INL/CON-12-25623
PREPRINT

Fifty Years of THERP and Human Reliability Analysis

PSAM11

Ronald L. Boring

June 2012

Fifty Years of THERP and Human Reliability Analysis

Ronald L. Boring
Idaho National Laboratory, Idaho Falls, Idaho, USA

Abstract: In 1962 at a Human Factors Society symposium, Alan Swain presented a paper introducing a Technique for Human Error Rate Prediction (THERP). This was followed in 1963 by a Sandia Laboratories monograph outlining basic human error quantification using THERP and, in 1964, by a special journal edition of Human Factors on quantification of human performance. Throughout the 1960s, Swain and his colleagues focused on collecting human performance data for the Sandia Human Error Rate Bank (SHERB), primarily in connection with supporting the reliability of nuclear weapons assembly in the US. In 1969, Swain met with Jens Rasmussen of Risø National Laboratory and discussed the applicability of THERP to nuclear power applications. By 1975, in WASH-1400, Swain had articulated the use of THERP for nuclear power applications, and the approach was finalized in the watershed publication of the NUREG/CR-1278 in 1983. THERP is now 50 years old, and remains the most well known and most widely used HRA method. In this paper, the author discusses the history of THERP, based on published reports and personal communication and interviews with Swain. The author also outlines the significance of THERP. The foundations of human reliability analysis are found in THERP: human failure events, task analysis, performance shaping factors, human error probabilities, dependence, event trees, recovery, and pre- and post-initiating events were all introduced in THERP. While THERP is not without its detractors, and it is showing signs of its age in the face of newer technological applications, the longevity of THERP is a testament of its tremendous significance. THERP started the field of human reliability analysis. This paper concludes with a discussion of THERP in the context of newer methods, which can be seen as extensions of or departures from Swain's pioneering work.

Keywords: HRA, THERP, history

1. INTRODUCTION TO HUMAN RELIABILITY ANALYSIS

Human reliability analysis (HRA) is the predictive study of human errors, typically in safety-critical domains like nuclear power generation (Swain and Guttman, 1983). Human error, in this context, describes any action or inaction on the part of an individual that decreases the safety of the system with which he or she is interacting. The term human error carries with it negative connotations (Dekker, 2005), often implying blame may be ascribed to an individual. Generally, however, HRA does not view human error as the product of individual shortcomings but rather as the culmination of contextual and situational factors that impinge on human performance. These factors are commonly referred to as performance shaping factors (PSFs), which serve to enhance or degrade human performance relative to a baseline (Gertman and Blackman, 1994).

HRA is often depicted as consisting of three distinct phases (Boring, 2009):

1. *Modeling of the potential contributors to human error.* This phase typically enlists some variety of task analysis to decompose an overall sequence of events into smaller units suitable for analysis. There is, however, no universally agreed standard for the best level of decomposition, which can lead to confounds when attempting to compare analyses from different HRA methods. HRA is often modeled as part of an overall probabilistic risk assessment (PRA) model of human, system, and environmental failures, in which the human contribution is called the human failure event.

2. *Identification of the potential contributors to human error.* At this phase, relevant performance shaping factors are selected. As with task decomposition, there is no standard list of performance shaping factors, and there is considerable variability between HRA methods. Recently, the standard-like *Good Practices for Implementing Human Reliability Analysis* (Kolaczkowski et al., 2005) produced a set of 15 PSFs to consider

at a minimum. Still, current HRA methods range from a single PSF (e.g., available time in the time-reliability methods) up through 50 PSFs.

3. *Quantification of human errors.* At this phase, a human error probability (HEP) is calculated. Each HRA method features a slightly different approach to quantification, including expert estimation, the use of PSF multipliers, Bayesian approaches, and simulations. Quantification determines the likelihood that the particular action modeled in the previous steps will fail. This error likelihood ranges from 1/100,000 (1E-5) for high reliability up to 1.0 for guaranteed failure. Nominal values for skilled actions are around 1/1000 (1E-3) to 1/100 (1E-2). Note that it is possible to have only qualitative HRA, which addresses insights particularly from the error identification phase but does not produce HEPs.

Probabilistic risk assessment (PRA, also known as probabilistic safety analysis or PSA) is an integration of failure modes and effects analysis, fault tree analysis, event tree analysis, and other techniques to assess the likelihood of failure and the consequence of that failure, and to help find ways to reduce risk to acceptable levels. The goal of human reliability analysis (HRA) is to support PRA in identifying and assessing what can go wrong, how likely it is, and what contribution human error makes to the risks associated with complex systems. PRA, in conjunction with HRA, affords analysts the ability to look at sequential as well as parallel pathways that generate risk, including the human contribution to that risk. Insights are gained by applying event frequencies to hardware failure models and reviewing expected frequencies for various hazardous end-states by condition assessments. Because data for human performance during infrequent, high consequence events at nuclear power plants have been thankfully scarce, HRA often has had to employ expert judgment methods to supplement existing data.

