Int. J. Human-Computer Studies 71 (2013) 1056–1068

Contents lists available at ScienceDirect

Int. J. Human-Computer Studies

journal homepage: www.elsevier.com/locate/ijhcs

Ecological interface design and sensor noise

Olivier St-Cyr*, Greg A. Jamieson, Kim J. Vicente

Department of Mechanical and Industrial Engineering, University of Toronto, 5 King's College Road, Toronto, Ont., Canada M5S 3G8

ARTICLE INFO

Article history:
Received 31 August 2011
Received in revised form
25 July 2013
Accepted 15 August 2013
Communicated by D. Boehm-Davis
Available online 27 August 2013

Keywords:
Sensor noise
Ecological interface design
Display design
Human–computer interaction
Instrumentation

ABSTRACT

This paper investigates the effects of the presence and magnitude of sensor noise on operators' performance and control stability when they use an Ecological Interface Design (EID) interface and a non-EID interface. Sensor noise was gradually increased in selected low-level physical sensors of DURESS III, a representative thermal-hydraulic process simulation. There are two important findings. First, participants in the EID condition achieved target goals significantly faster across all magnitudes of sensor noise. Second, participants in the EID condition exhibited more stable control; experiencing fewer and shorter oscillations around the target goals. This is the first study to empirically investigate the impact of the presence and magnitude of sensor noise on the robustness and effectiveness of an EID interface. These findings are important if EID is to be applied in industrial settings.

1. Introduction

Modern process facilities are instrumented with thousands of sensors that measure low-level state variables and deliver them to sophisticated control systems. Although there are standard industrial practices for verification and validation of sensors, information transmitted contains noise (Stein, 1969). This noise can affect the quality of the information provided in the display as well as the ways operators and automation control the process (Reising and Sanderson, 2002a). For this reason, graphical interface designers employ (at least) basic checks on the quality of data available to be presented in the interface – and mitigate against the possibility that "bad" or "nonsense" data will disrupt the information communication (K. Christoffersen, personal communication, December 19, 2012).

Ecological Interface Design (EID: Vicente and Rasmussen, 1992; Burns and Hajdukiewicz, 2004; Bennett and Flach, 2011) is a framework that helps designers develop displays that allow workers to cope with novelty and change. EID uses the Abstraction Hierarchy (AH; Rasmussen, 1985) as a modeling tool to identify the work domain constraints in terms of information content and structure. These constraints are then mapped on to interface forms that take into account the capabilities and limitations of the human operator. Over the past two decades, EID has gained currency as a framework for constructing human–computer interfaces for a variety of complex sociotechnical systems (see Vicente, 2002, for a comprehensive review), particularly in process monitoring and control (Burns, 2000a, 2000b; Burns et al., 2008; Drivalou and Marmaras, 2009; Ham and Yoon, 2001a, 2001b; Hajdukiewicz and Vicente, 2002; Ham et al., 2008; Jamieson, 2007; Jamieson et al., 2007; Jamieson and Vicente, 2001; Kim et al., 2012; Lau et al., 2008a, 2008b,2008c; Upton and Doherty, 2008). This paper investigates the potential effects of sensor noise on interfaces designed using EID. The results should contribute to theory and should also be relevant for the applicability of EID in industrial settings (Watanabe, 2001).

1.1. Emergent features in ecological interfaces

Computer interfaces that use emergent features in a graphical representation of a process have several advantages (Pomerantz and Pristach, 1989; Bennett and Flach, 1992, 2011; Bennett et al., 1993; Sanderson et al., 1989). For example, emergent features can represent the constraints of a system in ways that allow operators to easily see whether the constraints are intact or broken. When operators directly perceive emergent features, they can be relieved of the more cognitively demanding task of computing derived quantities (Bennett and Flach, 1992). Moreover, emergent features can provide a visual signal of a departure from normality and in some cases may help the operator diagnose the nature of a failure.

Fig. 1 shows an example of an emergent feature. The property of the configuration of individual variables (in this case mass input and mass output) emerges on the display to signal a significant, task-relevant, integrated variable (slope of the line).

EID interfaces rely on emergent features to convey higher-order relations in the work domain. Implementing these emergent

* Corresponding author. Tel.: +1 416 978 3040; fax: +1 416 978 7753.
E-mail address: stcyr@mie.utoronto.ca (O. St-Cyr).

1071-5819/$ - see front matter © 2013 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.ijhcs.2013.08.005

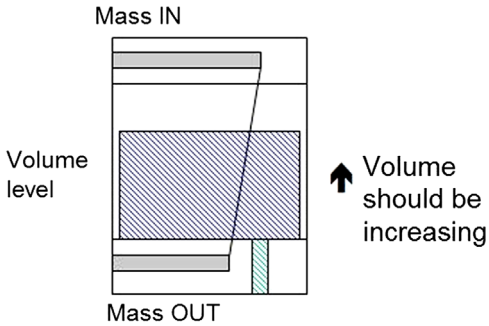

Fig. 1. Emergent feature showing the relationship between Mass Input and Mass Output (adapted from Vicente, 1999).

features often requires derivation of the higher-order relations from lower-level data normally obtained through sensors (cf. Reising and Sanderson, 2002a). Unavailability of sensors, whether through design efficiency or operational failure can therefore compromise the high-level constraints portrayed through the emergent features (Reising and Sanderson, 2002b).

1.2. Sensor noise in ecological interfaces

When EID was introduced, Vicente and Rasmussen (1992) pointed out that noisy sensors are a source of data uncertainty that could compromise the robustness of EID interfaces. Vicente et al. (1996) elaborated on this concern, suggesting that the distortion of emergent features in ecological interfaces could have a detrimental impact on the performance of operators. Vicente (2002) stated that empirically studying the presence and magnitude of sensor noise in EID interfaces was a high priority research item; pointing out that the negative impact of sensor noise on EID could be an obstacle to using the framework in industry. If sensor noise reflected in the emergent features does compromise the robustness of ecological interfaces, previously observed task performance and control stability advantages of EID over non-EID interfaces (see Vicente, 2002) could diminish, disappear, or even reverse.

Alternatively, Vicente (2002) also posited that an EID interface might be more robust to sensor noise on account of the redundant system constraints (Frank, 1990) discovered through the abstraction hierarchy (AH: Rasmussen, 1985) and portrayed in the interface (see also Vicente and Rasmussen, 1992; Jamieson and Vicente, 2005). These redundant system constraints describe the work domain at each level of the AH (Vicente, 1999). An interface that represents system constraints at multiple levels of abstraction may continue to provide a correct account of the constraints at one level of abstraction when those at another level are violated. For example, if settings of components at the Physical Function level do not yield expected system performance, the operator may be able to examine variables or relations at the Generalized or Abstract Function levels to detect the discrepancy and reason about the true state of the work domain and the source of the fault. Thus, an interface based on an AH may support multiple alternative ways of thinking about the control problem. However, if the redundant system constraints are derived from the same low-level sensors, the mechanism underlying this anticipated robustness might fail.

Only one study to date addresses the topic of sensors and EID. Reising and Sanderson (2002b, 2004) studied the impacts of sensor and system failures on diagnostic accuracy using the Pasteurizer II microworld (Reising and Sanderson, 2002c). They noted that sensor failures could produce both topographic failures, where downstream indications are borrowed from an upstream sensor; and derivational failures, where higher-level indications

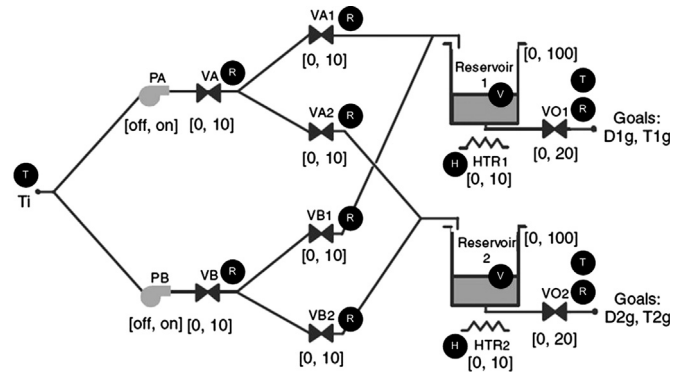

Fig. 2. DURESS III with locations and types of sensors. *T* indicates a temperature sensor, *R* indicates a flow rate sensor, *V* indicates a level sensor, and *H* indicates a heat transfer sensor. (Volume indications (*V*) are, in application, typically derived from level sensors. A heater sensor (*H*) can be thought of as a current sensor for an electrical element heater.)

involve calculations that use a lower-level sensor value. Instrumenting a system with a minimal rather than maximal number of sensors leads to more topographic and derivational failures, with potential consequences for operators' ability to diagnose failures, distinguish system from sensor failures, and control the system effectively.

Reising and Sanderson's (2004) results suggest that the EID framework will support better sensor failure diagnoses only when the interface is based on enough sensors to ensure that derivational information is adequate and available to help operators in their problem-solving strategies (i.e., maximal instrumentation). Although Reising and Sanderson (2002b, 2004) tested sensor noise as one of several ways that sensor failures could potentially compromise the advantages of EID interfaces, they manipulated sensor failures rather than the magnitude of sensor noise. Hence, to date, no experiments have systematically investigated the impact of the presence or magnitude of sensor noise on performance with an EID interface. The research reported in this paper fills part of this gap by exploring the effects of the presence and magnitude of sensor noise on operators' performance and control stability using an EID and a non-EID interface.

