Designing for Adaptation to Novelty and Change: Functional Information, Emergent Feature Graphics, and Higher-Level Control

John R. Hajdukiewicz and **Kim J. Vicente,** University of Toronto, Toronto, Ontario, Canada

Ecological interface design (EID) is a theoretical framework that aims to support worker adaptation to change and novelty in complex systems. Previous evaluations of EID have emphasized representativeness to enhance generalizability of results to operational settings. The research presented here is complementary, emphasizing experimental control to enhance theory building. Two experiments were conducted to test the impact of functional information and emergent feature graphics on adaptation to novelty and change in a thermal-hydraulic process control microworld. Presenting functional information in an interface using emergent features encouraged experienced participants to become perceptually coupled to the interface and thereby to exhibit higher-level control and more successful adaptation to unanticipated events. The absence of functional information or of emergent features generally led to lower-level control and less success at adaptation, the exception being a minority of participants who compensated by relying on analytical reasoning. These findings may have practical implications for shaping coordination in complex systems and fundamental implications for the development of a general unified theory of coordination for the technical, human, and social sciences. Actual or potential applications of this research include the design of human-computer interfaces that improve safety in complex sociotechnical systems.

INTRODUCTION

The very idea of deliberately designing to help workers adapt to novelty and change in complex work domains seems paradoxical. On the one hand, there is a growing recognition that the need to support worker adaptation is increasing (Hirschhorn, 1984; Nickerson, 1995; Rasmussen, Pejtersen, & Goodstein, 1994; Vicente, 1999). The pace of change in society is such that the job demands experienced by workers are evolving rapidly in unpredictable ways. Blum (1996, p. 160) summarized the implication succinctly: "Change is the invariant, not knowledge." One the other hand, it can seem impossible to design information support for events that are unknown and unforeseeable at the time of design (Shepherd, 1993). If future demands cannot be predicted, how can anyone possibly design a system to support those demands? From this perspective the idea of designing for worker adaptation to novelty and change seems unattainable, by definition.

Ecological interface design (EID) is a theoretical framework for human-computer interface design that aims to support worker adaptation to novelty and change (Vicente & Rasmussen, 1990, 1992). EID relies on Rasmussen's (1985) abstraction hierarchy (AH) to identify the information requirements for interface design. The physical and functional information identified by an AH model is then displayed in a form that allows experienced workers to rely on their powerful perception-action capabilities, rather than on their more resource-limited analytical capabilities, to control the work domain.

Address correspondence to Kim J. Vicente, Cognitive Engineering Laboratory, Department of Mechanical and Industrial Engineering, University of Toronto, 5 King's College Rd., Toronto, Ontario, Canada M5S 3G8; vicente@mie.utoronto.ca.
HUMAN FACTORS, Vol. 44, No. 4, Winter 2002, pp. 592–610.

The empirical literature on EID (see Vicente, 2002, for a review) has emphasized representativeness over experimental control. This choice has enhanced the generalizability of results to operational settings. However, that contribution to applied research has been obtained at a cost to basic research; the ability to develop defensible causal explanations for the observed effects is limited. To redress this imbalance, complementary research emphasizing experimental control should be conducted to foster fundamental theory building (Vicente, 2000). This article addresses that objective.

Background

The literature on EID suffers from several limitations when it comes to evaluating rigorously the framework's comparative ability to support adaptation to novelty and change. First, many studies that have investigated this issue have injected faults into dynamic scenarios with interactive simulations (e.g., Christoffersen, Hunter, & Vicente, 1997; Ham & Yoon, 2001; Reising & Sanderson, 2000). These scenarios benefit from being representative of the types of novel events encountered by workers in operational settings. However, in an interactive simulation the demands imposed by a fault depend on the work domain state and on the task that the participant is pursuing when a fault occurs. Furthermore, because participants may have the simulation in a different state or configuration when the fault occurs, "the very same fault can be very difficult for one participant, relatively easy for another, and not even experienced by a third (e.g., if they are not using the component that failed)" (Christoffersen et al., 1997, p. 19). As a result, the use of fault scenarios, though representative, can result in less experimental control.

Second, the studies cited have compared an EID interface with a baseline interface for novel scenarios and have consistently obtained a main effect in favor of EID. This result is important from a practical perspective because it demonstrates that EID can lead to better performance with unanticipated scenarios than can design approaches being used in industry. However, this finding does not allow one to conclude that EID provides greater support for adaptation to novelty and change. That inference can be based only on an interaction, not a main effect. One would have to compare the change in performance between familiar and novel conditions for each interface type and to demonstrate that the introduction of novel scenarios has a significantly smaller impact on the EID interface than does the baseline. Moreover, the scenarios in the familiar and novel conditions should be identical in all other ways (e.g., the values of the goal variables). As far as we know, a comparison satisfying these conditions has yet to be conducted.