A node in the sequence of events represents a specific function, subtask, or activity that can have two different outcomes: success or failure. A node can either represent the function of a technical system or component, or the interaction between an operator and the process. PRA is used to determine whether an event will likely succeed or fail, and to determine the probability of failure in order to calculate the combined probability that a specific outcome or end state will occur. If the node represents the function of an electronic or mechanical component, the failure probability is usually calculated from engineering knowledge alone. If the node represents the interaction between an operator and the process, the calculation must include the probability that the human will fail. HRA is used to calculate this probability, typically the HEP.

The roots of this approach can be found in the Technique for Human Error Rate Prediction (THERP), a method first presented 50 years ago and evolved over a 20-year period to become the *ur*-method of HRA. This paper discusses some of the history of HRA and THERP, and attempts to frame this history in the context of future directions for HRA.

2. ORIGINS OF THERP

At the Sixth Annual Meeting of the Human Factors Society, November 28-30, 1962, a symposium was held on the topic of "Human Error Quantification" (Swain, Altman, and Rook, 1963). This meeting, led by Dr. Alan Swain (see Figure 1) introduced an early version of THERP, the first formal method for HRA. This was followed in 1963 by a Sandia Laboratories monograph outlining basic human error quantification using THERP and, in 1964, by a special journal edition of Human Factors on quantification of human performance. Throughout the 1960s, Swain and his colleagues focused on collecting human performance data for the Sandia Human Error Rate Bank (SHERB; Rigby, 1967), primarily in connection with supporting the reliability of nuclear weapons assembly in the US. In 1969, Swain met with Jens Rasmussen of Risø National Laboratory and discussed the applicability of THERP to nuclear power applications (Swain, 1969). By 1975, in WASH-1400, Swain had articulated the use of THERP for nuclear power applications, and the approach was finalized in the watershed publication of the NUREG/CR-1278 in 1983.

The precepts of HRA have remained largely unchanged: HRA seeks to identify sources of human error (although the term *human error* has since come to be viewed as pejorative) and quantify the likelihood of such errors (Boring, 2009; Sheridan, 2008). Importantly, the 1962 Human Error Quantification symposium established a research agenda for HRA, including the collection of empirical data from human factors studies to inform the ability of HRA to predict human performance.

Early HRA emerged amid weapons reliability efforts in the 1950s (Meister, 1999). Subsequent to the 1962 symposium, human factors forums were the preferred venue for HRA research, and there was a healthy interplay between research in the two closely related fields. The predictive nature of HRA was noted as being somewhat at odds with conventional thinking in human factors (Swain, Altman, and Rook, 1963). Conventional human factors research centered strongly on system design and system improvement to optimize human performance. Because the evaluation of human performance for the purposes of system design collects data for specific system tasks, a comprehensive collection of empirical data was not available to document human performance. Thus, HRA required a fair amount of generalization of those data that were available. This approach challenged assumptions about the external validity of psychological findings, by using available human performance data in ways beyond the context in which they were originally gathered. Still, the approach held merit, and the gradual accumulation of human performance data allowed HRA to make plausible and valid predictions about human errors.

Figure 1. Alan Swain providing a lecture by candlelight during a power outage in 1972

Meister notes that the ability to predict is one of the hallmarks of science, and HRA uniquely filled that role within human factors (1999). HRA effectively set the stage for other predictive endeavors in human factors—namely, human performance modeling, cognitive modeling, and artificial intelligence techniques—that chronologically followed HRA. This lineage is clearly evident in the topical coverage of THERP, which more closely resembled the emerging cognitive movement than the then dominant behaviorist movement in psychology.

HRA has continued to evolve in methods and applications since the first versions of THERP. THERP was applied to the domain of nuclear power plants (Swain, 1969). Subsequent to the Three Mile Island nuclear power plant meltdown in the US in 1979 (Osif, Baratta, and Conkling, 2004), a refined version of THERP (Swain and Guttman, 1983) became a cornerstone of plant safety assurance and risk-informed decision making by the US Nuclear Regulatory Commission. HRA was subsequently adopted widely across the international nuclear industry, where HRA became an important component of PRA/PSA, combining hardware and human reliability into plant system models (Gertman and Blackman, 1994).