2. Analytical method: the DURESS III microworld

2.1. DURESS III

To study the effects of sensor noise on EID, a variant of the DURESS II (DUal REservoir Simulation System; Vicente, 1996; Pawlak and Vicente, 1996) microworld was implemented (see Fig. 2). DURESS III allowed the experimenter to add sensor noise at varying magnitudes to all 15 of its simulated sensors. For the purpose of this paper, sensor noise is what Anyakora and Lees (1974) call *reading erratic*. Erratic readings disturb the measurement signal and may erroneously send an indication that a malfunction is present and so potentially confuse operators. A random number algorithm computed white normally-distributed Gaussian noise and added it to the true sensor readings, forming an accuracy range (e.g., $\pm 2\,^{\circ}$C). The noise covered all frequencies available, given a channel bandwidth of .5 Hz that was limited by the sampling rate (2 s; see next paragraph). The maximum absolute value of the noise (accuracy value) was set to three standard deviations from the mean (3σ). Baseline sensor accuracy values were obtained for each type of sensor by averaging accuracy ranges of commercial sensors from different vendors. Table 1 shows the industrial average accuracy values for the different sensors of DURESS III.

Table 1
Representative industrial average noise values for sensors of DURESS III.

Category	DURESS III sensors	Range	Accuracy
Temperature	*T*0, *T*1, *T*2	0–50 units	±1 °C
Flow rate	FVA, FVB FA1, FA2, FB1, FB2	0–10 units	±2% of range=.2 units
	FO1, FO2	0–20 units	±2% of range=.4 units
Heat transfer	FHTR1, FHTR2	0–10 units	±2% of range=.2 units
Volume	*V*1, *V*2	0–100 units	±1% of range=1 unit

Noise values were generated every 2 s and updated on the interfaces. This was thought to be representative of contemporary process instrumentation and control systems.[1] Moreover, as pointed out by Moray (1986), computer displays should be updated at a higher rate than the dynamics of the underlying process (i.e., the time constants of the system's components). Time constants in DURESS III are on the order of 5–15 s, depending on the components. Hence, updating the display every 2 s is reasonable.

2.2. Sensor-annotated abstraction hierarchy

Constraints and relations underlying DURESS II have been modeled using an AH (Rasmussen, 1985; Bisantz and Vicente, 1994; Vicente, 1999). This model can be adopted for DURESS III because its work domain is identical to that of DURESS II. To analyze the effects of the sensor noise on the quality of information provided to the interfaces, a sensor-annotated AH of DURESS III was created following the notation of Reising and Sanderson (2002b). The sensor-annotated AH for DURESS III is shown in Fig. 3. Table 2 summarizes the sensor data available in the DURESS III simulation. The sensor-annotated AH was developed as follows:

- Means-ends links between AH levels are removed for clarity and replaced by lines connecting certain sensors at adjoining levels.
- Below some AH nodes are small circles representing sensors in the DURESS III system. Sensor nodes include: temperature (*T*), volume (*V*), flow rate (*R*), heat transfer (*H*), mass flow rate (*M*), and energy flow rate (*E*).
- *Black circles* are sensor nodes that are directly sensed as opposed to being derived mathematically from lower-level sensor data. That is, there exists a physical sensor for every sensor node represented by a black circle.
- *Gray circles* indicate variables of DURESS III that have been derived mathematically from lower-level sensor data rather than sensed directly. Equations for the derivations can be found in Vicente (1999).
- *Black solid lines* indicate redundant sensor data portrayed in more than one node of the AH. For example, data from the flow rate output valve sensors (FO1, FO2) are also used to portray mass output flow rates (MO1, MO2) and the output demand goal flow rates (*D*1, *D*2). Likewise, data from the volume sensors (*V*1, *V*2) are also used to portray mass inventories of each reservoir (*M*1, *M*2).
- *Black dashed lines* indicate all the sensors that contribute to a particular mathematical derivation (gray circles) from lower-level sensor data. For example, the mass input one flow rate (MI 1) is mathematically derived from the flow rates FA1 and FB1.

[1] However, the 2 s refresh rate is a departure from the .1 s refresh rate used in prior DURESS studies.

Following the notation of Reising and Sanderson (2002b, 2002c, 2004), the sensor set for DURESS III is *topographically inadequate* in that sensor readings from one location (e.g., FA1 and FB1) are used (although not displayed) at another location (e.g., Reservoir 1 inlet). Because the Mass and Energy Inputs for Reservoir (MI 1 and MI 2, respectively) are derived from valve flow rates, the sensor set can also be categorized as being *derivationally inadequate*. Thus, an interface based on this sensor set would be characterized as a PID.min or EID.min interface according to Reising and Sanderson's classification.

2.3. DURESS III interfaces

Two interfaces (P+S and P+F) were developed for the DURESS III simulation. The P+S (i.e., non-EID) interface (Fig. 4) portrays all 15 physical sensors available in the DURESS III simulation. In contrast to the P (i.e., non-EID) interface designed for previous DURESS studies (see Vicente, 2002), the P+S interface created for the current study includes flow information for the valves and heaters. The original P interface did not include flow information. Both Maddox, (1996) and O'Hara et al. (2000) have commented on the representativeness of the P interface design, asserting that industrial process control systems typically display flow information. We included flow information in our non-EID interface to ensure that both interfaces were based on the same underlying sensor set. Consequently, both interfaces (P+S and P+F) display the same inventory of physical sensors. However, the P+S interface does not display any high-level properties of the system and has no emergent features based on derived values. Therefore physical sensors do not contribute to any mathematical derivations in the P+S interface.

In contrast, the P+F interface (designed under EID principles; Fig. 5) displays both physical and functional information about the work domain by means of emergent features. As in the P+S interface, all 15 physical sensors are portrayed on the P+F interface. Eight additional state variables (MI 1, MI 2, EI 1, *E*1, EO 1, EI 2, *E*2, and EO 2) are mathematically derived from lower-level sensor data (see Table 2 for details). On the right side of the display, each of the two blue-filled funnel displays shows mass input (MI 1&2, top yellow bar) and mass output (MO 1&2, bottom yellow bar). Similarly, each of the orange-filled funnel displays shows energy input (EI 1& 2, top yellow and red bar) and energy output (EO 1&2, bottom red bar). In between the mass and energy graphics an integrated graphic displays the relationship between mass, energy, and temperature. The expected rates of change of key variables (*M*1, *M*2, *E*1, *E*2, *V*1, *V*2) are shown directly (emergent features) as instantaneous sloped lines on the right side of Fig. 5.

2.4. Noise magnitude and pixel area

The effect of noise on the DURESS III interface elements depends on the scale, the number of pixels taken up by each sensor indicator, the sensor accuracy and the noise multiplier (see Table 3). For example, both input flow rate valves and output flow rate valves occupy similar amounts of space on the interface

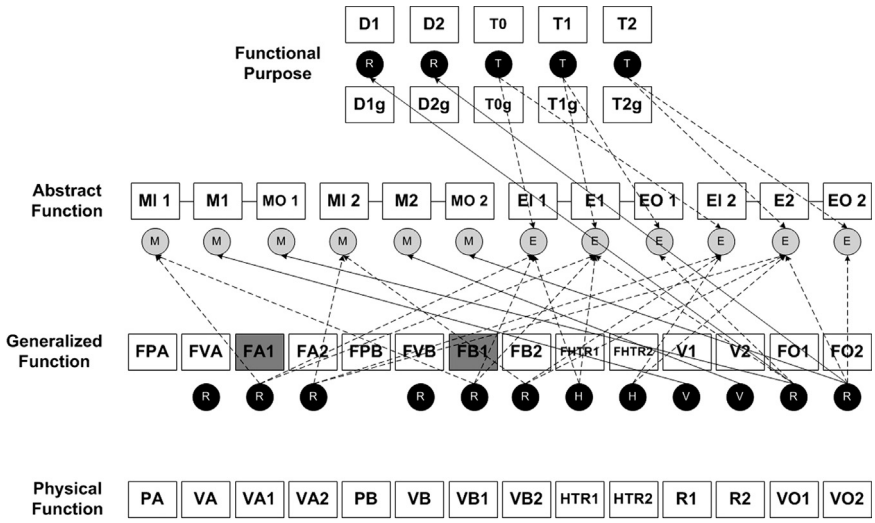

Fig. 3. Sensor-annotated AH of the DURESS III simulation. In the current study, increases in sensor noise magnitude were introduced for sensors FA1 and FB1.

Table 2
Physical sensors available in the DURESS III simulation as well as mathematically derived variables from lower-level sensor data.

