Special Issue

Levels of Automation in Human Factors Models for Automation Design: Why We Might Consider Throwing the Baby Out With the Bathwater

Greg A. Jamieson, University of Toronto, and **Gyrd Skraaning, Jr.**, Institutt for energiteknikk

This paper responds to Kaber's reflections on the empirical grounding and design utility of the levels-of-automation (LOA) framework. We discuss the suitability of the existing human performance data for supporting design decisions in complex work environments. We question why human factors design guidance seems wedded to a model of questionable predictive value. We challenge the belief that LOA frameworks offer useful input to the design and operation of highly automated systems. Finally, we seek to expand the design space for human–automation interaction beyond the familiar human factors constructs. Taken together, our positions paint LOA frameworks as abstractions suffering a crisis of confidence that Kaber's remedies cannot restore.

Keywords: level of automation, design methods, function allocation, automation, human–automation interaction

INTRODUCTION

Professor David B. Kaber's (2018 [this issue]) position paper invites a welcome exchange of ideas about a central construct in human factors engineering: levels of automation (LOA). We applaud the *Journal of Cognitive Engineering and Decision Making* for airing a range of responses to his views in the same issue. This approach offers a promising contrast to slow-motion, ping-pong exchanges between journals with competing perspectives on human factors science.

Address correspondence to Greg A. Jamieson, University of Toronto, 5 King's College Road, Toronto, ON M5S 3G8, Canada, jamieson@mie.utoronto.ca.

Journal of Cognitive Engineering and Decision Making
2018, Volume 12, Number 1, March 2018, pp. 42–49
DOI: 10.1177/1555343417732856

Differing Perspectives on Common Ground

The authors collaborate on human–automation interaction (HAI) research as both empiricists (Burns et al., 2008; Lau, Jamieson & Skraaning, 2014, 2016a, 2016b; Lau, Jamieson, Skraaning, & Burns, 2008; Lau, Veland, et al., 2008) and designers (Hurlen, Skraaning, Myers, Jamieson, & Carlson, 2015; Jamieson, Hurlen, & Skraaning, 2014; Skraaning, Hurlen, LeDarz, & Jamieson, 2016). We adopt an inductive approach to research, building knowledge through prototyping and experimentation in complex work environments. We seek to support designers by aligning that knowledge with the richness of the design problem and not primarily through models of questionable predictive value. This perspective has evolved through 20 years of realistic simulator studies of HAI on complex process-control tasks, where human factors models have shown little utility in predicting performance outcomes (Skraaning & Jamieson, 2017). It is from this perspective that we respond to Kaber's comments on the empirical basis for, and design relevance of, LOA models.

We concur with much of Kaber's critical consideration of the empirical evidence for LOA predictions. We agree that human factors research should move away from constrained artificial situations and toward the problems faced by designers of real operating environments. As he proposes, existing models of types and levels of automation based on human information processing may not be descriptive of operator behavior in complex systems, yielding a frustrating discrepancy between the models' predictions and real-life observations. As Kaber points out, HAI outcomes are likely a function

of more than human information processing. We heartily agree to his opening up of the model space to encompass teamwork, trust, and other metaphors.

ON THE NEED FOR, AND UTILITY OF, HUMAN PERFORMANCE DATA IN SUPPORT OF LOA-BASED PREDICTIONS

Fitts (1951) recognized a need for a research program to create the human performance data to effectively compare human and machine execution of tasks. Yet nearly 35 years after the introduction of MABA-MABA and LOA, Price (1985) cited a "general weakness of applicability" (p. 35) of function allocation methods. These methods

> presumed that human performance data would exist from which the performance of humans could be predicted. . . . [They] would depend on the availability of large quantities of quantified data on human performance, data that could be calibrated to the specific conditions of a new design. *Such data do not now exist, and they probably never will*. (Price, 1985, pp. 35–36; emphasis added)

Today, more than 65 years since the genesis of the Fitts List, the theory of LOA, and the recognition of the need for large quantities of human performance data, Price's prediction has proved prescient. LOA models depend on data that we largely do not have. And like him, although we share the appetite for more data, we cannot envision the human factors community mustering the resources needed to supply it.