Third, and relatedly, to evaluate an interaction between interface type and scenario novelty, a common metric must be used. As noted, however, most studies have used fault scenarios as stimuli, so the usual metric for evaluating performance under novel conditions has been fault diagnosis accuracy. This measure may be difficult to apply to familiar (i.e., nonfault) scenarios.

Finally, most empirical research on EID has focused on comparing the performance impact of this approach with that of traditional design approaches (Vicente, 2002). Less attention has been devoted to investigating rigorously *why* EID leads to performance improvements. Two issues in particular are outstanding. First, there is a suggestion that the benefits of EID are obtained because EID leads to higher-level control, whereas traditional designs lead to lower-level control (Yu, Lau, Vicente, & Carter, 2002). *Higher-level control* focuses on higher-level functions in the abstraction hierarchy by using any feasible lower-level components in a context-specific way. In this case the higher-level functions and outcome (i.e., products) are consistent across situations, whereas the lower-level actions vary from one situation to the next. This phenomenon is known as *context-conditioned variability* in the motor control literature (Turvey, Fitch, & Tuller, 1982).

In contrast, *lower-level control* focuses on a specific way of using components at the bottom of the abstraction hierarchy in a context-independent way, much like a rote recipe. In this case the lower-level actions (i.e., processes) are consistent across situations. The hypothesis that EID leads to better adaptation to novelty and change because of a greater reliance on higher-level control has yet to be tested rigorously.

Second, there is also a suggestion that the benefits of EID and the reliance on higher-level control are possible because of the emergent features graphics (displays that reveal higher-level invariant properties via the relationships between lower-level data; Buttigieg & Sanderson, 1991) that are typically found in interfaces based on EID (Vicente & Rasmussen, 1990, 1992). In an EID display participants may become perceptually coupled to the emergent feature and thereby pick up the functional information that is said to be important for adaptation to change and novelty. In contrast, in the absence of emergent feature graphics, participants would have to derive that functional information using resource-intensive analytical processes. The hypothesis that EID leads to better adaptation to novelty and change because of perceptual coupling to emergent feature graphics has also not been tested rigorously.

The research in this article addresses all of these limitations. Performance on novel and familiar events was compared rigorously using an interface based on EID and an interface based on traditional design practices. The novel events, referred to as *perturbations,* were experienced in the same way by all participants, thereby avoiding the problems associated with fault trials in interactive simulations. The Scenario × Interface interaction was investigated using common metrics so that inferences could be made about the degree of adaptation to novelty and change. Also, the hypotheses that higher-level control and emergent feature graphics might be responsible for the advantages of EID were investigated directly in a controlled fashion.

EXPERIMENTAL TEST BED

Two experiments were conducted in the context of a representative microworld simulation of a thermal-hydraulic process, shown in Figure 1. The process consists of two redundant feedwater