Physical Sensors	Mathematically derived variables	
Flow rate valve A (FVA)	**Variables**	**Contributing sensors**
Flow rate valve A1 (FA1)	Output demand goal 1 (*D*1)	Directly from FO1
Flow rate valve A2 (FA2)	Output demand goal 2 (*D*2)	Directly from FO2
Flow rate valve B (FVB)	Mass input 1 (MI 1)	FA1 and FB1
Flow rate valve B1 (FB1)	Mass inventory I (*M*1)	Directly from *V*1
Flow rate valve B2 (FB2)	Mass output 1 (MO 1)	Directly from FO1
Volume reservoir 1 (V1)	Mass input 2 (MI 2)	FA2 and FB2
Volume reservoir 2 (V2)	Mass inventory 2 (M2)	Directly from *V*2
Heat transfer heater 1 (FHTR1)	Mass output 2 (MO 2)	Directly from FO2
Heat transfer heater 2 (FHTR2)	Energy input 1 (EI 1)	FA1, FB1, FHTR1, *T*0
Flow rate output valve 1 (FO1)	Energy inventory I (E1)	FA1, FB1, FHTR1, *T*1, FVO1
Flow rate output valve 2 (FO2)	Energy output 1 (EO 1)	*T*1, FO1
Temperature Thermometer 0 (*T*0)	Energy input 2 (EI 2)	FA2, FB2, FHTR2, *T*0
Temperature Thermometer 1 (*T*1)	Energy inventory 2 (E2)	FA2, FB2, FHTR2, *T*2, FO2
Temperature Thermometer 2 (*T*2)	Energy output 2 (EO 2)	*T*2, FO2

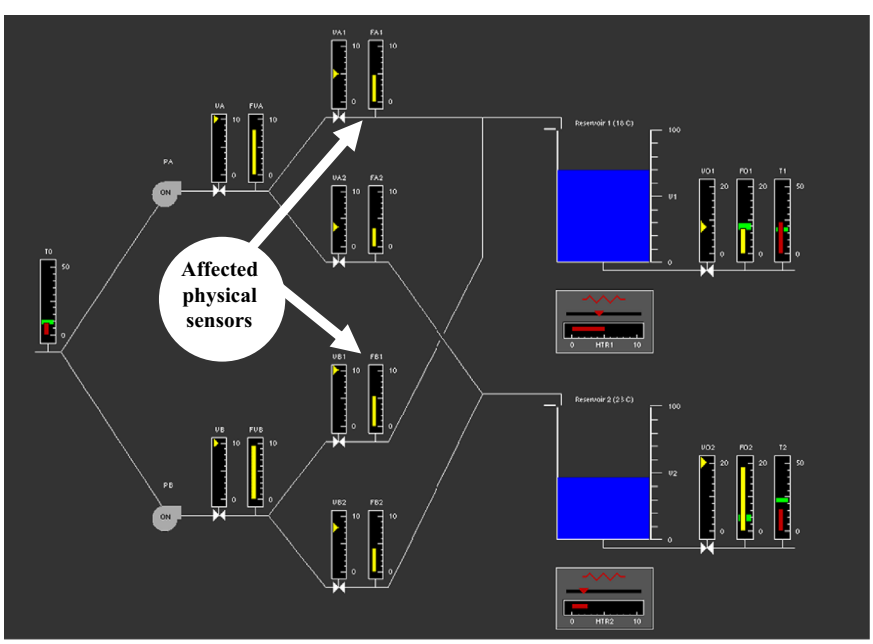

Fig. 4. P+S Interface for DURESS III – displays primarily physical information about the work domain as well as data from all 15 physical sensors available in the DURESS III simulation. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

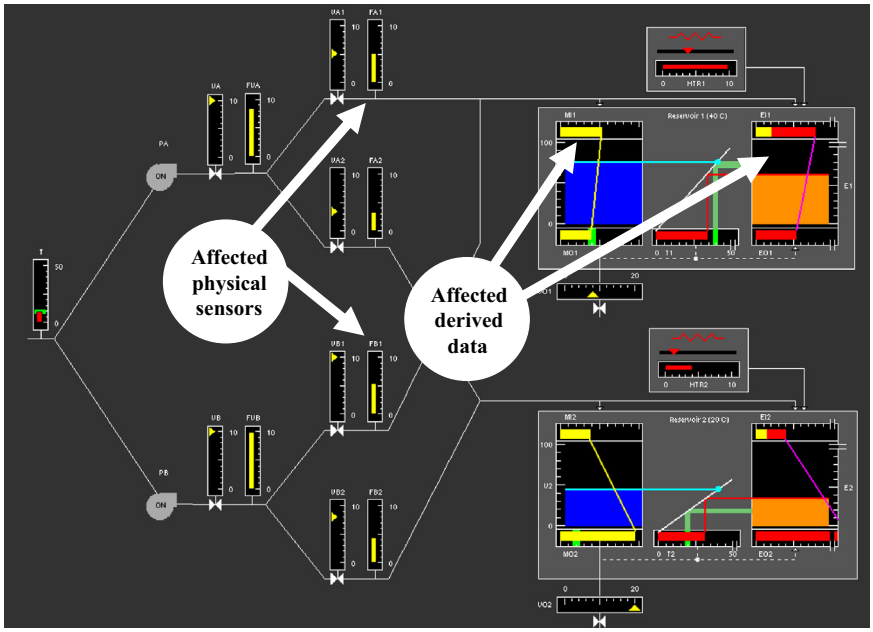

Fig. 5. P+F Interface for DURESS III – displays both physical and functional information about the work domain as well as data from all 15 physical sensors available in the DURESS III simulation. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 3
Relationships between scale and pixels for all sensors of the DURESS III microworld. Calculations are based on industrial average levels of sensor noise.

Sensors	Scale	Scalar pixels	Accuracy (3σ)	Noise	No. of pixels reflecting noise	Noise multiplier			
						1 ×	5 ×	25 ×	50 ×
T0, *T1*, *T2*	0–50	100	± 1 °C	± 1 °C	2 px	2 px	10 px	50 px	100 px
FVA, FVB; **FA1,** FA2;[a] FB1, FB2	0–10	83	± 2%	± .2	1.652 px	2 px	8 px	42 px	83 px
FO1, FO2	0–20	100	± 2%	± .4	2 px	2 px	10 px	50 px	100 px
FHTR1, FHTR2	0–10	95	± 2%	± .2	1.9 px	2 px	10 px	48 px	95 px
V1, *V2*	0–100	103	± 1%	± 1	1.03 px	1 px	5 px	26 px	52 px

[a] In the current study, only sensors FA1 and FB1 were subjected to increases in noise magnitude using noise multiplier. For all other sensors, industry average noise was present (see Sections 3.3 and 3.4 for more details).

(83 pixels for input valves and 100 pixels for output valves). However, they have different scales (0–10 for input valves and 0–20 for output valves). Thus, the calculated number of pixels reflecting noise in each indicator differs. For example, the input flow rate indicator allocates 2 pixels (100 × .2/10 = 2) to standard noise whereas the output flow rate indicator allocates 1.652 pixels (83 × .2/10 = 1.652). With rounding to the nearest pixel, at the industrial average level of noise, the variation in both indicators would equal 2 pixels. At higher levels of noise (e.g., 25 ×), the number of pixels reflecting noise equals 42 for the input flow rate and 50 for the output flow rate. Therefore, for a given sensor and noise magnitude, operators could perceive different degrees of variation. However, with the exception of the reservoir volume (which is not central to the study described in the next section), these differences are relatively small.

3. Empirical method

3.1. Apparatus

The study was conducted using the DURESS III microworld described in the previous section. The simulation ran on SGI IRIS INDIGO R4400 and SGI OCTANE R10000 machines. Participants received feedback about the state of the system through 21" high-resolution color monitors (1024 × 768 resolution; 72.5 Hz refresh rate) and controlled the simulation using a computer mouse.

3.2. Participants

The 20 participants (8 female, 12 male) were mechanical and industrial (10 each) engineering undergraduate students with a mean age of 21 years. They were solicited through local advertisements and selected based on their willingness to participate, their expertise with computer systems, and their cognitive style (see Torenvliet et al., 2000 for more details). All participants had taken at least two courses in physics. Each participant was paid at a maximum rate of $10 per hour ($5 for each hour, $3 for completing the study, and $2 for good performance). The participants were required to sign a consent form and were fully debriefed. All data remained confidential and no participant was identified individually. Participants were allowed to withdraw from the study at any time and were paid for the time they invested up until that point.

3.3. Experimental design

The objective of the study was to assess the effects of the presence and magnitude of sensor noise on performance and control stability using EID (P+F) versus non-EID (P+S) interfaces.

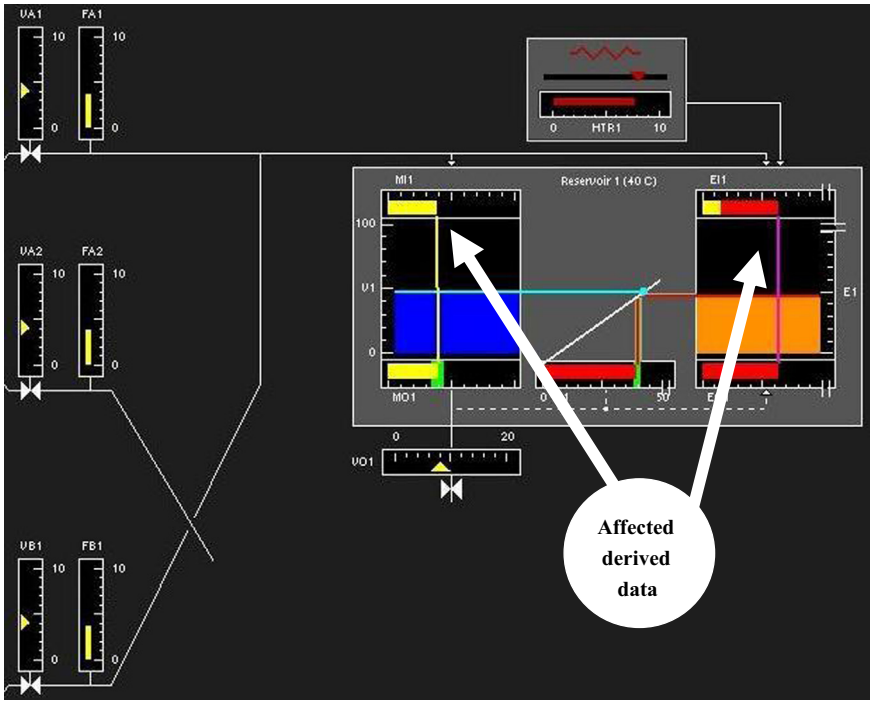

Fig. 6. P+F interface in steady state condition with industry average sensor noise. Under this condition, all sensors exhibit normal oscillations around their true values, with updates to sensor values every 2 s. These oscillations, although visible, are minimal and have minor impacts on the emergent features related to mass and energy balance. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

The experiment followed a mixed design with interface as a between-participants factor and noise magnitude as a within-participants factor. Ten participants each controlled the simulation using the P+S and P+F interfaces. Each participant completed 80 trials grouped into four identical blocks of 20 trials, averaging to 25 one-hour daily sessions. For each trial, participants were presented with a system in its shut-down state and were asked to bring the system to a steady-state condition in which the four systems goals (two output demands – **D1** and **D2** – and two temperature demands – **T1** and **T2**) were within target regions for five consecutive minutes within a 30-minute time limit. The four system goals were determined in advance by the investigator, varied from trial to trial within a block and were the same for all participants.

