufficiently transparent ('observability') that the human operator could change the algorithm according to the immediate context ('directability'), allowing the operator to simultaneously access the information-processing capabilities of the computer and the context-sensitivity of the human.

Despite the limitations of comparison between some of these works and anaesthesia alarms, they highlight the separation between human and machine that exists in current alarm design. Furthermore, they shed light on what the application of the principles summarised in this review could bring to modern anaesthesia alarm design at this point in history, where the problem has been amply recognised but the way forward not fully elucidated.

5. Conclusion: towards a more theory-driven alarm technology

Physiologic monitors have increased safety for patients undergoing anaesthesia. The alarms embedded within those monitors, however, remain problematic. Alarms in anaesthesia have been applied to the clinical setting with little apparent attention to the more complex aspects of human–machine interaction. The prominence of the industry-led agenda has played an important role in creating the alarms that anaesthesiologists use today. The prioritisation of 'never missing a hit' creates a high false positive rate, leading to a range of behaviours in humans that can have a negative impact on overall performance. To date, methods to address the alarm problem in anaesthesia have been carried out within the existing framework of alarm operation: improving physiologic data detection and streamlining user interface. Though key to the development of alarms that could respond more accurately and less intrusively, these granular level efforts presume linearity in human–machine interaction (Figure 4).

The role of the alarm in a dual monitoring system calls for a more nuanced approach to design and development than the simple alerting of the operator to out-of-range variables. Several literatures applicable to this problem have not been exploited and could provide a theoretical base from which to build alarms that function as a 'team players' within a machine-supported human monitoring system. The authors suggest that future research be directed towards understanding how dual-staged signal detection theory and CSE could apply to anaesthesia alarms. Together, these well-developed theoretical frameworks could provide a foundation from which a more appropriate strategy for alarm design could be built. Some important aspects in designing future alarm systems for anaesthetic care could include:

- Understanding the goals and consequences of choosing a particular decision criterion for when to activate a specific alarm. This includes understanding the tension(s) that can exist between the single-staged case and the

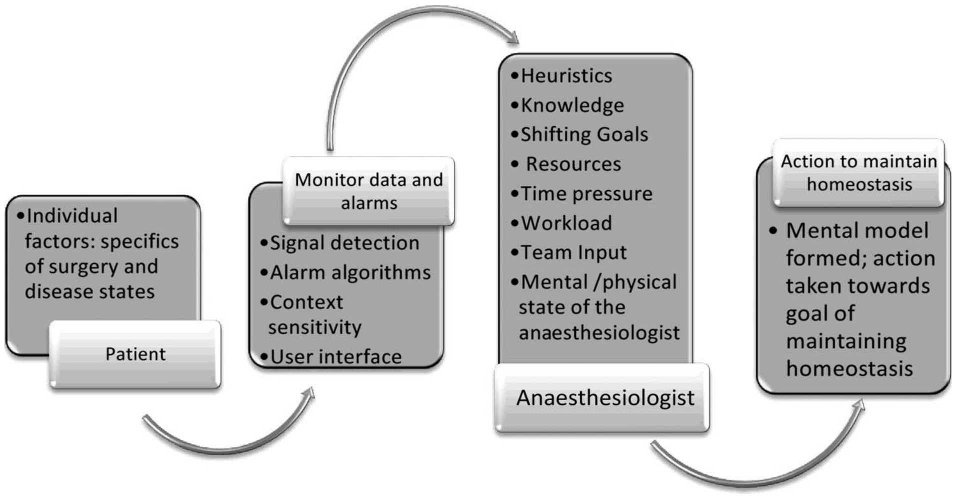

Figure 4. Linear modelling of the interaction between the patient, machine and anaesthesiologist.

dual-staged case as well as exploring the use of non-binary (fuzzy) sets to more closely mirror human decision-making.

- Incorporation of central concepts from the field of CSE in design and development activities. These include:
  - Joint cognitive activity.
  - User image.
  - Data overload.
  - Co-ordination.