Table 1. Enduring features of THERP in subsequent HRA methods

- Error identification
  - Task analysis approach
  - Human failure events
  - Performance shaping factors
- Quantification
  - Nominal HEP and error factor
  - Basic HEP: Adjustment for PSFs
  - Conditional HEP: Adjustment for dependence
- PRA Integration
  - HRA event tree designed for integration into PRA
  - Recovery
  - Pre- and post-initiating events

A recent review (Bell and Holroyd, 2009) identifies 39 distinct HRA methods. Additional HRA methods are identified in other reviews (e.g., Adhikari et al., 2008; Forester et al., 2006; Lois et al., 2008). Importantly, much of HRA in practice in in newer methods is derived from a common set of HRA features introduced in THERP (See Table 1). HRA methods continue to become more refined and more usable, and they address increasingly nuanced aspects of human performance. While the majority of HRA methods continue to support the nuclear industry, new methods are being applied to other safety-critical domains such as the automotive sector (Sträter, 2000), the oil industry (Aven, Sklet, and Vinnem, 2006), and air-traffic control (Kirwan and Gibson, 2007), all building on the foundations of THERP.

3. CLASSIFYING DIFFERENT HRA METHODS

3.1 Different Funding Sources

Figure 2 depicts a simplified history of some major HRA methods according to sponsorship of the development of the methods. This depiction suggests three parallel developments—one group of methods emerging through the US NRC, a parallel group emerging through sponsorship through the EPRI, and yet another group emerging through non-US sponsors. THERP produced direct, simplified descendants in the form of the Accident Sequence Evaluation Program (ASEP), developed by Swain himself (1987), and later by the Standardized Plant Analysis Risk-Human Reliability Analysis (SPAR-H) method developed by Gertman et al. (2005). In parallel, variations of the Success Likelihood Index Method (SLIM) and its failure-centric FLIM counterpart, were developed (Embrey et al., 1984). More recently, A Technique for Human Error Analysis (ATHEANA) has been developed within the US NRC (2000) to address perceived shortcomings of THERP and its descendants. EPRI developed a series of HRA methods to address industry needs, including the Human Cognitive Reliability/Operator Reliability Experiments (1992) and the Cause Based Decision Trees (1992), as well as the standardized framework for incorporating HRA into probabilistic risk assessment, entitled Systematic Human Action Reliability Procedure (SHARP; 1992). Most of the EPRI and US NRC methods have been implemented in software as the EPRI HRA Calculator. Subsequent to THERP and in parallel to other method developments, various international entities also sponsored HRA development efforts.

Figure 2. Parallel HRA method developments

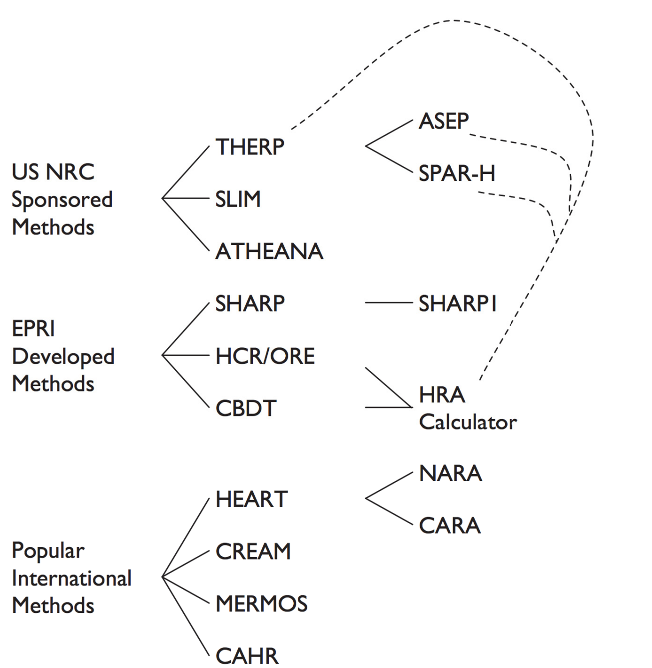

ACRONYM LIST

ATHEANA: A Technique for Human Error Analysis

ASEP: Accident Sequence Evaluation Program

CAHR: Connectionist Assessment of Human Reliability

CARA: Controller Action Reliability Assessment

CBDT: Cause Based Decision Tree

CREAM: Cognitive Reliability Error Analysis Method

HCR/ORE: Human Cognitive Reliability/Operator Reliability Experiments

HEART: Human Error Assessment and Reduction Technique

MERMOS: Method d'Evaluation de la Realisation des Missions Operateur pour la Surete

NARA: Nuclear Action Reliability Assessment

SHARP: Systematic Human Action Reliability Procedure

SLIM: Success Likelihood Index Method

SPAR-H: Standardized Plant Analysis Risk-Human

THERP: Technique for Human Error Rate Prediction

3.2 Generational Approaches

For a number of years, there has existed a distinction between first and second generation HRA methods (see Figure 3). The guidance for classifying a particular method as first or second generation has not been entirely consistent. For example, Hollnagel's Cognitive Reliability Error Analysis Method (CREAM) HRA method (1998) makes a strong argument for considering the HRA methods' use of cognitive factors. Hollnagel argues that the so-called first generation HRA methods did not consider cognition among their PSFs. More modern methods—the so-called second generation HRA methods—explicitly consider and model cognitive PSFs. The delineation fits nicely with the ascent of the cognitive psychological movement. First generation HRA methods coincided with pre-cognitive movements in psychology; second generation HRA methods harnessed findings and insights from the then nascent cognitive movement.