We interpret Kaber's call for finer-grained LOA models as an effort to more precisely specify the data needed to validate new deductive models for increasingly complex systems. Like Rouse (1988), we have reservations about the practicality of using empirical data to deductively predict task performance for any reasonably complex system. We also share Perrow's (1984) perspective that simple and complex technological systems differ in kind as opposed to degree. Such systems are essentially unpredictable; their behaviors emergent. In such settings, we argue, an inductive research approach is preferred to build knowledge and extract design principles.

ON THE STRENGTH OF EVIDENCE IN SUPPORT OF LOA-BASED PREDICTIONS IN COMPLEX WORK SETTINGS

Kaber expresses concern about the inconsistent response trends in individual LOA experiments. Onnasch, Wickens, Li, and Manzey (2014) applied reasonable selection criteria to a much larger body of research to identify studies for inclusion in their meta-analysis. Those 18 studies provide a proxy for a discussion about the strength of the empirical support for LOA models. However, we argue that the composition of the studies themselves fails to support the application of the results to complex work systems.

Most of the studies included in the metastudy were conducted in laboratory multitask contexts. According to Onnasch et al. (2014, p. 485), only four of the 18 studies included nonstudent participants; one of which employed military personnel not engaged in tasks pertaining to their expertise (i.e., Calhoun, Draper, & Ruff 2009). Similarly, Cummings and Mitchell (2007) recruited active-duty military personnel with overlapping subject matter expertise with respect to an anticipated future-world experimental task (M. L. Cummings, personal communication, March 31, 2017). The participants in Metzger and Parasuraman (2005) and Sarter and Schroeder (2001) were en route controllers and commercial aircraft pilots, respectively.

Table 1 extracts from Onnasch et al. (2014, p. 481) the studies that employed nonstudent participants and introduces results from an experiment that we are currently preparing for dissemination (Skraaning & Jamieson, 2017). We have also added a column noting the experimental characteristics of each study, including the participants, the simulation environment, and the experimental tasks or scenarios. Table values are Kendall's rank correlation coefficients, referred to as Kendall's tau, a nonparametric measure of correlation. LOAs in each study were converted to sequential rank values and compared with dependent metavariable indicators distinguished in rank by statistical significance. For Skraaning and Jamieson (2017), we followed the ranking method as described in the metastudy (although we have

TABLE 1: Selected (and Supplemented) Results From Onnasch, Wickens, Li, and Manzey (2014) Pertaining to Findings From Complex Work Environments

Study	Experiment Characteristics	Routine Primary Task Performance	Return-to-Manual Primary Task Performance	Routine Secondary Task Performance	Return-to- Manual Secondary Task Performance	Subjective Workload	SA
Calhoun, Draper, & Ruff (2009)	Military personnel (non-SME); commercial simulator; complex scenarios	−.816		0			0
Cummings & Mitchell (2007)	Active-duty military personnel; laboratory simulator; futuristic scenarios	0					0
Metzger & Parasuraman (2005)	En route controllers; medium-fidelity task simulator; multitask scenarios	0	0	0	0	0	0
Sarter & Schroeder (2001)	Commercial aircraft pilots; full-scope simulator; complex scenarios	1					
Skraaning & Jamieson (2017)	Licensed nuclear operators; full-scope simulator; complex scenarios	0	0			0	1[a]

Note. SA = situation awareness; SME = subject matter expert.
[a]Linear contrast analysis showed a large experimental effect (partial eta squared for four levels of procedure automation, $\eta^2 = .30$).

lingering questions about the validity of Onnasch et al.'s [2014] ranking method to support inferences about LOA effects on performance, workload, or situation awareness).

In stark contrast to the conclusions of Onnasch et al. (2014), Table 1 reveals little empirical evidence for a predictive model of LOA effects on task performance, situation awareness, or workload for complex work settings. Ironically, removing these studies from the Onnasch et al. metastudy would presumably improve the predictive power of LOA characterizations for laboratory tasks that can be executed by students after a few hours of training (cf. Li, Wickens, Sarter, & Sebok, 2014).

Onnasch et al.'s (2014) findings do support the position that human factors engineers might use LOA predictions to make valid, reliable, and useful predictions about automation design for simple work environments. However, their article offers no caution regarding limiting the application of the predictions to complex work settings given the current paucity of data.

ON THE ADHERENCE TO LOA-BASED PREDICTIONS DESPITE THEIR POOR PREDICTIVE POWER

Although Kaber acknowledges the conflicting results from individual LOA experiments, he gives two reasons for standing behind predictive

modeling of LOA as the general research strategy to support automation design. First, he argues that they are "handy." But how handy can models of poor predictive value be? Do we have sufficient confidence in these models to encourage their use in the design of safety-critical applications?