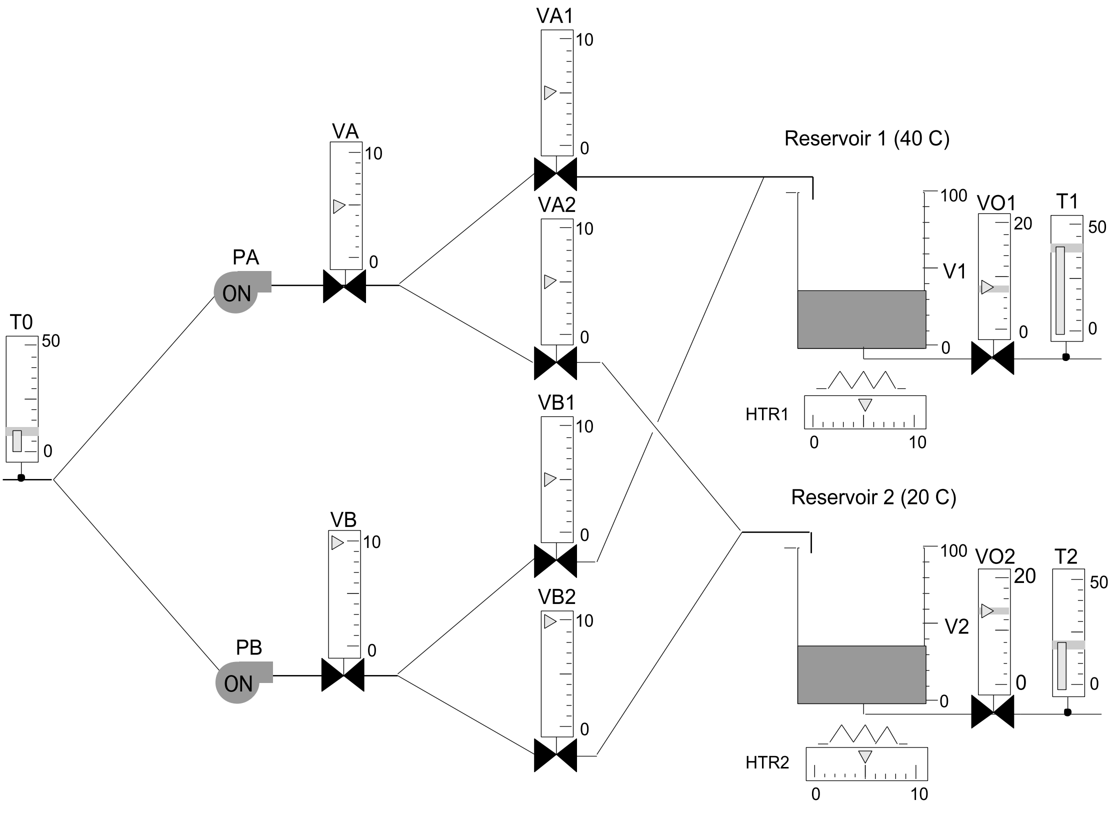

Figure 1. The P interface for the thermal-hydraulic process simulation microworld used in this research. Adapted from the *International Journal of Human-Computer Studies,* Vol. 44, W. S. Pawlak and K. J. Vicente, “Inducing effective operator control through ecological interface design,” pp. 653–688. Copyright 1996, with permission from Elsevier Science.

streams (FWSs) that can be configured to supply water to either of the two reservoirs. Associated with each reservoir is an externally determined demand for water that can change over time. The simulation purposes are twofold: to keep each of the reservoirs at a prescribed temperature (40°C and 20°C) and to satisfy the current mass (water) output demand rates. To satisfy these purposes, participants were given control over eight valves (VA, VA1, VA2, VO1, VB, VB1, VB2, and VO2), two pumps (PA, PB), and two heaters (HTR1, HTR2). Normally, the time constants for the heaters and the remaining components were 15 and 5 s, respectively.

Two interfaces were constructed for this microworld (Pawlak & Vicente, 1996; Vicente & Rasmussen, 1990). The P interface in Figure 1 presents only a subset of the information identified by an abstraction hierarchy analysis, focusing on physical information. The form of the interface uses a mimic diagram format, the industry standard in process control. The P+F interface in Figure 2 presents all the information identified by an AH analysis. The physical information is presented using a mimic diagram format, whereas the functional information takes advantage of the emergent feature graphics shown on the right side of Figure 2.