Figure 3. Generations of HRA

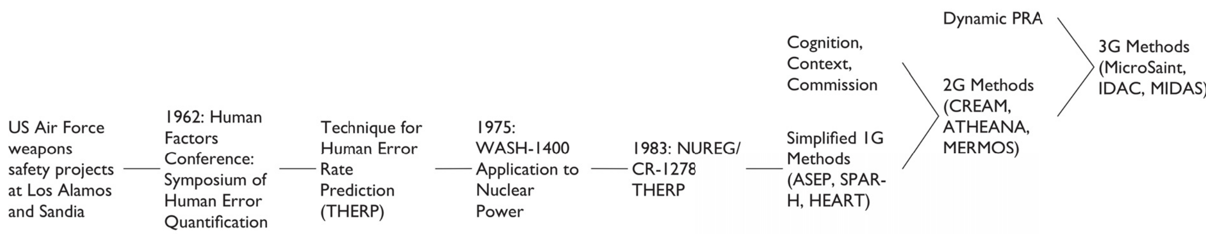

In contrast to the cognitive focus of CREAM, ATHEANA (US NRC, 2000) has in practice developed a differentiation between first and second generation HRA methods on other lines. In ATHEANA, context becomes the key to demarcation between first and second generation HRA. Earlier, first generation methods largely failed to consider the context in which humans made errors, while later, second generation methods carefully consider and model the influences of context on the error.

Other distinctions have been drawn based on the consideration of errors of commission in second generation methods, as opposed to a heavy focus on errors of omission in first generation methods. ATHEANA, for example, has a strong emphasis on and explicit guidance for identifying potential errors of commission.

More generally, the HRA community has been inclined to refer to the HRA generational gap simply in terms of chronology. The oldest, first developed HRA methods are colloquially considered first generation methods, while subsequent methods—the descendants of the earlier methods—are considered second generation methods. Not so coincidentally, these latter or second generation methods tend to have a broader coverage than earlier methods, sometimes at the cost of simplicity. The de facto defining characteristics of second generation methods are the methods' relative novelty (at least chronologically speaking) and their comprehensiveness.

These tidy distinctions of the four classificatory Cs—cognition, context, commission, and chronology—are blurred when one considers an HRA method like SPAR-H (Gertman et al., 2005). SPAR-H was developed as a simplified quantification method built upon THERP, an unambiguously first generation method. SPAR-H augments THERP with an information processing framework, a theoretical model akin to cognitive psychology. Using the CREAM definition of first and second generation, one would clearly consider SPAR-H a second generation HRA method due to its consideration of cognition. However, if one considers context as a defining characteristic of second generation methods, SPAR-H falls short and might be considered a first generation method or even a hybrid (1.5th generation) method. SPAR-H as a method is largely indifferent to errors of omission and commission, suggesting it might be more a first generation method. Yet, SPAR-H is newer and represents at least an iterative modification to its first generation ancestor. So, chronologically, it doesn't seem quite right to call SPAR-H a first generation method, as one might if context or errors of commission are the deciding factors. Clearly there is room for debate, which may not always prove an entirely illuminating endeavor in terms of determining the suitability or quality of a particular HRA method. First generation methods like THERP are still widely and successfully employed, while some second generation methods have remained underutilized.

It should be noted that THERP does not exclude cognition, context, or commission in its discussion of human reliability. NUREG/CR-1278 contains discussions of all these areas, which later served as springboards for newer methods. Far from being causes célèbre for the shortcomings of THERP, these topics simply proved to be introductions in THERP that were later expanded into new methods. First generation HRA methods do and will continue to play a role in classifying and quantifying human performance. First generation methods and notably THERP should continue to be implemented wherever needed; second generation methods should continue to be researched and improved to ensure an efficient, accurate, and complete capture of human performance.

There exist developments—namely in human performance simulation—that do not fit the classification of first or second generation HRA methods. Human performance simulation utilizes virtual scenarios, virtual environments, and virtual humans to mimic the performance of humans in actual scenarios and environments. What sets this form of HRA apart is that it provides a dynamic basis for HRA modeling and quantification. First and second generation methods, by any definition, have featured largely static task analyses of operating events as the underlying basis of performance modeling. These methods have also relied on performance estimations mapped to similar previous performance derived through empirical data or expert opinion. Simulation-based HRA differs from its antecedents in that it is a dynamic modeling system that reproduces human decisions and actions as the basis for its performance estimation. Simulation-based HRA may be called third generation HRA by looking at those features and limitations that are unique to it.