The trials were divided into two phases: *learning* and *perturbation*. In the learning phase (first 60 trials), participants were given extended practice at operating DURESS III. During this phase, industry average noise (see Table 1) was present on all sensors. During the perturbation phase (trials 61 to 80), the noise magnitude for sensors **FA1** and **FB1** increased every second trial based on scaling multipliers (5, 10, 15, 20, 25, 30, 35, 40, 45, and 50). For instance, a multiplier of 10 increased the magnitude of industry average sensor noise by 10 times.

3.4. Selection of sensors to carry noise

Sensors FA1 and FB1 were chosen to carry noise because they have the greatest impact on the mass and energy balance emergent features, as calculated by the equations outlined in Vicente (1999) and as shown in Fig. 3. Noise in FA1 and FB1 propagates to the Mass Input 1 (MI 1), the Energy Input 1 (EI 1), and the Energy Inventory 1 (E1) derived values. Increasing noise for sensors that bear the *greatest* impact, such as FA1 and FB1, constitutes the worst-case scenario for an EID interface. Note that, if there were flow rate sensors immediately prior to the Reservoirs, DURESS III would be topographically and derivationally adequate in the feedwater streams.

Also, noise was added to both FA1 and FB1 to ensure that participants would not subvert the experimental manipulation. Hajdukiewicz and Vicente (2004) reported that, under some low-demand conditions, some DURESS users adapted to perturbations introduced in valve VA1 by shutting it off and maintaining a constant flow of water to both reservoirs by using valve VB1 instead. Adding noise to sensor FB1 ensured the experimental manipulation presented a more uniform challenge across possible strategies for managing the sensor noise.

3.5. Effect of noise on interfaces

The effect of sensor noise on the flow, temperature, volume and heater output indicators was largely[2] consistent between the P+S and P+F interfaces. Under steady state conditions, as the sensor values updated every two seconds, the indicators would jump up/down or left/right (depending on the orientation of the indicator in question) reflecting the varying noise component in that sensor (see Fig. 6). In keeping with the normally distributed noise function, these jumps would typically be $<1\sigma$ for each sensor. However, jumps of up to 3σ were possible, but increasingly rare occurrences. At high or low values on each scale – and increasingly at higher noise multipliers in the two affected flow sensors – the summed signal and noise could exceed the range of the display. In these situations, the indicator would show maximum or minimum values, respectively.

The noise also impacted the graphical elements for the eight additional state variables and the emergent features in the P+F interface. The indicators for the additional state variables

[2] There are, indeed, minor differences in the graphical forms between these elements of the P+S and P+F interfaces. For example, the output flow rate indicators in P+S are oriented vertically whereas they are oriented horizontally in the P+F interface.

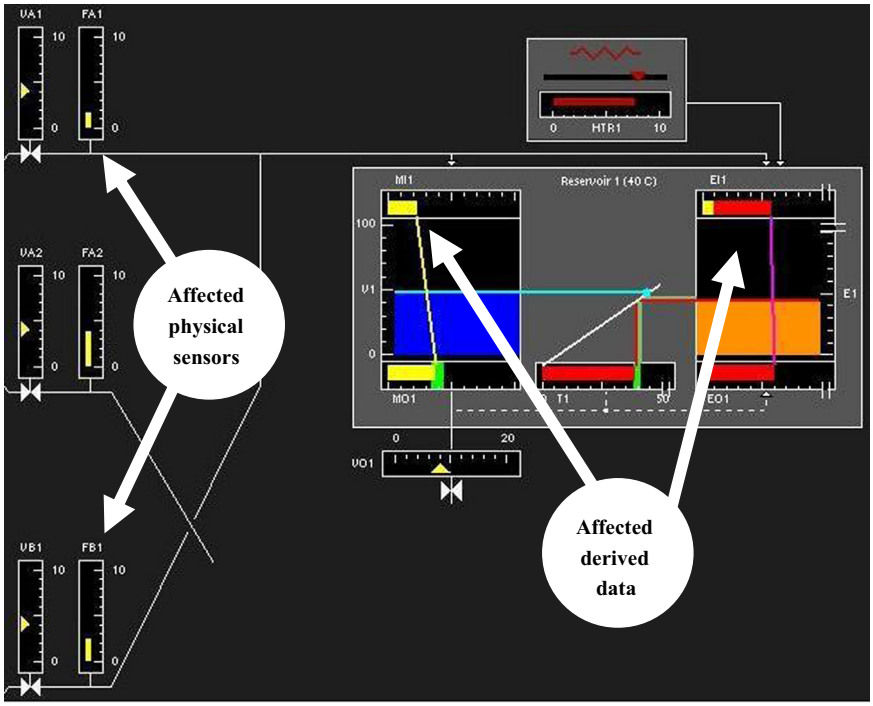

Fig. 7. P+F interface in steady state condition with a noise multiplier of 25 ×. Under this condition, sensors FA1 and FB1 exhibit oscillations around their true values 25 times greater than under industrial average noise. Sensor values update every 2 s and are more pronounced than values for sensors not subjected to increases in sensor noise magnitude. These oscillations also impact the emergent features related to mass and energy balance. In this specific example, the total mass input into reservoir 1 shows a value below its true state. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

responded to noise in a similar fashion as the other indicators described thus far. However, the most salient impact of noise was on the emergent features (i.e., the yellow and red lines) connecting the mass and energy input and output indicators. As depicted in Fig. 6, even small values of noise could disrupt the vertical line reflecting a balanced inflow and outflow of mass or energy in the reservoirs. This offset from the vertical would give P+F participants the impression that the reservoirs were imbalanced when they were, in fact, balanced. More extreme cases of this imbalance at noise multipliers of 25 × (for FA1 and FA2 only) are shown in Fig. 7.

3.6. Procedure

3.6.1. Introductory session (2 h)

Each participant was introduced to the purposes and benefits of the study, completed a consent form, and an initial questionnaire. Then participants completed the Spy Ring History test (Pask and Scott, 1972; Pask, 1976). Previous research identified a task performance effect of the interaction between Spy Ring History test scores and EID versus non-EID interfaces (Torenvliet et al., 2000). The test requires participants to learn lists of ordered pairs that represent spies and their connections over a number of years (Pask, 1976; Howie, 1996). Patterns of answers were scored on three dimensions: Holist, Serialist, and Versatile, producing a summary of the cognitive style tendency of each person. Spy Ring scores were used to balance the groups. Using a minimum distance algorithm, pairs of participants with the lowest difference between them were computed, and the members of each pair were assigned to separate interface groups.

3.6.2. Training session (2 h)

Participants received a preliminary oral tutorial explaining the basics of DURESS III (independent of the interface). After completing the tutorial, each participant completed a brief activity to test his or her understanding of the system. The training and assessment material was the same as used in previous DURESS II studies. Participants were then introduced to the procedures for their respective interfaces (P+S or P+F). Participants were told that sensor readings have some normal "noise" around the true value. Finally, participants completed a brief activity to test their understanding of the operating procedures of DURESS III.

3.6.3. Experimental sessions (1 h each)

Participants completed the 80 experimental trials at a pace of 1 h per business day. The investigator did not ask any questions during the study, allowing the participants to explore different ways of controlling DURESS III. Each participant operated DURESS III for approximately 25–30 h in total, resulting in approximately 600 h of data across the 20 participants.

3.6.4. Debriefing session (1 h)

At the end of the study, a debriefing session captured participant comments on how they controlled the simulation and what effect the perturbations caused by sensor noise had on their actions, strategies, and performance. The experimenter also explained the hypotheses and goals of the study.

3.7. Measures

Five dependent variables were collected to assess the impact of sensor noise on task performance and control stability (see Table 4 and Fig. 8). Details on the calculation of these measures can be found in Yu et al. (1997). Trial Completion Time has served as the primary performance measure in prior DURESS studies (see Vicente, 2002 for a review). Hajdukiewicz and Vicente (2002) employed the control stability measures in their study of operation adaptation to novelty and change using EID and non-EID interfaces.

Table 4
Measures.