The need for a theoretical foundation for anaesthesia alarm design and the desire for speedy, practical improvements are not entirely mutually exclusive goals. The latter goal supports the former, in that much of what we do know about alarm design is currently of limited utility to the anaesthesia application, given their uniquely high false positive rates. Nonetheless, while the recent attention drawn to medical device alarms is deserved, it runs the risk of driving the agenda towards quick fixes if the importance of fundamental principles of human–machine interaction is not recognised.

Anaesthesia alarms and anaesthesiologists participate in a joint activity whose goal is to ensure that the human actor is aware of all instances in which the patient's physiologic parameters extend outside safe limits during the dynamic conditions that characterise illness, surgery and anaesthesia. Monitoring of physiologic variables under anaesthesia involves uncertainty at each step in the process of generating numeric data from organ function. Learning how to design monitors that can support such complex clinical work can have pragmatic and analytic benefits. Such a design agenda would not only improve the monitoring of patients; it would also be a contribution to CSE literature and the design of technologies outside medical settings.

Notes

1. Monitored physiological variables include heart rate (HR) and rhythm, ST segment analysis, blood pressure, oxygen saturation, end-tidal carbon dioxide tension, temperature and bispectral index.
2. Much has changed since the early days of anaesthesiology, including the fact that in many jurisdictions, nurse anaesthetists monitor the patient, under the supervision of an anaesthesiologist who may or may not be present. Indeed, other healthcare providers, including the surgeon, may be in the position of responding to anaesthesia monitor alarms. The authors recognise this diversity, but for simplicity use 'anaesthesiologist' as a collective term.
3. Whittingham (2004, pp. 171–175) described a fatal rail crash where the conductor repeatedly 'acknowledged' an alarm that was indicating the increasing proximity of a train ahead, without taking the corrective action of slowing the train. The authors concluded that his action of depressing the acknowledge button had become a conditioned response, a phenomenon that was not uncommon amongst conductors. They also found that the alarm system was 'inadequate for its intended purpose'. The conductor survived, but served 18 months in jail for manslaughter.
4. Tachycardia is the term for an elevated HR (in adults, >100 beats per minute).
5. It is not easy to make improvements in d', which is often fixed by both hardware limitations and our own limited understanding of the ways in which a true event differs from normal with respect to the patterns it will present to the monitor. This is an area in which more research clearly needs to be done if clinically appropriate alarms are to be developed.

6. The criterion determination is influenced by the a-priori probability of a signal and the risk/benefit ratio of each of the decision outcomes elaborated in Figure 1.
7. Pollack and Maddans (1964) examined the performance of a two-stage detection system, finding that optimal system performance occurred when each stage had equal discriminability and equal criteria. Interestingly, they found that system performance was 1.2d', just slightly less than the theorised maximal performance (cited in Sorkin and Woods 1985).
8. The machine itself is a cognitive system, as is the operator. Working together, they are a JCS.
9. The six challenges outlined by Klein *et al.* (2005) for joint human–machine activity are:

   (1) Achievement of basic compact.
   (2) Mutual predictability.
   (3) Goal negotiation.
   (4) Phase co-ordination.
   (5) Attention management.
   (6) Control of co-ordination costs.

References

Ansermino, J.M., Dosani, M., Amari, E., Choi, P.T., and Schwartz, S.K.W., 2008. Defining rules for the identification of critical ventilatory events. *Canadian Journal of Anesthesia*, 55 (10), 702–714.

Association for the Advancement of Medical Instrumentation (AAMI), 2011. *A siren call to action. Priority issues from the 2011 medical device alarms summit* [online]. Available from: http://www.aami.org/alarms/Materials/2011_Alarms_Summit_publication.pdf [Accessed 24 February 2012].

Baig, M.M, GholamHosseini, H.G., Kouazani, A., and Harrison, M.J., 2011. Anaesthesia monitoring using fuzzy logic. *Journal of Clinical Monitoring and Computing*, 25, 339–347.