3.3 The Divergence of Human Factors and HRA

Importantly to the human factors community, HRA has become closely linked with reliability engineering. The result of this alignment is that many of the advances in HRA are disseminated outside the mainstream human factors publication channels. A major paper in HRA is more likely to appear in the journal *Reliability Engineering and System Safety* than in *Human Factors*. Late-breaking research more frequently appears at American Nuclear Society conferences than at the annual meetings of the Human Factors and Ergonomics Society.

HRA has become seamlessly integrated in the safety engineering domains it supports. But, HRA as a field must be careful not to distance itself from its human factors roots. The quality of HRA depends on human performance insights, and these insights are most readily built on a human factors foundation. Moreover, recent work using HRA for human performance modeling (Laux and Plott, 2007) and HRA for design (Stanton and Baber, 2002) suggests that the benefit is reciprocal. Human factors stands much to gain still from HRA, just as HRA depends on human factors for improved insights on human performance.

Historically, the study of human error and performance was one of the hallmarks of human factors (Chapanis, 1959). Some (e.g., Roscoe, 1997; Meister, 1999) have credited the birth of human factors to the early efforts by research psychologists to improve human performance, specifically to develop better landing controls on fighter airplanes during World War II. Originally tasked with providing accelerated training for fighter pilots, research psychologists encountered barriers to their pedagogical efficacy due to the limitations of the engineered systems for which the pilots were trained. Despite good training, pilots consistently committed errors in their use of flight controls in the cockpit, resulting in damage to the aircraft and injury to the pilots. In time, research psychologists began offering recommendations for improving the human-machine interface, ultimately decreasing training time and improving human performance on the engineered systems by actually improving the underlying system. Over the years, additional psychological research has continued to be funded to develop and maintain standards for the design of military equipment and training, much of which is widely used in the design of human-machine interfaces in other industrial and commercial applications. The need to improve human reliability and decrease human error rates led to the expansion in psychological research to support human factors. Thus, from the very beginning, HRA and human factors were interrelated, both with a shared basis in psychology and system design.

Although analysts have conducted assessments of human reliability as part of system evaluations since the 1960s (Swain, 1963), its formal foundations came in the WASH-1400 (US Nuclear Regulatory Commission, 1975) study, which purpose was to address the safety of nuclear power plants. This method was further developed with the release of the THERP method (Swain and Guttman, 1983), in which a systematic method for identifying, modelling, and quantifying human errors was documented. THERP and subsequent HRA methods emerged against the backdrop of the Three Mile Island incident in the US, with a corresponding call for risk-informed decision making (Kadak and Matsuo, 2007), specifically through PRA/PSA and HRA. This application firstmost required assessment of existing systems, with less of an emphasis on design than had typically been the case in human factors.

HRA has emerged in two primary roles, that of retrospective and prospective analysis. Retrospective HRA focuses on assessing the risk of something that has already happened, such as an incident or accident. The purpose is to determine the likelihood that something could have happened the way it actually did—was it an anomalous activity, or would it be expected that such an activity could occur again given the same situation? Prospective HRA attempts to assess the risk of something that hasn't actually happened, such as determining the characteristics of an extremely rare event like human performance during a seismic event. Note that while prospective HRA holds tremendous opportunity to anticipate breakdowns in the human-system interface, prospective applications of HRA have not commonly centered on incorporation of such information into the early-stage design of a system. However, as noted in Hirschberg (2004), HRA and PRA/PSA are actively used for improvement of existing processes and systems. HRA and PRA/PSA pinpoint weaknesses and allow prioritization of fixes. Thus, HRA is typically used not in the initial system design phase but rather in the assessment and improvement of existing technologies. It should be noted though, that recent regulatory design guidance such as the *Human Factors Engineering Program Review Model* (NUREG-0711; O'Hara et al., 2004) or the *Human-Rating Requirements* (NASA, 2005) suggest HRA as part of the design process to complement existing human factors design best practices (Boring, 2007). The so-called second-generation HRA methods like ATHEANA (US NRC, 2000) hold considerable promise for redressing the lack of HRA in design. However, in practice, HRA has not to date been extensively used in early design.

The historical differences in the origins between human factors and HRA have resulted in different research approaches. Human factors has typically maintained a strong basis in empirical observation, a reflection of the need to collect data to generate design recommendations. Conversely, HRA has typically relied heavily on process expertise to identify problems and HRA methods themselves or expert estimation for quantification. As a consequence, HRA has tended not to be tied closely to the collection of empirical data. Moreover, human factors has tended not to focus on the probabilistic prediction of human performance, relying instead on the findings from specific empirical observations. While human factors and HRA clearly represent different research approaches in practice, they are part of the same overall process of studying the interface between humans and technology.