Construct	Measures	Definition
Performance	Trial completion time (TCT)	Time (in seconds) it took for participants to reach the commanded steady state condition (control target) starting from a shut-down state.
Control stability	Rise time (RT)	Elapsed time from the start of a trial until all a goal variable reached the target range.
	Number of Oscillations (NO)	Number of times the goal variables crossed above and below the target regions.
	Oscillation time (OT)	Time (in seconds) to stabilize the goal variables after they reached the target regions and before steady state was achieved.
	Maximum deviation (MD)	Maximum value by which the goal variables exceeded the target regions.

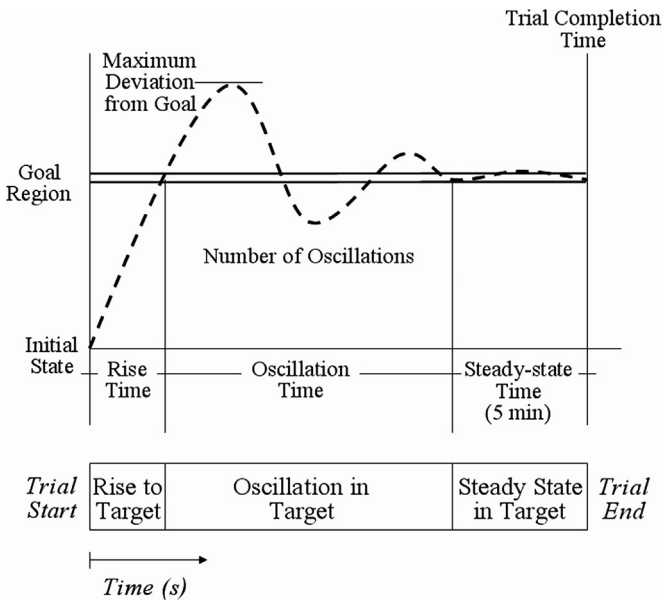

Fig. 8. Graphical depiction of task performance and control stability measures.

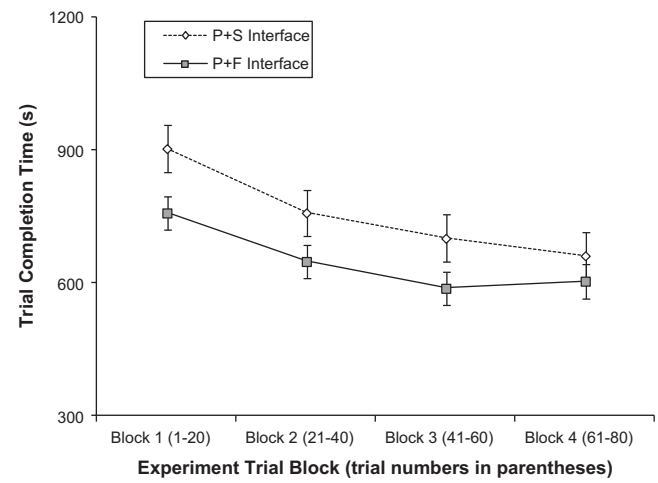

Fig. 9. Average TCT for each of the four experimental blocks. Error bars reflect plus and minus one standard error of the mean.

3.8. Hypotheses

Previous DURESS studies have consistently shown no difference in TCT between EID and non-EID interfaces under normal operating conditions (Vicente, 1996; Vicente, 2002). However, Hajdukiewicz and Vicente (2002) reported faster TCTs for the EID group compared to the non-EID group in one block of normal trials. They also reported control stability benefits (i.e., faster rise times, fewer oscillations and lower oscillation times) for the EID group in the same learning trials. In other studies, EID interfaces have been shown to support faster detection and more accurate diagnosis performance under fault conditions (Vicente, 2002).

The present study differs from prior DURESS studies in two important ways. First, the non-EID (i.e., P+S) interface contains more sensors than the non-EID interface used in prior studies. However, these sensors are also present in the P+F interface, so no differential benefit is anticipated for P+S participants. Second, the learning phase includes the presence of industry average noise whereas prior DURESS studies have not. However, given the low magnitudes of noise in these blocks, the performance impact is expected to be negligible in both interface conditions. Thus, we predict no effect of interface on TCT or control stability measures.

As the magnitude of sensor noise increases, performance and control stability is expected to worsen for both interface groups as the control task becomes more difficult. However, given the additional disturbances to the emergent features of the P+F interface, the decrements are expected to be more pronounced in the EID versus non-EID condition. This should manifest as an interaction effect for the performance and control stability between Blocks 3 and 4.

3.9. Data analysis

Data were stored in time-stamped logs that were collected by the computer in the background while participants completed the study. Previously-developed Matlab tools (see Yu et al., 1997) as well as other software tools (e.g., Microsoft Excel) were used to calculate the dependent measures. The measures were subjected to one-way ANOVA by block, two-way ANOVA between Blocks 3 and 4, and to linear regression for block 4 with interface and noise multiplier as factors. For all significance tests, we interpreted p values below .05 as offering strong indication of experimental effect, and p values between .05 and .1 as offering weak indication of experimental effect. Thus, we only report p values below 0.1.

4. Results

4.1. Task performance

Fig. 9 shows the average trial completion time across Blocks 1–4 for each Interface condition. In all four blocks, participants in the P+F condition completed trials significantly faster than participants in the P+S condition [Block 1: $F(1, 396)=17.75$, $p<.001$; Block 2: $F(1, 395)=17.79$, $p<.001$; Block 3: $F(1, 398)=23.09$, $p<.001$; Block 4: $F(1, 398)=7.11$, $p=.008$].

The two-way ANOVA on Blocks 3 and 4 showed a significant main effect for Interface, $F(1, 796)=28.56$, $p<.001$ and no significant main effect for Block. It also provided a weak indication of an interaction effect between Interface and Block, $F(1, 796)=3.05$, $p=.08$. TCTs in the P+S condition declined slightly between Blocks

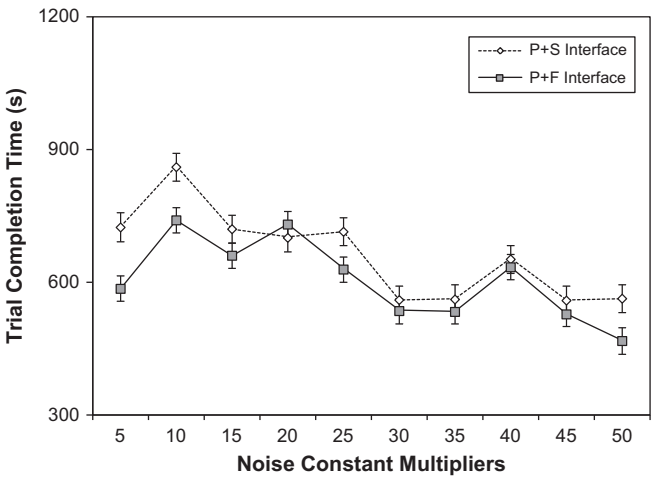

Fig. 10. Average TCT as a function of noise constant multipliers in block 4. Error bars reflect plus and minus one standard error of the mean.

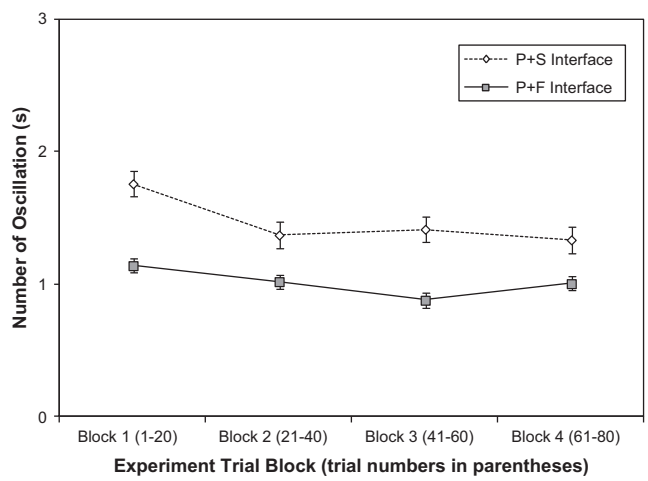

Fig. 11. Average Number of Oscillations for each of the four experimental blocks. Error bars reflect plus and minus one standard error of the mean.

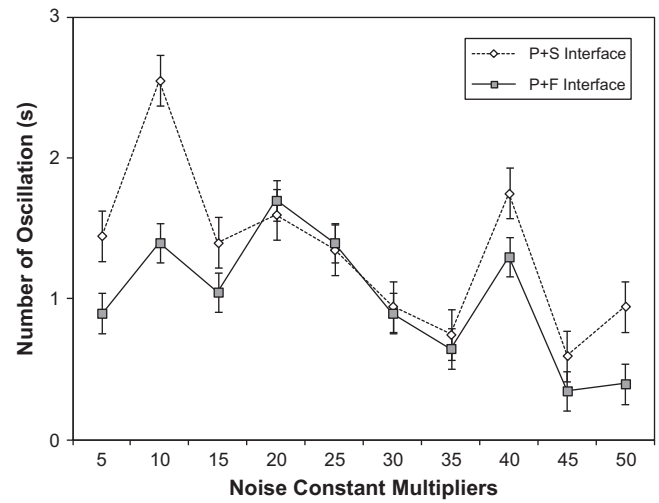

Fig. 12. Average Number of Oscillations in Block 4. Error bars reflect plus and minus one standard error of the mean.