By admitting that LOA models developed and evaluated for seven decades are imprecise and unreliable (as Kaber does), it becomes hard to believe in the practical usefulness of the approach. Engineers will continue to make difficult HAI design decisions without the trustworthy technical basis promised by advocates of LOA models. We sympathize with the desire to offer consistent and universal human response trends that are truly useful to designers, but suspect that the LOA approach has overplayed its role in this regard.

Kaber, on the other hand, calls for a research program to establish fine-grained LOA models under the assumption that more sophisticated classification will lead to descriptive performance predictions and thereby greater utility to designers. We are skeptical of the assumption that more detail in these models will yield more accurate predictions than they have to date.

The second reason that Kaber remains committed to LOA look-up tables is that they are the best human factors researchers have to offer to designers and we should be reluctant to dispose of them without a suitable replacement at hand. In our view, this argument is false and impedes the search for better alternatives. It is an argument in favor of function allocation that has long been refuted by a minority in our community (e.g., Fuld, 2000):

> When a scientific discipline finds itself in a dead end, despite hard and diligent work, the dead end should probably not be attributed to lack of knowledge of facts, but to the use of faulty concepts which do not enable the discipline to order the facts properly. The failure of human factor engineering to advance in the area of allocation of function seems to be such a situation. (Jordan, 1963, pp. 161–162)

Kaber (in press) himself asserts, "If actual performance observations on LOAs cannot be connected to theoretical descriptions, then such descriptions have little import in terms of systems design." We agree with Kaber's appraisal of the quality of the empirical efforts to generate knowledge about HAI. We see no systematic flaw in the research methods employed to create the body of evidence. Rather, like Jordan (1963), we suspect that the poor predictive power of LOA-based predictions of human behavior in the presence of automation is based on faulty concepts. We do not see how greater refinement of these concepts will lead to more predictive models.

From our perspective, Kaber adopts a remarkably high threshold for rejecting the LOA concept. We should hesitate to throw this LOA baby out with the bathwater, he says metaphorically. We concur with Kaber that LOA modeling is in its infancy—a 65-year-old infancy. The LOA concept is the Benjamin Button of human factors research. After 65 years of imprecise and unreliable predictions, what more reason do we need to at least entertain the idea of throwing out this baby that fails to thrive?

ON THE USE AND REJECTION OF LOA TAXONOMIES IN SYSTEM DESIGN AND OPERATION

Kaber points to the SAE's (2014) recent adoption of a taxonomy and definitions of driving automation systems as evidence for the utility of LOA taxonomies. However, other practitioner communities have expressed difficulty in applying the LOA concept. For example, the Federal Aviation Administration's (FAA) report on the operational use of flight path management systems cites a limited utility in hierarchical LOA descriptions of flight deck automation.

> The [working group] found that several operators started with a policy that used explicit definitions of levels of automation described as a simple hierarchy in a rigid and prescribed fashion. After gaining operational experience with training and operational use of these rigid definitions, several airlines concluded that such a description assumed a linear hierarchy that does not exist. The various features of the autoflight system (autopilot, flight

> director, autothrottle/autothrust, [flight management system], etc.), can be, and are, selected independently and in different combinations that do not lend themselves to simple hierarchical description. As a result of this experience, those operators revised their policies to allow the pilot to use the appropriate combination of automation features for the situation, *without rigidly defining them in terms of levels*. (Abbott, McKenney, & Railsback, 2013, p. 55; emphasis added)

This example is noteworthy for two reasons. First, the critique emphasizes how the blend of automation capabilities in complex operating environments overwhelms simplistic hierarchical descriptions. We have observed similar difficulties in the process industries as well, where the integration of individual automatic devices, such as protections, controllers, scripts, routines, and programs, makes assignment of automation configurations to ordinal levels of a hierarchy a speculative endeavor (Skraaning & Jamieson, 2017). Instead of a single hierarchy, human factors researchers might appeal instead to multiple LOA hierarchies defined by stages of automation (Parasuraman, Sheridan, & Wickens, 2000), system functions (autopilot, flight director, etc.), or otherwise. Furthermore, they might adopt the notion of adaptive selection of LOAs depending on the context of dynamic operational circumstances or operator state (Byrne & Parasuraman, 1996). In doing so, we risk gravitating toward a descriptive theory of unique function allocation situations instead of providing general LOA predictions that are useful to designers.