As Table 2 depicts, human factors typically exists at the intersection of design and empirical data collection, while HRA intersects risk assessment and probabilistic prediction. It may be argued that human factors and HRA should be seen as part of the same overarching approach, focused in slightly different directions by different requirements. Importantly, there is nothing to preclude a merger, or at least a convergence, of these two research approaches in future applications, although some differences in, e.g., the design of experiments may be needed. Empirical data collection in the style of human factors can inform risk assessment, and probabilistic prediction in the style of HRA can inform design.

Table 2. The intersection of humans and technology in human factors and human reliability analysis

Human \ Technology		Technological Application: Design	Technological Application: Risk Assessment
Human Performance	Empirical Data Collection	HF →	
Human Performance	Probabilistic Prediction		← HRA

From the earliest attempts at HRA, including THERP, there have been attempts to collect data to inform quantification in HRA (Rigby, 1967). These efforts have continued to the present time, with efforts like the US Nuclear Regulatory Commission's Human Event Repository and Analysis (HERA) system (Hallbert et al., 2007), a human performance database developed for NASA (Boring et al, 2006), or the UK's CORE Database (Kirwan, Basra, and Taylor-Adams, 1997). However, as the primary author of THERP, Alan Swain, has noted in reference to the development of quantification in THERP (Personal Communication):

> There should have been many more changes had the research been done to develop more hard data on HEPs for human tasks. That failure has been a disappointment to me. ... I always said that the data tables in [THERP] were not written in stone, and I was not Moses coming down from a mountain with these tables so inscribed.

Despite the desire for more empirical data, HRA researchers have not been entirely successful in leveraging performance data from human factors. As the now retired lead of human factors research at the US Nuclear Regulatory Commission, Julius Persensky, notes (Boring et al., 2009):

> HRA has not provided a clear prescription for human factors on what data are needed. What values are needed under what conditions? Moreover, HRA deals with extremely small probabilities of human actions, which can be extremely difficult to observe in human factors studies. To observe such low frequency human errors, it is necessary to run many more human participants in studies or to run them longer than we have. Clear guidance on HRA data gathering requirements would facilitate research that is more attuned to helping HRA methods develop empirically driven HEPs.

These difficulties notwithstanding, the key to using human factors to inform HRA resides in understanding the data that is produced and that is required from the two domains.

One of the key distinctions between human factors and HRA concerns the product of these respective domains. By product is meant the output—the type of data generated by human factors and HRA, respectively. As noted, much of human factors centers around the collection of empirical data, typically from carefully controlled laboratory studies but also from ethnographic or naturalistic observation. The data produced by such research are invariably tied to a particular research question, such as the usability of a particular computer interface. The data may be quantitative (e.g., descriptive performance measures or inferential statistics when comparing performance on different interfaces) or qualitative (e.g., narrative accounts or specific performance insights), according to the specific question or product under investigation. HRA also produces quantitative and qualitative data, although these are always tied to specific modeled events in an accident sequence. Quantitative data in HRA take the form of HEPs, while the qualitative data typically consist of PSFs and other performance insights.

It has been noted that there is a disconnect particularly in the quantitative data produced within human factors and HRA (Tran et al., 2007). While quantitative human factors data focus on variables such as workload or situation awareness, these measures often do not translate directly into human error data, making them of limited applicability to informing HRA. Similarly, quantitative HRA data, which focus entirely on the HEP or residual uncertainty, do not always inform the performance measures desired by human factors. For example, knowing the error likelihood does not actually tell the human factors researcher or practitioner the expected performance or the level of performance degradation that may precede an actual error.

One may argue that the data of human factors are primarily performance driven while the data of HRA are primarily safety driven. This point holds true to the extent that human factors has been applied broadly outside safety-critical systems (e.g., the usability of consumer products) and HRA has been used primarily in risk analysis applications. These traditional applications need and should not be seen as prescriptive for future applications. Human factors is widely employed for safety-critical applications, and HRA has found its way into applications such as consumer usability. Human factors, which is not always accepted as a reasonable cost of the consumer product design cycle, is nevertheless readily cost justified for safety-critical systems, where human-in-the-loop studies ensure that the system interface may be used effectively without costly unintended and harmful consequences to the operator or user of the system.

4. CONCLUSION

This paper has briefly outlined the emergence of the THERP approach, beginning more than 50 years ago. All subsequent HRA methods are derived as a refinement of THERP or as an attempt to address perceived shortcomings with the original technique. Since the advent of THERP, there have been dozens of new HRA methods. These new methods have included refinements to THERP in a number of key areas: inclusion of a wider spectrum of human performance influences—especially cognitive factors, inclusion of greater consideration of situational context, and better modeling of errors of commission (instead of THERP's heavy emphasis on errors of omission such as skipping a step in a process). It is appropriate on the occasion of THERP's golden anniversary to reflect on its legacy and its overall contribution to the fields of HRA and human factors. THERP remains a highly used and relevant method. As THERP approaches its second half century, it is important to consider what ways HRA methods might be improved and, indeed, what enhancements to THERP will continue to ensure it is a relevant method for contemporary applications.