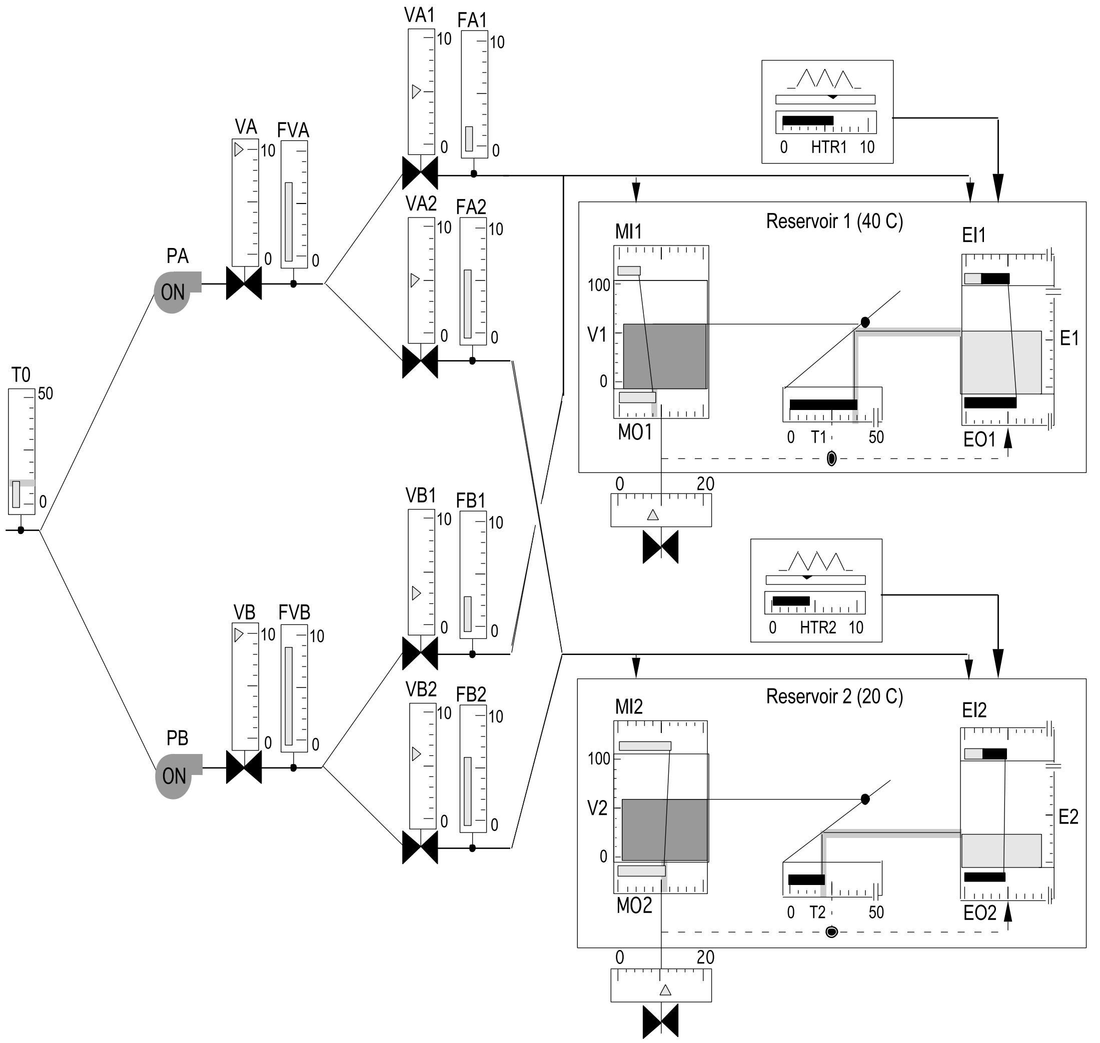

Figure 2. The P+F interface for the thermal-hydraulic process simulation microworld used in this research. Adapted from the *International Journal of Human-Computer Studies,* Vol. 44, W. S. Pawlak and K. J. Vicente, "Inducing effective operator control through ecological interface design," pp. 653–688. Copyright 1996, with permission from Elsevier Science.

According to EID, the P+F interface should be more likely than the P interface to lead to successful adaptation to novelty and change. Because the emergent feature graphics in the former interface revealed the functional levels of information in the abstraction hierarchy, it should have been easier for participants to exhibit higher-level control by relying on perceptual coupling. In contrast, those functional levels of information were not visible in the P interface, so participants had to engage in analytical reasoning to exhibit higher-level control. If such cognitive demands exceeded the P participants' limited resources, then those participants would tend to exhibit lower-level control and, thus, less success in adapting to novelty and change. These predictions need not hold. The P+F interface could lead to worse performance because it is visually more complex than the P. Also, the P interface could lead to higher-level control if participants were to use their knowledge to compensate for the missing functional information. Conversely, the P+F interface could lead to lower-level control if participants were to use heuristics rather than functional information to control the process.

EXPERIMENT 1

The purpose of the first experiment was to investigate the impact of dynamic perturbations on success in adaptation and on level of control using the P and P+F interfaces with experienced participants. EID predicts that participants with the P+F interface should be better able to adapt to the perturbations than should the P participants. Also, EID predicts that those participants who are able to adapt successfully with the P+F interface would exhibit signs of higher-level control. In contrast, the P participants who are not able to adapt successfully should exhibit signs of lower-level control.

Method

Experimental design. A repeated-measures, mixed design was adopted with interface (P or P+F) as a between-participants variable and trial type (normal and perturbation) as a within-participants variable. The study consisted of 80 trials divided into four matched blocks of 20 trials. The first three blocks were the same (with one exception, noted later) and provided an opportunity for participants to learn how to control the microworld under normal conditions. During the last block the time constants of all of the components were uniformly but randomly changed from trial to trial by multiplying their normal values by specific factors (0.1, 1.0, 2.5, 5.0, 7.5, and 10.0×). The 0.1× multiplier corresponded to more responsive component dynamics, whereas the 10.0× multiplier corresponded to less responsive component dynamics. Two perturbation trials were also included in the first block (7.5× in Trial 2; 0.1× in Trial 4).