Second, the FAA example highlights the value of *operational experience* with rigid LOA descriptions. In contrast, the SAE standard describes a simplified abstraction agreed to by stakeholders. In Kaber's words, the standard is a social construct. Looking beyond the six levels of automation reveals trade-offs in the taxonomy's negotiated formulation. The taxonomy excludes active safety systems and driver assistance systems, such as automated emergency braking and lane keeping assistance, respectively (SAE, 2014, p. 2). Thus, although the SAE's notion of the dynamic driving task includes longitudinal and lateral motion control, automation systems that intervene in these functions are excluded from the taxonomy. Although this may be a useful simplifying assumption, a skeptic might wonder how effective such an abstraction can be in providing a framework for design specification and regulatory practice.

A second example of rejection of the LOA notion was provided by Clinton D. Chapman, production group chief software architect at Schlumberger. He spoke in a session on automation case studies at the September 2014 Society of Petroleum Engineers workshop "Implementation of Drilling Systems Automation." Dr. Chapman reviewed both the Endsley and Kaber (1999) and Sheridan and Verplank (1978) LOA descriptions and recounted how the Drilling Systems Automation Technical Section had been "reviewing this perspective over the previous 5 to 6 meetings." He then asked if any of the workshop attendees was finding the concept useful. Not a single positive response was offered. The first author [Jamieson] was nonplussed. He rose to ask Dr. Chapman if he had understood correctly that the Technical Section was finding the LOA concept to be of no use in supporting the development and implementation of drilling automation. Chapman's reply was that the LOA concept was considered an abstraction that was not "fit for purpose" for drilling automation design.

It should come as no surprise that practitioners are split on the utility of the LOA concept. Like other human factors notions, LOA-based taxonomies offer an intuitive description and an apparently concise aid to cope with a wicked design problem. However, implementing the taxonomy can lead to the realization that the abstraction does not survive beyond the conceptual design phase. It remains to be seen whether the SAE taxonomy will enjoy greater success in surface transportation than reported in aviation and drilling.

ON THE PROSPECTS FOR NEW THINKING IN HAI DESIGN

We agree with Kaber that HAI research asks some of the right design questions. But the LOA tradition prematurely closes off other questions. For example, it implies that the answer to the

question, "How should automation interact with humans?" should be "By throttling the LOA (i.e., adjusting function allocation) and monitoring task performance, situation awareness, and workload effects," as if doing so were the only alternative. From our perspective, the teamwork metaphor (Christoffersen & Woods, 2002) and Dekker and Woods' (2002) "abracadabra" accusation are responses to this unnecessary narrowing of the design space.

There are many more questions that LOA research does not address—questions that are discovered both in operational settings via automation "surprises" (Sarter & Woods, 1997) and through design experiences (Guerlain, Jamieson, Bullemer, & Blair, 2002). Questions arising from an effort to develop concepts of operation for highly automated nuclear plants (Jamieson et al., 2014) include the following:

- Should we design for human interaction with many individual automated agents working at the component or subsystem level or through a single meta-agent?
- How should the designer allocate dialogue with machine agents among human agents? Should these dialogues be discretely held between one crewmember and the agent or be held in the open? Should dialogue with automated agents be the role of one operator or shared among the crew?

We found no useful guidance on these questions in existing LOA research. Designers cannot wait for empiricists to discover the relevant questions through theory-driven research.

Of greater concern is that Kaber leaps from the assurance that we have the right questions to the assumption that we are answering them in the right way. To the extent that he is willing to consider new approaches to HAI modeling, he assumes that starting over means returning to human information-processing models. Why make this assumption? After trying a single strategy for a long time with questionable success, should researchers fiddle with the same basic recipe? Or should we instead demonstrate some creativity in coming up with new ways of supporting designers of HAI in complex work environments?

At the same time, let us also not succumb to nouveau-folk models of HAI that fixate on a single construct or mechanism. Although constructs such as teamwork, trust, and the like almost certainly serve as factors in predicting HAI; if taken in isolation, they also threaten to close off degrees of design freedom. Engineers working on underspecified design problems must adopt a pragmatic approach. Like Sheridan's (2000) bridge builder, they must start out with subjective intuitions before they can move on to normative/objective tools. Likewise, we human factors researchers should also be pragmatic and not subscribe to narrow models.