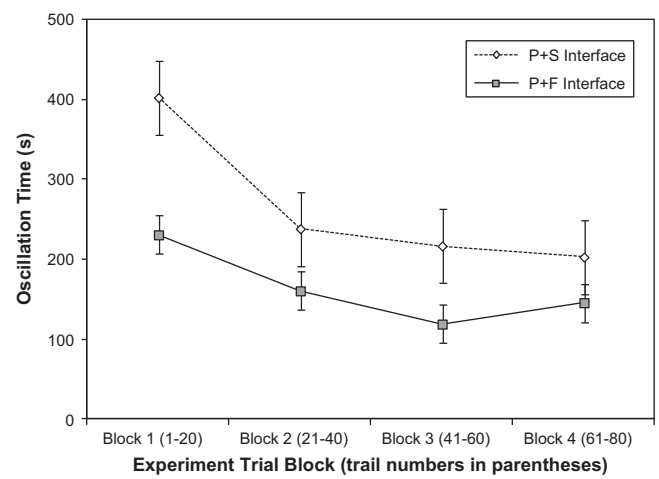

Fig. 13. Average Oscillation Time for each of the four experimental blocks. Error bars reflect plus and minus one standard error of the mean.

3 and 4 whereas TCTs in the P+F condition increased slightly across the two Blocks.

Multiple regression was used to test whether Interface, noise Multiplier or their interaction significantly predicted TCT in Block 4 (see Fig. 10). The results of the regression indicated that the three predictors explained 12% of the variance (R^2=.12, F(3, 396)=17.91, $p<.001$). Noise Multiplier predicted TCT ($\beta=-4.09$, $p<.0001$), as did Interface ($\beta=94.15$, $p=.03$). The interaction was not a significant predictor of TCT. However, we retained the interaction term in the model on account of its theoretical importance. To summarize, the higher the noise magnitude, the faster the TCT. And, at a given level of noise Multiplier, TCTs for participants using the P+S interface were approximately 94 s longer.

4.2. Control stability

Fig. 11 shows the average Number of Oscillations across Blocks 1–4 for each interface condition. There was a significant effect of interface on Number of Oscillations in all four Blocks [Block 1: $F(1, 396)=13.41$, $p<.001$; Block 2: $F(1, 395)=4.35$, $p=.04$; Block 3: $F(1, 398)=7.21$, $p=.008$; Block 4: $F(1, 398)=4.69$, $p=.03$]. In all cases, participants in the P+F condition experienced significantly fewer oscillations around the target goals.

The two-way ANOVA showed no significant main effect of Block. The main effect of interface was, once again, significant, $F(1, 796)=11.90$, $p<.001$. The interaction effect between Block and Interface was non-significant.

Multiple regression was used to test whether Interface, noise Multiplier or their interaction significantly predicted Number of Oscillations in Block 4 (see Fig. 12). The results of the regression indicated that the three predictors explained 5% of the variance (R^2=.05, $F(3, 396)=6.82$, $p<.001$). Noise Multiplier predicted Number of Oscillations ($\beta=-.02$, $p=.02$). Neither Interface nor the interaction were significant predictors of Number of Oscillations. However, we retained both terms in the model on account of their theoretical importance. To summarize, the higher the noise Magnitude, the fewer oscillations around the target goals.

Fig. 13 shows the average Oscillation Time across Blocks 1–4 for each interface condition. There was a significant effect of Interface on Oscillation Time in all four Blocks [Block 1: $F(1, 396)=21.87$, $p<.001$; Block 2: $F(1, 395)=7.96$, $p=.005$; Block 3: $F(1, 398)=18.37$, $p<.001$; Block 4: $F(1, 398)=6.80$, $p=.009$]. In all cases, participants in the P+F condition experienced significantly less Oscillation Time around the target goals.

The two-way ANOVA showed no significant effect of Block. The main effect of Interface was, once again, significant, $F(1, 796)=23.98$, $p<.001$. The interaction effect between Block and Interface was non-significant.

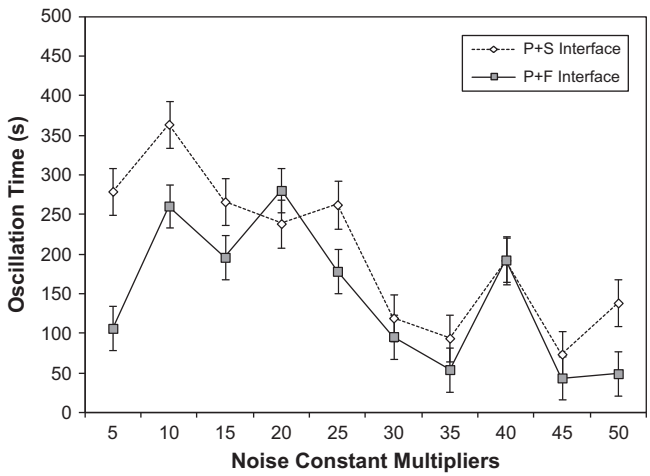

Fig. 14. Average Oscillation Time in Block 4. Error bars reflect plus and minus one standard error of the mean.

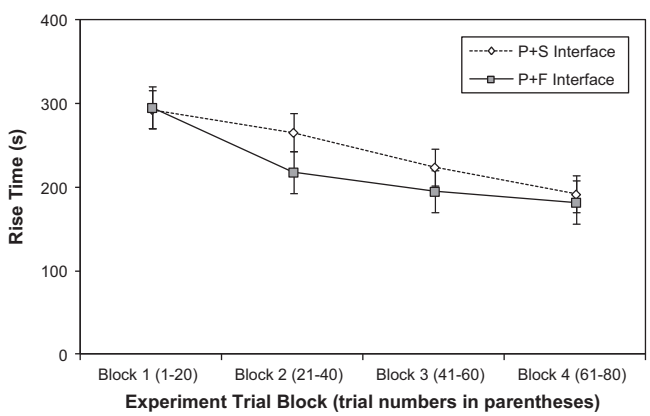

Fig. 15. Average Rise Time for each of the four experimental blocks. Error bars reflect plus and minus one standard error of the mean.

Multiple regression was used to test whether Interface, noise Multiplier or their interaction significantly predicted Oscillation Time in Block 4 (see Fig. 14). The results of the regression indicated that the three predictors explained 9.8% of the variance ($R^2=.098$, $F(3, 396)=14.33$, $p<.001$). Noise Multiplier predicted Oscillation Time ($\beta=-3.41$, $p=.001$), as did Interface ($\beta=104.81$, $p=.02$). The interaction was not a significant predictor of Oscillation Time. However, we retained the interaction term in the model on account of its theoretical importance. In summary, the higher the noise Magnitude, the less time was spent oscillating around the target goals. And, at a given level of noise Multiplier, Oscillation Times for participants using the P+S interface were approximately 105 s longer.

Fig. 15 shows the average rise time across Blocks 1–4 for each Interface condition. There was a significant effect of Interface on rise time in Blocks 2 and 3 [Block 2: $F(1, 395)=11.28$, $p<.001$; Block 3: $F(1, 398)=5.63$, $p=.02$]. In both cases, participants in the P+F condition initially reached the target range faster than participants in the P+S condition.

There was no significant effect of Interface on Maximum Deviation in any Block.

5. Discussion

The study presented in this paper addresses a high-priority research item for EID (Vicente, 2002). Since the introduction of the EID framework, Vicente and Rasmussen (1992), Vicente et al. (1996), and Vicente (2002) have pointed out that the presence and magnitude of sensor noise could have a negative impact on the robustness of EID interfaces. Our results show that an EID interface was robust to sensor noise in terms of impacts on both task performance and control stability. This was observed for trials in both the perturbation phase where noise magnitude sharply increased and the learning phase with industry average sensor noise. We discuss each of these findings in turn before considering several related perspectives.

5.1. Task performance and control stability in the perturbation phase

We hypothesized that the support provided by the EID interface would progressively degrade as increasing noise magnitude disrupted the emergent features. Our expectation was that these disruptions would impair the ability of operators to make use of the higher-order state variables and relations depicted in the EID interface. This did not obtain. Participants using an EID interface completed trials 8.6% faster on average throughout the perturbation phase, completing trials nearly one minute faster than their counterparts using the non-EID interface. As well, participants using an EID interface exhibited more stable control performance under elevated noise conditions.

In some ways this finding mirrors Hajdukiewicz and Vicente (2002). In their perturbation trials they found that an EID interface was more robust to work domain disturbances in the form of changing time constants as compared to a non-EID interface. The finding is also consistent with other EID studies (Vicente, 2002), where the benefits of the higher-order state variables and emergent features have obtained in fault or beyond-design basis trials and scenarios (Burns et al., 2008; Jamieson, 2007; Lau et al., 2008b; Reising and Sanderson, 2004). Taken as a whole, results from the perturbation phase show that the robustness of the EID interface was not compromised by increases in the magnitude of sensor noise, even for exceptionally high noise magnitudes.

5.2. Task performance and control stability in the learning phase

The task performance (i.e., TCT) benefit in the learning phase stands in clear contrast to prior DURESS studies, which have reported no statistically significant advantage for EID interfaces in terms of mean task completion times under normal conditions[3] (Vicente, 2002). Our findings are however consistent with the findings from the learning phase of Hajdukiewicz and Vicente (2002), which did not include sensor noise. Using DURESS II they report faster trial completion times, fewer oscillations around the target range and less time spent oscillating around the target range for the EID interface group on the last normal block before introducing perturbations. Moreover, we have analyzed the performance data from the learning phase of that study and find that trial completion times were faster for the EID interface group across the entire learning phase.[4] Thus, the same task performance and similar control stability (cf. rise times were faster in the non-EID condition for Hajdukiewicz and Vicente, 2002) benefits were obtained from the two studies, both of which stand in contrast to other EID research.