Bitan, Y. and Meyer, J., 2007. Self-initiated and respondent actions in a simulated control task. *Ergonomics*, 50 (5), 763–788.

Bliss, J.P. and Dunn, M.C., 2000. Behavioural implication of alarm mistrust as a function of task workload. *Ergonomics*, 43 (9), 1283–1300.

Bliss, J.P., Gilson, R.D., and Deaton, J.E., 1995. Human probability matching behaviour in response to alarms of varying reliability. *Ergonomics*, 38, 2300–2312.

Bustamante, E.A., Bliss, J.P., and Anderson, B.L., 2007. Effects of varying the threshold of alarm systems and workload on human performance. *Ergonomics*, 50 (7), 1127–1147.

Cooper, J.B., Newbower, R.S., Long, C.D., and McPeek, B., 2002. Preventable anesthesia mishaps: a study of human factors. *Quality and Safety in Health Care*, 11 (3), 277–282.

Dadashi, N., Wilson, J.R., Golightly, D., Sharples, S., and Clarke , T., 2012. Practical use of work analysis to support rail electrical control rooms: a case of alarm handling. *Journal of Rail and Rapid Transit*. (In press).

Dalal, N.P. and Kasper, G.M., 1994. The design of joint cognitive systems: the effect of cognitive coupling on performance. *International Journal of Human-Computer Studies*, 40 (4), 677–702.

Dixon, S.R., Wickens, D., and McCarley, J.S., 2007. On the independence of compliance and reliance: are automation false alarms worse than misses? *Human Factors*, 49 (4), 564–572.

Dorneich, M.D., Ververs, P.M., Mathan, S., Whitlow, S., and Hayes, C.C., 2012. Considering etiquette in the design of an adaptive system. *Journal of Cognitive Engineering and Decision Making*, 6 (2), 243–265.

ECRI Institute, 2011. Top 10 Health Technology Hazards for 2012. *Health Devices*, 40 (11), 1–18. Available from: http://www.safeincommon.org//sites/default/files/field_document/health_devices_top_10_hazards_2012.pdf [Accessed 6 September 2012].

Edworthy, J. and Hellier, E., 2006. Alarms and human behaviour: implications for medical alarms. *British Journal of Anaesthesia*, 97, 12–17.

Ferris, T.K. and Sarter, N., 2011. Continuously informing vibrotactile displays in support of attention management and multitasking in anesthesiology. *Human Factors*, 53 (6), 600–611.

Gorges, M., Markewitz, B.A., and Westenskow, D.R., 2009. Improving alarm performance in the medical intensive care unit using delays and clinical context. *Anesthesia & Analgesia*, 108, 1546–1552.

Graham, K.C. and Cvach, M., 2010. Monitor alarm fatigue: standardizing use of physiological monitors and decreasing nuisance alarms. *American Journal of Critical Care*, 19 (1), 28–34.

Hagenouw, R.R.P.M., 2007. Should we be alarmed by our alarms? *Current Opinion in Anesthesiology*, 20, 590–594.

Healey, A.N., Sevdalis, N., and Vincent, C.A., 2006. Measuring intra-operative interference from distraction and interruption observed in the operating theatre. *Ergonomics*, 49 (5–6), 589–604.

Heeger, D., 1997. Signal detection theory. *Journal of Physiology London*, 12 (3), 1–10.

Heeger, D., 2007. *Signal detection theory* [online]. Available from: http://www.cns.nyu.edu/~david/handouts/sdt/sdt.html [Accessed 10 January 2012]. From New York University.

Hollnagel, E. and Woods, D.D., 1983. Cognitive systems engineering: new wine in new bottles. *International Journal of Man-Machine Studies*, 18, 583–600.

Imhoff, M. and Kuhls, S., 2006. Alarm Algorithms in critical care monitoring. *Anesthesia and Analgesia*, 102, 1525–1537.