CONCLUSION

In our view, the LOA paradigm has lost its momentum and is approaching a crisis. New HAI challenges are emerging along with continual technological development. It is not evident that automation is evolving in such a way that the LOA framework accurately describes the new design problems. We suspect that it is too late to save the concept through refinement or by hunting for stronger evidence of its predictive power. The failure of LOA theory thus far to account for human behavior and experience in operational environments is as likely attributable to faulty concepts as to insufficient facts (Jordan, 1963). If the evidence does not support the predictions of LOA models for safety- and production-critical work, we must be willing to at least consider rejecting the theory. We cannot persist in our unwillingness to be guided by accumulated evidence that forms a clear pattern over time just because we lack a contingency plan. And we cannot dismiss the critics for having nothing better (or even nothing at all) up their magician's sleeves.

ACKNOWLEDGMENT

Professor Kaber shared many of the ideas expressed in his paper as a speaker in a Distinguished Seminar Series at the University of Toronto in December of 2016. He graciously agreed to Jamieson sharing an audio recording of his presentations with Skraaning. We extend our thanks to him for that collegial gesture; our comments are offered in the same spirit.

REFERENCES

Abbott, K., McKenney, D., & Railsback, P. (2013). *Operational use of flight path management systems: Final report of the*

Performance-based operations Aviation Rulemaking Committee/Commercial Aviation Safety Team, Flight Deck Automation Working Group. Washington, DC: Federal Aviation Administration.

Burns, C. M., Skraaning, Jr., G., Jamieson, G. A., Lau, N., Kwok, J., Welch, R., & Andresen, G. (2008). Evaluation of ecological interface design for nuclear process control: Situation awareness effects. *Human Factors*, *50*, 663–679.

Byrne, E. A., & Parasuraman, R. (1996). Psychophysiology and adaptive automation. *Biological Psychology*, *42*, 249–268.

Calhoun, G. L., Draper, M. H., & Ruff, H. A. (2009). *Effect of level of automation on unmanned aerial vehicle routing task*. In *Proceedings of the Human Factors and Ergonomics Society 53rd Annual Meeting* (pp. 197–201). Santa Monica, CA: Human Factors and Ergonomics Society.

Christoffersen, K., & Woods, D. D. (2002). How to make automated systems team players. In E. Salas (Series Ed.), *Advances in human performance and cognitive engineering research* (pp. 1–12). Bingley, UK: Emerald Group.

Cummings, M. L., & Mitchell, P. J. (2007). Operator scheduling strategies in supervisory control of multiple UAVs. *Aerospace Science and Technology*, *11*, 339–348.

Dekker, S. W., & Woods, D. D. (2002). MABA-MABA or abracadabra? Progress on human–automation co-ordination. *Cognition, Technology & Work*, *4*, 240–244.

Endsley, M. R., & Kaber, D. B. (1999). Level of automation effects on performance, situation awareness and workload in a dynamic control task. *Ergonomics*, *42*, 462–492.

Fitts, P. M. (Ed.). (1951). *Human engineering for an effective air navigation and traffic control system*. Washington, DC: National Research Council.

Fuld, R. B. (2000). The fiction of function allocation, revisited. *International Journal of Human–Computer Studies*, *52*, 217–233.

Guerlain, S., Jamieson, G. A., Bullemer, P., & Blair, R. (2002). The MPC elucidator: A case study in the design for human–automation interaction. *IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems and Humans*, *32*, 25–40.

Hurlen, L., Skraaning, G., Myers, W. P., Jamieson, G. A., & Carlson, H. (2015). The plant panel: Feasibility study of an interactive large screen concept for process monitoring and operation. In *Proceedings of the 9th International Topical Meeting on Nuclear Plant Instrumentation, Control & Human–Machine Interface Technologies*. La Grange Park, IL: American Nuclear Society.

Jamieson, G. A., Hurlen, L., & Skraaning, G., Jr. (2014). *Highly automated plants: Perspectives, methods and prototypes* (HWR-1128). Halden, Norway: OECD Halden Reactor Project.

Jordan, N. (1963). Allocation of functions between man and machines in automated systems. *Journal of Applied Psychology*, *47*, 161.