Disclaimer

This work of authorship was prepared as an account of work sponsored in part by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately-owned rights. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a nonexclusive, paid-up, irrevocable, world- wide license to publish or reproduce the published form of this manuscript, or allow others to do so, for United States Government purposes. INL is a multi-program laboratory operated by Battelle Energy Alliance LLC, for the United States Department of Energy under Contract DE-AC07-05ID14517.

References

Adhikari, S., Bayley, C., Bedford, T., Busby, J., Cliffe, A., Devgun, G., Eid, M., French, S., Keshvala, R., Pollard, S., Soane, E., Tracy, D., & Wu, S. (2008). *Human reliabilty analysis: A review and critique. Final Report for EP/E017800/1.* UK Engineering and Physical Sciences Research Council.

Aven, T., Sklet, S., & Vinnem, J.E. (2006). Barrier and operational risk analysis of hydrocarbon releases (BORA-Release), Part I. Method description. *Journal of Hazardous Materials, A137*, 681-691.

Bell, J., & Holroyd, J. (2009). *Review of human reliability assessment methods. RR679.* Buxton, UK: Health and Safety Executive.

Boring, R.L. (2007). Dynamic human reliability analysis: Benefits and challenges of simulating human performance. In T. Aven & J.E. Vinnem (Eds.), *Risk, Reliability and Societal Safety, Volume 2: Thematic Topics. Proceedings of the European Safety and Reliability Conference (ESREL 2007)* (pp. 1043-1049). London: Taylor & Francis.

Boring, R.L. (2009). Human reliability analysis in cognitive engineering. *Frontiers of Engineering: Reports on Leading-Edge Engineering from the 2008 Symposium* (pp. 103-110). Washington, DC: National Academy of Engineering.

Boring, R.L., Gertman, D.I., Whaley, A.M., Joe, J.C., Tran, T.Q., Keller, M.D., Thompson, R.P., & Gerszewski, I.R. (2006). A taxonomy and database for capturing human reliability and human performance data. *Proceedings of the 50th Annual Meeting of the Human Factors and Ergonomics Society,* 2217-2221.

Boring, R.L., Roth, E., Straeter, O., Laumann, K., Blackman, H.S., Oxstrand, J., & Persensky, J.J. (2009). Is human reliability relevant to human factors? *Proceedings of the 53rd Annual Meeting of the Human Factors and Ergonomics Society,* 610-614.

Chapanis, A. (1959). *Research Techniques in Human Engineering*. Baltimore: Johns Hopkins Press.

Dekker, S.W.A. (2005). *Ten Questions about Human Error: A New View of Human Factors and System Safety*. Mahwah: Lawrence Erlbaum Associates.

Electrical Power Research Institute. (1992). *An Approach to the Analysis of Operator Actions in Probabilistic Risk Assessment, TR-100259*. Palo Alto: Electrical Power Research Institute.

Electrical Power Research Institute. (1992). *SHARP1—A Revised Systematic Human Action Reliability Procedure, TR-101711*. Palo Alto: Electrical Power Research Institute.

Embrey, D.E., Humphreys, P., Rosa, E.A., Kirwan, B., & Rea, K. (1984). *SLIM-MAUD: An Approach to Assessing Human Error Probabilities Using Unstructured Expert Judgment, Volume I: Overview of SLIM-MAUD,* NUREG/CR-3518. Washington, DC: US Nuclear Regulatory Commission.