Participants. There were eight paid participants in each of the two interface groups. All were university engineering students who were selected primarily based on their willingness to participate, their degree of relevant formal training (i.e., one to three courses in fluid mechanics and thermodynamics), and their cognitive style (discussed in the procedure section).

Apparatus. The interactive, dynamic simulation was programmed on a Silicon Graphics, Inc. (Mountain View, CA) computer with an IRIX operating system. Information was presented on a 19-inch (48-cm), high-resolution color graphics monitor. Control actions were input using a computer mouse. Each time an input was made, the state of all of the simulation variables was automatically logged for data analysis.

Task. For each trial participants were required to start up the process and achieve steady state by satisfying the four target conditions simultaneously (i.e., two outflow demand and two temperature goals) for 5 consecutive min within a 30-min time limit.

Procedure. In the introductory session participants read an explanation of the purposes of the experiment and filled out a consent form and an initial questionnaire. They then completed the Spy Ring History Test (Pask & Scott, 1972) so that their cognitive style could be assessed. The results of this test were used to assign participants to the two interface groups to control for cognitive style (see Torenvliet, Jamieson, & Vicente, 2000).

Participants were then given a tutorial on the components of the microworld (independent of the interface). After completing the tutorial, each participant was given a brief exercise. Next, the specific interface assigned to each participant (P or P+F) was introduced; the discussion centered only on the elements of the display, not the functioning of the simulation.

In each subsequent session participants performed several trials requiring them to control the simulation. Each session lasted approximately 1 hr/day. Each participant operated the simulation for approximately 25 hr in total, resulting in approximately 400 participant hours of data collection. Before their first trial, and immediately after both the 60th and 80th trials, participants were asked to write out a control recipe (discussed later). At the end of the experiment each participant was debriefed individually.

Dependent variables. The main measure of adaptation success was trial completion time. Adaptation and stability of control were assessed using four goal variable measures. *Rise time* was the elapsed time from the start of a trial until the process first reached the target area, where all four goal variables were in their respective set point regions. The goal variables sometimes went beyond, and oscillated around, the target regions before reaching steady state. *Oscillation time* was the time to stabilize the goal variables after they had reached the target regions but before steady state was achieved. During this period the *normalized maximum deviation* was the maximum value by which the goal variables exceeded the target regions, divided by the target value. *Number of oscillations* was the number of times the goal variables crossed above and below the target regions during this period. *Steady state time* was the period when all goal variables were in the target regions for 5 consecutive min. To be conservative, only the maximum value within each trial for the set of four goal variables was used in the analysis of each measure.

Level of control in the abstraction hierarchy was assessed quantitatively by using a set of novel multidimensional measures of variance (Yu et al., 2002). Each level of the AH (i.e., functional purpose, abstract function, generalized function, and physical function) is a different representation of the same work domain, so each level provides a different state space of variables for describing the quantitative behavior of the participant-environment system (i.e., goal variables, mass and energy variables, flow rate and heat transfer rate variables, and component settings, respectively).

The quantitative behavior during each trial can be plotted as a trajectory over time. For example, the state of the goal variables over time for one trial constitutes one trajectory. However, because each level of the AH defines a different state space, the same trial will be revealed as a different trajectory at each level of the AH. For each block of trials for each participant, we obtained a series of trajectories in the same state space, one for each trial. It was then possible to calculate the multidimensional variance of these trajectories. This variance is a quantitative measure of consistency (or, conversely, variability) at a particular level of the AH. For example, if the path that the goal variables take is exactly the same for each trial, then the variance would be zero. However, if the path that the goal variables take is wildly different from trial to trial, then the variance would be large.

Because each level of the abstraction hierarchy represents a different state space for the same work domain, the same block of trials is represented as a different set of trajectories at each level of the AH. Thus it is possible for the same participant to exhibit high variance at one level and low variance at another. Such a pattern of results allows us to make inferences about level of control; low variance at a functional level of the AH would be evidence of higher-level control, whereas low variance at a physical level of the AH would be evidence of lower-level control.

A knowledge elicitation measure, a control recipe (Irmer & Reason, 1991), was also administered. Participants were asked to write down a set of instructions describing how they controlled the process. These instructions were to be sufficiently detailed to allow someone who had never seen the process before to control it in the same manner as did the participant.