Kaber, D. (2018). Issues in human–automation interaction modeling: Presumptive aspects of frameworks of types and levels of automation. *Journal of Cognitive Engineering and Decision Making, 12*, 7–24.

Lau, N., Jamieson, G. A., & Skraaning, G., Jr. (2014). Inter-rater reliability of query/probe-based techniques for measuring situation awareness. *Ergonomics*, *57*, 959–972.

Lau, N., Jamieson, G. A., & Skraaning, G., Jr. (2016a). Empirical evaluation of the process overview measure for assessing for assessing situation awareness in process plants. *Ergonomics*, *59*, 393–408.

Lau, N., Jamieson, G. A., &Skraaning, G., Jr. (2016b). Situation awareness acquired from monitoring process plants: The process overview concept and measure. *Ergonomics*, *59*, 976–988.

Lau, N., Jamieson, G. A., Skraaning, G., Jr., & Burns, C. M. (2008). Ecological interface design in the nuclear domain: An empirical evaluation of ecological displays for the secondary subsystems of a boiling water reactor plant simulator. *IEEE Transactions on Nuclear Science*, *55*, 3597–3610.

Lau, N., Veland, O., Kwok, J., Jamieson, G. A., Burns, C. M., Braseth, A. O., & Welch, R. (2008). Ecological interface design in the nuclear domain: An application to the secondary subsystems of a boiling water reactor plant simulator. *IEEE Transactions on Nuclear Science*, *55*, 3579–3596.

Li, H., Wickens, C. D., Sarter, N., & Sebok, A. (2014). Stages and levels of automation in support of space teleoperations. *Human Factors*, *56*, 1050–1061.

Metzger, U., & Parasuraman, R. (2005). Automation in future air traffic management: Effects of decision aid reliability on controller performance and mental workload. *Human Factors*, *47*, 35–49.

Onnasch, L., Wickens, C. D., Li, H., & Manzey, D. (2014). Human performance consequences of stages and levels of automation: An integrated meta-analysis. *Human Factors*, *56*, 476–488.

Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2000). A model for types and levels of human interaction with automation. *IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems and Humans*, *30*, 286–297.

Perrow, C. (1984). *Normal accidents: Living with high risk systems*. New York, NY: Basic Books.

Price, H. E. (1985). The allocation of functions in systems. *Human Factors*, *27*, 33–45.

Rouse, W. B. (1988). Adaptive aiding for human/computer control. *Human Factors*, *30*, 431–443.

SAE. (2014). *Taxonomy and definitions for terms related to on-road motor vehicle automated driving systems* (Standard J3016). Warrendale, PA: SAE International.

Sarter, N. B., & Schroeder, B. (2001). Supporting decision making and action selection under time pressure and uncertainty: The case of in-flight icing. *Human Factors*, *43*, 573–583.

Sarter, N. B., & Woods, D. D. (1997). Team play with a powerful and independent agent: Operational experiences and automation surprises on the Airbus A-320. *Human Factors*, *39*, 553–569.

Sheridan, T. B. (2000). Function allocation: Algorithm, alchemy or apostasy? *International Journal of Human–Computer Studies*, *52*, 203–216.

Sheridan, T. B., & Verplank, W. L. (1978). *Human and computer control of undersea teleoperators*. Cambridge, MA: MIT Man–Machine Systems Lab.

Skraaning, G., Jr., Hurlen, L., LeDarz, P., & Jamieson, G. A. (2016). *Feasibility study of an interactive large screen concept for automated plant start-up* (HWR-1179). Halden, Norway: OECD Halden Reactor Project.

Skraaning, G., & Jamieson, G. A. (2017). *Twenty years of HRP research on human–automation interaction: Insights on automation transparency and levels of automation*. OECD Halden Reactor Project, Halden, Norway.

Greg A. Jamieson is a professor and Clarice Chalmers Chair of Engineering Design at the University of Toronto. He received his PhD in mechanical and industrial engineering from the University of Toronto

in 2003. He is a licensed Professional Engineer in Ontario.

Gyrd Skraaning Jr. received his PhD in experimental psychology from the Norwegian University of Science and Technology in 2003. He has worked as a researcher at the Institute for Energy Technology (IFE) and OECD Halden Reactor Project since 1996, focusing on the design and evaluation of control environments in the nuclear domain, air traffic control, and the petroleum industry. His current position at IFE is principal scientist.