Imhoff, M., Kuhls, S., Gather, U., and Fried, R., 2009. Smart alarms from medical devices in the OR and ICU. *Best Practice & Research Clinical Anaesthesiology*, 23, 39–50.

Kestin, I.G., Miller, B.R., and Lockhard, H., 1988. Auditory alarms during anaesthesia monitoring. *Anesthesiology*, 69, 106–109.

Kiefer, N. and Hoeft, A., 2010. Display of information in the operating room. *Current Opinion in Anesthesiology*, 23, 772–777.

Klein, G., Feltovich, P.J., Bradshaw, J.M., and Woods, D.D., 2005. Common ground and coordination in joint activity. *In:* W. Rouse and K. Boff, eds. *Organizational simulation*. Hoboken, NJ: John Wiley & Sons, 139–184.
Kowalcyzk, L., 2011, February 13. *Alarm fatigue blamed in hospital deaths* [online]. Available from: http://www.boston.com/lifestyle/health/articles/2011/02/13/patient_alarms_often_unheard_unheeded/?page=full [Accessed 24 February 2012]. From The Boston Globe.
Masalonis, A. and Parasuraman, R., 2003. Fuzzy signal detection theory: analysis of human and machine performance in air traffic control, and analytic consideration. *Ergonomics*, 46 (11), 1045–1075.
McKenna, B.P. and Elm, W.C., 2006. *Decision support for the intelligence community: design solutions to real world problems* [online]. Available from: http://www.resilientcognitivesolutions.com/wordpress/wp-content/uploads/2007/11/solvingcommonproblems-2006.pdf [Accessed 26 May 2012]. From Resilient Cognitive Solutions' website.
Meyer, J., 2001. Effects of warning validity and proximity on responses to warnings. *Human Factors*, 43 (4), 563–572.
Momtahan, K., Hétu, R., and Tansley, B., 1993. Audibility and identification of auditory alarms in the operating room and intensive care unit. *Ergonomics*, 36 (10), 1159–1176.
Mondor, T. and Finley, G., 2003. The perceived urgency of auditory warning alarms used in the hospital operating room is inappropriate. *Canadian Journal of Anesthesia*, 50, 221–228.
Niwa, Y. and Hollnagel, E., 2001. Enhancing operator control by adaptive alarm presentation. *International Journal of Cognitive Ergonomics*, 5 (3), 367–384.
Nyssen, A.S. and Blavier, A., 2006. Error detection: a study in anaesthesia. *Ergonomics*, 49 (5–6), 517–525.
O'Carroll, T.M., 1986. Survey of alarms in an intensive therapy unit. *Anaesthesia*, 41, 742–744.
Ortero, A., Felix, P., Barro, S., and Palicios, F., 2009. Addressing the flaws of current critical alarms: a fuzzy constraint satisfaction approach. *Artificial Intelligence in Medicine*, 47, 219–238.
Parasuraman, R., Hancock, P.A., and Olofinboba, O., 1997. Alarm effectiveness in driver-centered collision-warning systems. *Ergonomics*, 40, 390–399.
Phipps, D.L., Parker, D., Pals, E., Meakin, H., Nsoedo, C., and Beatty, P., 2008. Identifying violation-provoking conditions in a healthcare setting. *Ergonomics*, 51 (11), 1625–1642.
Pollack, I., and Maddans, A.B., 1964. On the performance of a combination of detectors. *Human Factors*, 6, 523–531.
Pott, C., Cnossen, F., and Ballast, A., 2005. More than psychologists' chitchat: the importance of cognitive modelling for requirements engineering in complex and dynamic environments. *In: Proceedings of SREP'05*, 29–30 August 2005. Limerick, Ireland: University of Limerick, 163–175.
Roth, E.M. and Pew, R.W., 2008. Integrating cognitive engineering in the systems engineering process: opportunities, challenges, and emerging approaches. Introduction to the special issue. *Journal of Cognitive Engineering and Decision Making*, 2 (3), 161–164.