Data analysis. Data were analyzed using several statistical software packages: Matlab (The Mathworks, Inc., Natick, MA), Excel (Microsoft Corp., Redmond, WA), and SAS (Cary, NC). Wherever meaningful, statistical significance was assessed using 95% confidence intervals because these provide the same information as traditional null hypothesis tests (e.g., analysis of variance) as well as additional important information (see Loftus, 1993, for details). If the confidence interval bars for two means do not overlap, then the difference between means is statistically significant at the $p < .05$ level.

Results

Trial completion times. Figure 3a shows the results for normal trials in Block 3 and perturbation trials in Block 4 by interface group. For both blocks of trials the P+F group had significantly faster trial completion times than did the P group. Within each group the perturbation trials in Block 4 were significantly slower than the normal trials in Block 3.

Figure 3b shows the differences between the two interface groups in terms of percentage change in trial completion times between matched trials from Block 3 and Block 4 (e.g., Trial 61 was compared with Trial 41). The completion times for the P+F group increased by 24% on average, whereas the completion times for the P group increased by 43% on average; this difference between the two groups was statistically significant. These results show that the P+F group was better able to adapt to the perturbations than was the P group.

Figure 4a illustrates how trial completion times varied for each interface group as a function of the magnitude of the perturbation. The completion times increased as the time constant multiple increased for both groups. There were statistically significant differences between the two groups at the multiples of 0.1, 2.5, 5.0, and 10.0×. Applying linear regression, the intercept, slope, and standard deviations of the P group were larger than those of the P+F group; the R^2 value was smaller for the P group compared with the P+F group.

Figure 4b shows the results of regression analyses performed on individual participant data; one regression analysis was conducted for each participant, followed by averaging of the slopes and intercepts. The slopes were significantly lower for the P+F group compared with the P group; the intercepts were lower for the P+F compared with the P group, but not significantly so. These results also show that the P+F group was better able to adapt to the perturbations than was the P group.

Goal variable measures. Goal variable measures provide an assessment of success in adapting to change as well as stability of control. Figure 5 shows the results comparing Block 3 (normal trials) and Block 4 (perturbation trials). The mean rise times within each block were roughly the same for the two groups. However, the perturbation trials in Block 4 were significantly slower than the normal trials in Block 3 (99% slower for P and 93% slower for P+F participants, from matched trials in Blocks 3 and 4). This result suggests that the negative impact of the perturbations on rise times did not differ for the two groups.

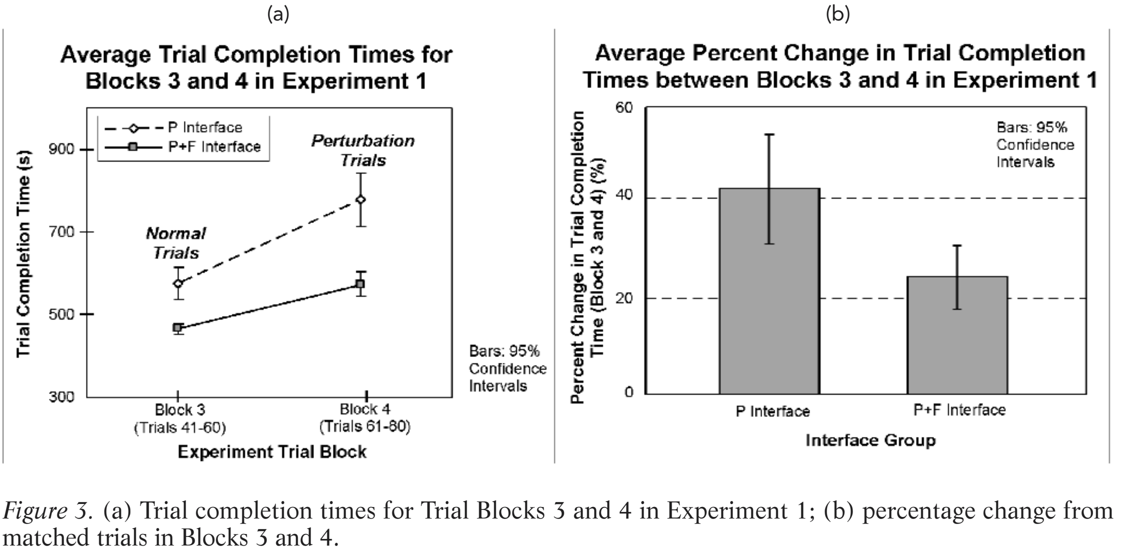

Figure 3. (a) Trial completion times for Trial Blocks 3 and 4 in Experiment 1; (b) percentage change from matched trials in Blocks 3 and 4.