Forester, J., Kolaczkowski, A., Lois, E., & Kelly, D. (2006). *Evaluation of human reliability analysis methods against good practices. NUREG-1842.* Washington, DC: US Nuclear Regulatory Commission.
Gertman, D.I., & Blackman, H.S. (1994). *Human reliability & safety analysis data handbook.* New York: Wiley-Interscience.
Gertman, D., Blackman, H., Marble, J., Byers, J., & Smith, C. (2005). *The SPAR-H Human Reliability Analysis Method*, *NUREG/CR-6883.* Washington, DC: US Nuclear Regulatory Commission.
Hallbert, B., Whaley, A., Boring, R., McCabe, P., & Lois, E. (2007). *Human Event Repository and Analysis (HERA): The HERA Coding Manual and Quality Assurance, NUREG/CR-6903, Vol. 2*. Washington, DC: US Nuclear Regulatory Commission.
Hirschberg, S. (2004). *CSNI Technical Opinion Papers No. 4, Human Reliability Analysis in Probabilistic Safety Assessment for Nuclear Power Plants. OECD NEA No. 5068.*
Hollnagel, E. (1998). *Cognitive Reliability and Error Analysis Method (CREAM).* Oxford: Elsevier.
Kadak, A.C., and Matsuo, T. (2007). The nuclear industry's transition to risk-informed regulation and operation in the United States. *Reliability Engineering and System Safety, 92,* 609-618.
Kirwan, B., Basra, G., and Taylor-Adams, S.E. (1997). CORE-DATA: A computerized human error database for human reliability support. *Proceedings of the 1997 IEEE Sixth Conference on Human Factors in Power Plants,* 9-7 – 9-12.
Kirwan, B., & Gibson, H. (2007). CARA: A human reliability assessment tool for air traffic safety management—Technical basis and preliminary architecture. In F. Redmill and T. Anderson (Eds.), *The Safety of Systems: Proceedings of the Fifteenth Safety-Critical Systems Symposium* (pp. 197-214). London: Springer Verlag.
Kolaczkowski, A., Forester, J., Lois, E., and Cooper, S. (2005). *Good Practices for Implementing Human Reliability Analysis (HRA), Final Report, NUREG-1792.* Washington, DC: US Nuclear Regulatory Commission.
Laux, L., & Plott, C. (2007). Using operator workload data to inform human reliability analyses. *Official Proceedings of the Joint 8th IEEE Conference on Human Factors and Power Plants and the 13th Annual Workshop on Human Performance/Root Cause/Trending/Operating Experience/ Self Assessment*, 309-313.
Lois, E., Dang, V.N., Forest, J., Broberg, H., Massaiu, S., Hildebrandt, M., Braarud, P.Ø., Parry, G., Julius, J., Boring, R., Männistö, I., & Bye, A. (2008). *International HRA empirical study—Pilot phase report, HWR-844*. Halden, Norway: OECD Halden Reactor Project.
Meister, D. (1985). *Behavioral Analysis and Measurement Methods.* New York: Wiley-Interscience.
Meister, D. (1999). *The history of human factors and ergonomics*. Mahwah, NJ: Lawrence Erlbaum Associates.
National Aeronautical and Space Administration (NASA). (2005). *Human-Rating Requirements for Space Systems, NPR 8705.2A*. Washington, DC: NASA Office of Safety and Mission Assurance.
O'Hara, J.M., Higgins, J.C., Persensky, J.J., Lewis, P.M., and Bongarra, J.P. (2004). *Human Factors Engineering Program Review Model, NUREG-0711, Rev. 2*. Washington, DC: US Nuclear Regulatory Commission.
Osif, B.A., Baratta, A.J., & Conkling, T.W. (2004). *TMI 25 years later: The Three Mile Island nuclear power plant accident and its impact*. University Park, PA: Pennsylvania State University Press.
Rigby, L.V. (1967). The Sandia human error rate bank (SHERB). *Man-Machine Effectiveness Analysis: A Symposium of the Human Factors Society,* 5-1 – 5-13.
Roscoe, S.N. (1997). *The Adolescence of Aviation Psychology.* Santa Monica: Human Factors and Ergonomics Society.
Stanton, N.A., and Baber, C. (2002). Error by design: methods for predicting device usability. *Design Studies,* 23**,** 364-384.
Swain, A.D. (1969). *Human reliability assessment in nuclear reactor plants*. *SC-R-69-1236.* Albuquerque: Sandia Laboratories.
Swain, A.D. (7 October 2008). Personal communication with author.
Swain, A.D., Altman, J.W., & Rook, L.W., Jr. (1963). *Human error quantification. A symposium. SCR-610.* Albuquerque: Sandia Corporation.
Swain, A.D., & Guttman, H.E. (1983). *Handbook of human reliability analysis with emphasis on nuclear power plant applications. Final report. NUREG/CR-1278.* Washington, DC: US Nuclear Regulatory Commission.
Sheridan, T.B. (2008). Risk, human error, and system resilience: Fundamental ideas. *Human Factors*, *50*, 418-426.
Tran, T.Q., Boring, R.L., Joe, J.C., & Griffith, C.D. (2007). Extracting and converting quantitative data into human error probabilities. *Official Proceedings of the Joint 8th IEEE Conference on Human Factors and Power Plants and the 13th Annual Workshop on Human Performance/Root Cause/Trending/Operating Experience/Self Assessment*, 164-169.
U.S. Nuclear Regulatory Commission. (1975). *Reactor Safety Study: An Assessment of Accident Risks in US Commercial Nuclear Power Plants, WASH-1400, NUREG-75/014.* Washington, DC: US Nuclear Regulatory Commission.
U.S. Nuclear Regulatory Commission. (2000). *Technical Basis and Implementation Guidelines for A Technique for Human Event Analysis (ATHEANA), NUREG-1624, Rev. 1.* Washington, DC: US Nuclear Regulatory Commission.