What methodological details common to Hajdukiewicz and Vicente (2002) and the present study – but uncommon to all prior DURESS II studies – might explain the emergence of a significant task performance benefit for users of EID interfaces over non-EID

[3] However, the non-EID interface led to significantly less *consistent* performance than did the EID interface.

[4] These results from Blocks 1 and 2 are not reported in Hajdukiewicz (2001) or Hajdukiewicz and Vicente (2002).

interfaces? The most relevant systematic difference in our view is that both studies applied the minimum distance calculator to the Spy Ring test scores to match participant pairs across interface conditions. This method of assigning participants to interface groups was first adopted by Hajdukiewicz (2001). Torenvliet et al. (2000) provided strong evidence that the interaction of high holist score and assignment to the EID interface predicted faster trial completion whereas the interaction of high serialist score and assignment to the non-EID interface predicted slower trial completion. It is possible that, in DURESS studies prior to Hajdukiewicz (2001), high holist score individuals were under-represented in EID interface conditions and/or that high serialist individuals were under-represented in the non-EID interface conditions. This would seem to be a testable hypothesis.

We are unable to find an alternate methodological detail that is present in the two studies showing the learning phase performance effect and absent from studies not showing the effect. To wit:

- Hajdukiewicz and Vicente (2002) included no industry average sensor noise in the learning phase in keeping with prior DURESS studies. The present study includes industry average sensor noise in the learning phase.
- Hajdukiewicz and Vicente (2002) used a .1 s refresh rate in keeping with prior DURESS studies. The present study used a refresh rate of 2 s.
- Hajdukiewicz and Vicente (2002) used constant temperature goals in keeping with prior DURESS studies. The present study varied temperature goals for each trial within a block.
- Hajdukiewicz and Vicente (2002) included two fault trials in the first block of their learning phase. The present study excludes fault trials from the learning phase.

In the absence of an alternate explanation for the observed effects, we return to the fundamental claim of EID; that providing redundant constraints in the interface supports effective control (Vicente and Rasmussen, 1992). Under this explanation, EID participants benefitted from the abstract higher-level information that is not portrayed in the non-EID interface. Taken as a whole, results from the learning phase show that the robustness of the EID interface was not compromised by the presence of industry average sensor noise.

5.3. Learning and adaptation

Both interface groups showed learning effects. Task performance and control stability measures improved throughout the learning phase and further in the face of increasingly high noise levels. We see two possible explanations for this finding. One is that participants in both interface conditions adopted control strategies in the learning phase that were robust to the increasing noise and could be incrementally honed as noise magnitude increased. A second explanation is that participants effectively modified their control strategies in response to increasing noise magnitude (see St-Cyr, 2006). Either way, participants learned how to cope with perturbations and adapted their control actions accordingly, even though sensor noise magnitude was increasing.

5.4. Topographic and derivational adequacy

These results may be seen as striking a contrast with the findings of Reising and Sanderson (2004), who hypothesized that a well-design EID interface would provide better support for fault diagnosis than a non-EID interface. This hypothesis held for a work domain with a topographically and derivationally adequate sensor set. However, they observed more pronounced drops in failure diagnosis performance for the EID interface than for a non-EID interface in transitioning from the maximal to a minimal sensor set (although performance with the EID interface was not significantly worse than that with the non-EID interface). In an important difference from the present study, however, participants in Reising and Sanderson (2004) were exposed to noisy sensors plus sensor faults, as opposed to only sensor noise around the true value.

In the current study, we have observed superior task performance and control stability for an EID interface over a non-EID interface for a work domain with a topographically inadequate sensor set. Had our study involved noisy sensors plus sensor faults, we may have observed a decrease in performance and control stability for the EID interface. Nevertheless, our results suggest that, for some control tasks, the robustness of the EID framework was not affected by a topographically inadequate sensor set. The comparison is, of course, oblique in that control task performance and stability are distinctly different from fault diagnosis. Elsewhere we have shown that EID interfaces can show benefits on some dependent variables and not others. Therefore, an important topic for future research is to explore fault trial performance under sensor noise conditions.

5.5. Toward a new response to Maddox (1996)

The results also provide a new, albeit partial, response to Maddox's (1996) critique of an earlier DURESS study (Christoffersen et al., 1996). Maddox pointed out that the original non-EID (i.e., the P) interface in DURESS II contained fewer sensors than the EID interface. He challenged Christoffersen and colleagues to justify their claim that the observed performance benefits of the EID over non-EID interface could be attributed to the abstract function information present in the P+F interface and not to the additional sensors. The P+S interface used in our study is effectively the non-EID interface for DURESS that Maddox (1996) specified in his critique.

As pointed out earlier, one would expect users in the non-EID interface (i.e., P+S in our study) to benefit from the additional sensors (i.e., they were provided with flow indicators in addition to component settings). Reising and Sanderson (2004) made the same addition to their non-EID interfaces. They observed fault diagnosis performance benefits for the EID condition under these conservative design decisions. Similarly, our results show that despite the non-EID interface having the same sensors as the EID interface, the task performance and control stability of participants in the EID interface condition were superior to that of participants in the non-EID condition. This provides some evidence in support of the analytical redundancy argument; namely that abstract function information built into the P+F interface benefits participants in the EID condition.

Further comparisons based on Maddox's (1996) critique are warranted, but outside the scope of this paper. We could combine the results of the current study with those of St-Cyr and Vicente (2005) and compare the performance and control stability for the original P interface, the P+S interface, and the P+F interface. However, the Christoffersen et al. (1996) study to which Maddox offered his critique examined performance on fault trials in terms of fault detection time and diagnosis accuracy. Our study did not include faults of this type. A complete response to Maddox's (1996) critique would require a replication of the study with all three interface conditions and the introduction of the fault trials from Christoffersen et al. (1996).

6. Conclusions

The purpose of this paper was to understand the impacts of the presence and magnitude of sensor noise on operators'

performance and control stability using an EID and non-EID interface. There are two main findings. First, participants in the EID condition performed significantly better and demonstrated more control stability than participants in the non-EID condition, both under the presence of industry average sensor noise and under increased noise magnitude. This clearly shows that the robustness of EID interfaces is not compromised by sensor noise. More broadly, these results add weight to the claim that EID interfaces are robust to a variety of process and instrumentation disturbances that might be expected to challenge the design framework.

Second, the task performance and control stability results were obtained despite the DURESS III apparatus exhibiting a topographically inadequate sensor set. Whereas Reising and Sanderson (2004) found that an inadequate sensor set compromised participants' ability to diagnose faults using the emergent features in an EID interface, we did not observe a similar deficit for control task performance or stability of control. It is possible that, with a more analytically adequate sensor set, either EID or non-EID users might show higher levels of control or task performance, but that comparison is not possible with the current DURESS III configuration.

Despite the consistency of our findings, this study has limitations that motivate future research. First, the current study was performed with a specific type of perturbations – gradual increases in the magnitude of sensor noise visibly manifesting as erratic readings on the interface. It is not known how the results generalize to other types of perturbations, more abrupt onsets, or perturbations that are not as readily visible (such as slow drifts or sensors reading at a constant value). Second, it is not known how participants in the two conditions would have performed had there been further sensors (i.e., maximal instrumentation) that reduced the impact of the noisy sensors on the derivations required to calculate the higher-order variables that were mapped onto display emergent features. Third, the simulation program, DURESS III, was limited in scale compared with industrial systems. This limits the extent to which the results can be generalized to larger scale systems. Fourth, other characteristics of process control systems such as sensor noise bandwidth, component time constants, and display refresh rates may have an impact on operators' ability to deal with noisy sensors and need to be investigated more fully.

The notion of sensor noise adopted in this study has its roots in analog technology that is becoming anachronistic in contemporary process facilities. The transition to digital instrumentation and control is well under way in the petrochemical and increasingly in the nuclear industry. Moreover, emerging computational and probabilistic reasoning engines allow for the placement of simulated or even virtual sensors in a plant model. These contemporary and emerging instrumentation and control technologies open up powerful opportunities for the designers of emergent feature displays (e.g., Lau and Jamieson, 2006). However, they also raise concerns about the cognitive engineering issues associate with degradation and failure of digital instrumentation and control systems (O'Hara et al., 2010). It seems likely, therefore, that sensor reliability may become an increasingly important topic in both microworld and simulator studies.

Acknowledgments

Funds to support this work came from the Natural Science and Engineering Council of Canada (NSERC), the E. W. R. Steacie Memorial Fellowship, and the University of Toronto Glynn Williams Fellowship foundation. The authors would like to thank Ridha Ben Mrad, David Woods, Gerard Torenvliet, and the three reviewers, Kevin B. Bennett, Penelope M. Sanderson, and Dal Vernon C. Reising, for their thoughtful comments on and discussions about the ideas presented herein.