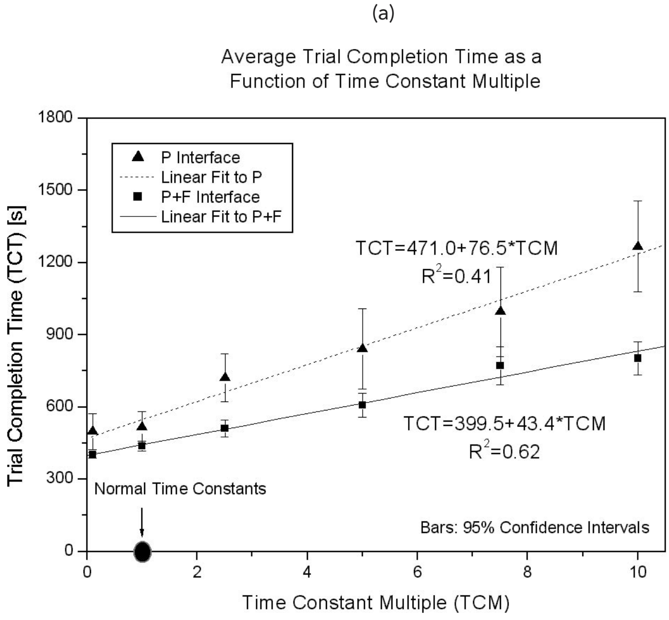

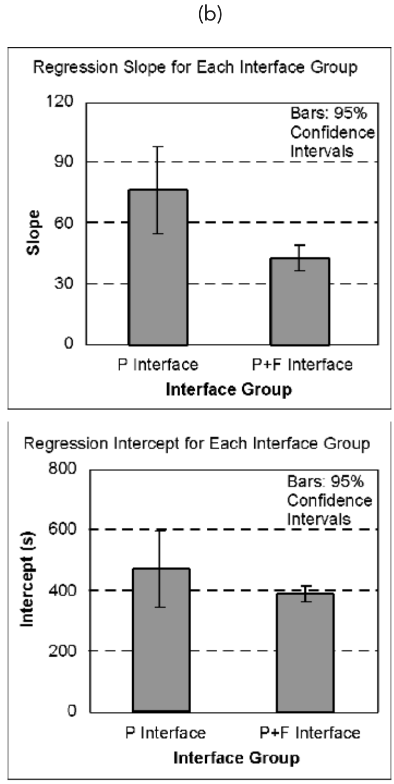

Figure 4. (a) Effect of specific global perturbations on trial completion times in Experiment 1; (b) results of regression analyses for individual participants.

The mean oscillation times in each block were significantly slower for the P group. Both groups experienced significantly slower oscillation times between Blocks 3 and 4 (112% slower for P and 58% slower for P+F participants, on average). The difference between the groups was much higher in Block 4 than in Block 3. This result suggests that the perturbations had a proportionately larger negative impact on stability for the P group than for the P+F group, based on oscillation times.

The mean number of oscillations per trial was significantly lower for the P+F group. Both groups experienced a small, nonsignificant decrease between Blocks 3 and 4 (6% decrease for P and 10% decrease for P+F participants, on average). This result suggests that the perturbations did not affect stability for either group, based on the mean number of oscillations per trial.

The mean normalized maximum deviations from the target regions was significantly lower for the P group in Block 3. However, in Block 4 this difference was not significant, with an increase for the P group and a decrease for the P+F group (51% increase for P and 26% decrease for P+F participants, on average). This result suggests that the perturbations had a larger negative impact on stability for the P group compared with the P+F group, based on the mean normalized maximum deviations from the target regions.

In sum, the P+F group exhibited more stable control in the face of perturbations than did the P group, in terms of both oscillation times and maximum deviations from the target regions. The perturbations had a similar negative impact on rise times for both groups. The perturbations had no significant impact on average number of oscillations per trial.

AH trajectory variability. The variability of state trajectories at each level of the abstraction hierarchy was calculated for each block of trials. These values were compared between participants at each level of the AH. Based on previous research (Yu et al., 2002), two AH levels – abstract function and physical function – were chosen for analysis to assess whether the participant group was engaged in higher-level or lower-level control, respectively.