Sanderson, P.M., Wee, A., and Lacherez, P., 2006. Learnability and discriminability of melodic medical equipment alarms. *Anaesthesia*, 61, 142–147.
Schmid, F., Goepfert, M.S., Kuhnt, D., Eichhorn, V., Diedrichs, S., Reichenspurner, H., Goetz, A.E., Reuter, D.A. 2011. The wolf is crying in the operating room: patient monitor and anesthesia workstation alarming patterns during cardiac surgery. *Anesthesia and Analgesia*, 112 (1), 78–83.
Seagull, F.J. and Sanderson, P.M., 2001. Anesthesia alarms in context: an observational study. *Human Factors*, 43, 66–78.
Siebig, S., Kuhls, S., Imhoff, M., Gather, U., Scholmerich, J., and Wrede, C.E., 2010. Intensive care unit alarms: how many do we need? *Critical Care Medicine*, 38, 451–456.
Smith, A.F., Goodwin, D., Mort, M., and Pope, C., 2003. Expertise in practice: an ethnographic study exploring acquisition and use of knowledge in anaesthesia. *British Journal of Anaesthesia*, 91, 319–328.
Sorkin, R.D. and Woods, D.D., 1985. Systems with human monitors: a signal detection analysis. *Human–Computer Interaction*, 1, 49–75.
Solet, J.M. and Barach, P.R., 2012. Managing alarm fatigue in cardiac care. *Progress in Pediatric Cardiology*, 33, 85–90.
Stevens, N., Giannareas, A.R., Kern, V., Viesca, A., Fortino-Mullen, M., King, A., and Lee, I., 2012. Smart alarms: multivariate medical alarm integration for post CABG surgery patients. *In: Proceedings of the 2nd ACM SIGHIT International Health Informatics Symposium*, 28–30 January 2012, Miami, FL. New York, NY: ACM, 533–542.
Stoelting, R.K. and Miller, R.D., 2006. *Basics of anesthesia*. New York, NY: Churchill Livingstone.
Takla, G., Petre, J.H., and Doyle, D.J., 2006. The problem of artefacts in patient monitor data during surgery: a clinical and methodological review. *Anesthesia and Analgesia*, 103, 1196–1204.
Watson, M., and Sanderson, P., 1998. Work domain analysis for the evaluation of human interaction wtih anesthesia alarm systems. *In: Proceedings of the 1998 Australasian computer human interaction conference*, Los Alamitos, CA: IEEE Computer Society, 228–235.
Wee, A.N. and Sanderson, P.M., 2008. Are melodic medical equipment alarms easily learned? *Anesthesia and Analgesia*, 106, 501–508.
Weil, K.M., 2009, September 18. *Alarming monitor problems* [online]. Available from: http://www.fda.gov/MedicalDevices/Safety/AlertsandNotices/TipsandArticlesonDeviceSafety/ucm222022.htm [Accessed 24 February 2012]. From The Food and Drug Administration.
Whittingham, R.B., 2004. *The blame machine: why human error causes accidents*. Burlington, MA: Elsevier.
Woods, D.D., 1985. The alarm problem and directed attention in dynamic fault management. *Ergonomics*, 38 (11), 2371–2393.
Woods, D.D., Patterson, E.S., and Roth, E.M., 2002. Can we ever escape from data overload? A cognitive systems diagnosis. *Cognition, Technology & Work*, 14, 22–36.
Wright, M.C., Dawodu, M., Segall, N., Taekman, J.M., and Mark, J.B., 2011. Anesthesia crisis alerting: a qualitative study toward understanding the complexity of crisis identification. *Anesthesiology*, A741.
Xu, X., Wickens, C., and Rantanen, E.M., 2007. Effects of conflict alerting system reliability and task difficulty on pilots' conflict detection with cockpit display of traffic information. *Ergonomics*, 50 (1), 112–130.