References

Anyakora, S.N., Lees, F.P., 1974. Detection of instrument malfunction by the process operator. In: Edwards, E., Lees, F.P. (Eds.), The Human Operator in Process Control. Taylor and Francis, London, UK, pp. 238–248.
Bennett, K.B., Flach, J.M., 1992. Graphical displays: implications for divided attention, focused attention, and problem solving. Human Factors 34, 513–533.
Bennett, K.B., Flach, J.M., 2011. Display and interface design: Subtle science, exact art. CRC Press, Boca Raton, FL.
Bennett, K.B., Toms, M.L., Woods, D.D., 1993. Emergent features and graphical elements: designing more effective configural displays. Human Factors 35, 71–97.
Bisantz, A.M., Vicente, K.J., 1994. Making the abstraction hierarchy concrete. International Journal of Human–Computer Studies 40, 83–117.
Burns, C.M., 2000a. Navigation strategies with ecological displays. International Journal of Human–Computer Studies 52, 111–129.
Burns, C.M., 2000b. Putting it all together: display integration in ecological displays. Human Factors 46, 226–241.
Burns, C.M., Hajdukiewicz, J.R., 2004. Ecological Interface Design. CRC Press, Boca Raton, FL.
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.
Christoffersen, K., Hunter, C.N., Vicente, K.J., 1996. A longitudinal study of the effects of ecological interface design on skill acquisition. Human Factors 38, 523–541.
Drivalou, S., Marmaras, N., 2009. Supporting skill, rule-, and knowledge-based behaviour through an ecological interface: an industry-scale application. International Journal of Industrial Ergonomics 39, 947–965.
Frank, P.M., 1990. Fault diagnosis in dynamic systems using analytical and knowledge-based redundancy – A survey and some new results. Automatica 20, 459–474.
Hajdukiewicz, J.R., 2001. Adapting to change in dynamic worlds: a study of higher-level control and key success factors in a process control microworld (Unpublished PhD Thesis). University of Toronto, Department of Mechanical and Industrial Engineering.
Hajdukiewicz, J.R., Vicente, K.J., 2002. Designing for adaptation to novelty and change: functional information, emergent feature graphics, and higher-level control. Human Factors 44, 592–610.
Hajdukiewicz, J.R., Vicente, K.J., 2004. what does computer-mediated control of a thermal-hydraulic system have to do with moving your jaw to speak? Evidence for synergies in process control. Ecological Psychology 16, 255–285.
Ham, D.-H., Yoon, W.C., 2001a. Design of information content and layout for process control based on goal-means domain analysis. Cognition, Technology & Work 3, 205–223.
Ham, D.-H., Yoon, W.C., 2001b. The effects of presenting functionally abstracted information in fault diagnosis tasks. Reliability Engineering and System Safety 73, 103–119.
Ham, D.-H., Yoon, W.C., Han, B.-T., 2008. Experimental study on the effects of visualized functionally abstracted information on process control tasks. Reliability Engineering and System Safety 93, 254–270.
Howie, D., 1996. Shaping expertise through ecological interface design: strategies, metacognition, and individual differences. Cognitive Engineering Laboratory Technical Reports Series: CEL 96-01. University of Toronto, Department of Mechanical and Industrial Engineering.
Jamieson, G.A., 2007. Ecological interface design for petrochemical process control: an empirical assessment. IEEE Transactions on Systems, Man And Cybernetics 37, 906–920.
Jamieson, G.A., Vicente, K.J., 2001. Ecological interface design for petrochemical applications: Supporting operator adaptation, continuous learning, and distributed, collaborative work. Computers and Chemical Engineering 25, 1055–1074.
Jamieson, G.A., Vicente, K.J., 2005. Designing effective human-automation-plant interfaces: a control-theoretic perspective. Human Factors 47, 12–34.
Jamieson, G.A., Miller, C.A., Ho, W.H., Vicente, K.J., 2007. Integrating task-and work domain-based work analysis in ecological interface design: a process control case study. IEEE Transactions on Systems, Man And Cybernetics 37, 887–905.
Kim, S.K., Suh, S.M., Jang, G.S., Hong, S.K., Park, J.C., 2012. Empirical research on an ecological interface design for improving situation awareness of operators in an advanced control room. Nuclear Engineering and Design 235, 226–237.
Lau, N., Jamieson, G.A., 2006. Numerical models in representation design: computing seawater properties in an ecological interface. In: Proceedings of the 50th Annual Meeting of the Human Factors and Ergonomics Society. Santa Monica, CA: Human Factors and Ergonomics Society, pp. 245–249.
Lau, N., Veland, O., Kwok, J., Jamieson, G.A., Burns, C.M., Welch, R., Braseth, A.O., 2008a. 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.
Lau, N., Jamieson, G.A., Skraaning, G., Burns, C.M., 2008b. 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.
Maddox, M.E., 1996. Critique of "A longitudinal study of the effects of ecological interface design on skill acquisition" by Christoffersen, Hunter, and Vicente. Human Factors 38, 542–545.
Moray, N., 1986. Monitoring Behavior and Supervisory Control. In: Boff, K.R., Kaufman, L., Thomas, J.P. (Eds.), Handbook of Perception and Human Performance, 1-51. Wiley, New York, p. 40.
O'Hara, J., Higgins, J., Kramer, J., 2000. Advanced information systems design: Technical basis and human factors review guidance (NUREG/CR-6633). U.S. Nuclear Regulatory Commission, Washington, DC.
O'Hara, J., Guenther, W., Martinez-Guridi, G., 2010. The Effects of Degraded Digital Instrumentation and Control Systems on Human-System Interfaces and Operator Performance. Brookhaven National Laboratory.
Pask, G., 1976. Styles and strategies of learning. British Journal of Educational Psychology 46, 128–148.
Pask, G., Scott, B.C., 1972. Learning styles and individual competence. International Journal of Man-Machine Studies 4, 217–253.
Pawlak, W., Vicente, K.J., 1996. Inducing effective operator control through ecological interface design. International Journal of Human-Computer Studies 44, 653–688.
Pomerantz, J.R., Pristach, E.A., 1989. Emergent features, attention, and perceptual glue in visual form perception. Journal of Experimental Psychology: Human Perception & Performance 15, 635–649.
Rasmussen, J., 1985. The role of hierarchical knowledge representation in decision making and system management. IEEE Transactions on Systems, Man and Cybernetics, SMC 15, 234–243.
Reising, D.V.C., Sanderson, P.M., 2002a. Work domain analysis and sensors I: principles and simple example. International Journal of Human-Computer Studies 56, 569–596.
Reising, D.V.C., Sanderson, P.M., 2002b. Work domain analysis and sensors II: pasteurizer II case study. International Journal of Human-Computer Studies 56, 597–637.
Reising, D.V.C., Sanderson, P.M., 2002c. Ecological interface design for Pasteurizer II: a process description of semantic mapping. Human Factors 44, 222–247.
Reising, D.V.C., Sanderson, P.M., 2004. Minimal instrumentation may compromise failure diagnosis with an ecological interface. Human Factors 46, 316–333.
Sanderson, P.M., Flach, J.M., Buttifieg, M.A., Casey, E.J., 1989. Object displays do not always support better integrated task performance. Human Factors 31, 183–198.
St-Cyr, O., 2006. Impact of Sensor Noise Magnitude on Emergent Features of Ecological Interface Design. In: Proceedings of the Human Factors and Ergonomics Society 50th Annual Meeting. Santa Monica, CA: Human Factors and Ergonomics Society, pp. 319–323.
St-Cyr, O., Vicente, K.J., 2005. Sensor noise and ecological interface design: effects of increasing noise magnitude on operators' performance. In: Proceedings of the Human Factors and Ergonomics Society 49th Annual Meeting. Santa Monica, CA: Human Factors and Ergonomics Society pp. 417–421.
Stein, P.K., 1969. The response of transducers to their environment: The problem of signal and noise. 17. Lf/MSE Publication, pp. 1–15.
Torenvliet, G.L., Jamieson, G.A., Vicente, K.J., 2000. Making the most of ecological interface design: the role of individual differences. Applied Ergonomics 31, 395–408.
Upton, C., Doherty, 2008. Extending ecological interface design principles: a manufacturing case study. International Journal of Human Computer Studies 66, 271–286.
Vicente, K.J., 1996. Improving dynamic decision making in complex systems through ecological interface design: a research overview. System Dynamics Review 12, 251–279.
Vicente, K.J., 1999. Cognitive Work Analysis: Towards Safe, Productive, and Healthy Computer-based Work. Lawrence Erlbaum Associates, Mahwah: NJ.
Vicente, K.J., 2002. Ecological interface design: Progress and challenges. Human Factors, 44; , pp. 62–78.
Vicente, K.J., Rasmussen, J., 1992. Ecological interface design: theoretical foundations. IEEE Transactions on Systems, Man and Cybernetics, SMC 22, 589–606.
Vicente, K.J., Moray, N., Lee, J.D., Rasmussen, J., Jones, B.G., Brock, R., Djemil, T., 1996. Evaluation of a rankine cycle display for nuclear power plant monitoring and diagnosis. Human Factors 38, 506–521.
Watanabe, O., 2001. Issues toward practical application of ecological interface to real world. Paper presented at the International Workshop on Ecological Approach to Human Work Interface, Tokai, Japan.
Yu, X., Farzad, K.S., Lau, E., Vicente, K.J., Carter, M.W., 1997. ADAPT user's manual: A data analysis tool for human performance evaluation in dynamic systems. Cognitive Engineering Laboratory Technical Reports Series: CEL 97-03. University of Toronto, Department of Mechanical and Industrial Engineering.