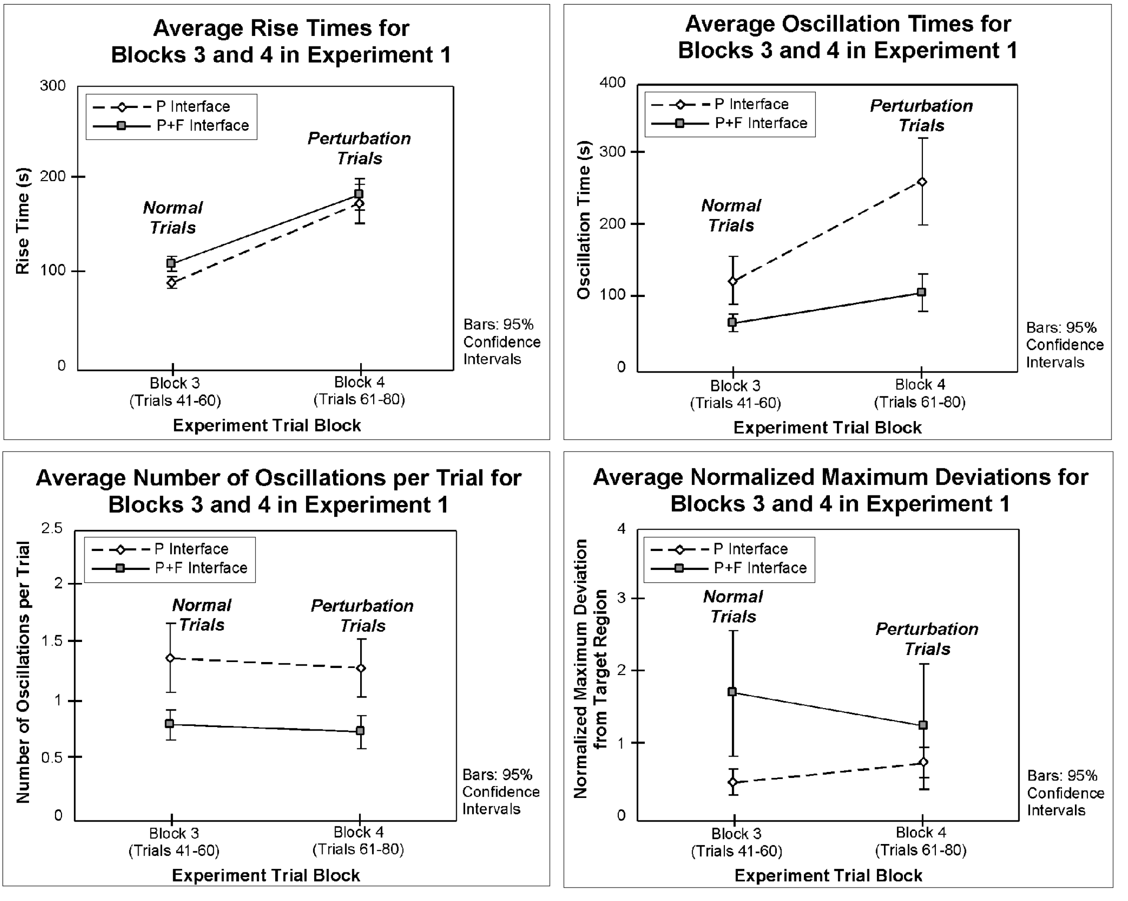

Figure 5. Goal variable measures for Experiment 1.

Figure 6 shows the variability distributions for participants in each interface group for the four blocks of the experiment. The horizontal axis represents the ranges in variability values; the vertical axis represents the number of participants who had a variability value in that range.

The left column of Figure 6 displays the results of the variability analyses for the abstract function AH level. As participants gained experience, both groups exhibited less variability at this level (from Blocks 1 to 3). In Block 4 the variability increased for both groups compared with Block 3. However, the P+F group had more participants with relatively lower variability than did the P group, although this difference was not statistically significant (i.e., 6/8 P+F and 2/8 P participants had values below the median for both groups; nonparametric Fisher test for a small sample size: $p = .13$ exact, two-tailed). This result suggests that the P+F participants may have been engaged in more higher-level control compared with the P participants.

The right column in Figure 6 shows the results of the variability analyses for the physical function AH level. As participants gained experience, both groups exhibited less variability at this level (from Blocks 1 to 3); by Block 3 the distributions were identical. In Block 4, however, the P group had more participants with relatively lower variability than did the P+F group, although this difference was also not statistically significant (i.e., 6/8 P and 2/8 P+F participants had values below the median for both groups; nonparametric Fisher test for a small sample size: $p = .13$ exact, two-tailed). This result suggests that P participants may have been engaged in more lower-level control compared with the P+F group.

Control recipes. Control recipes provided a subjective assessment of level of control. After

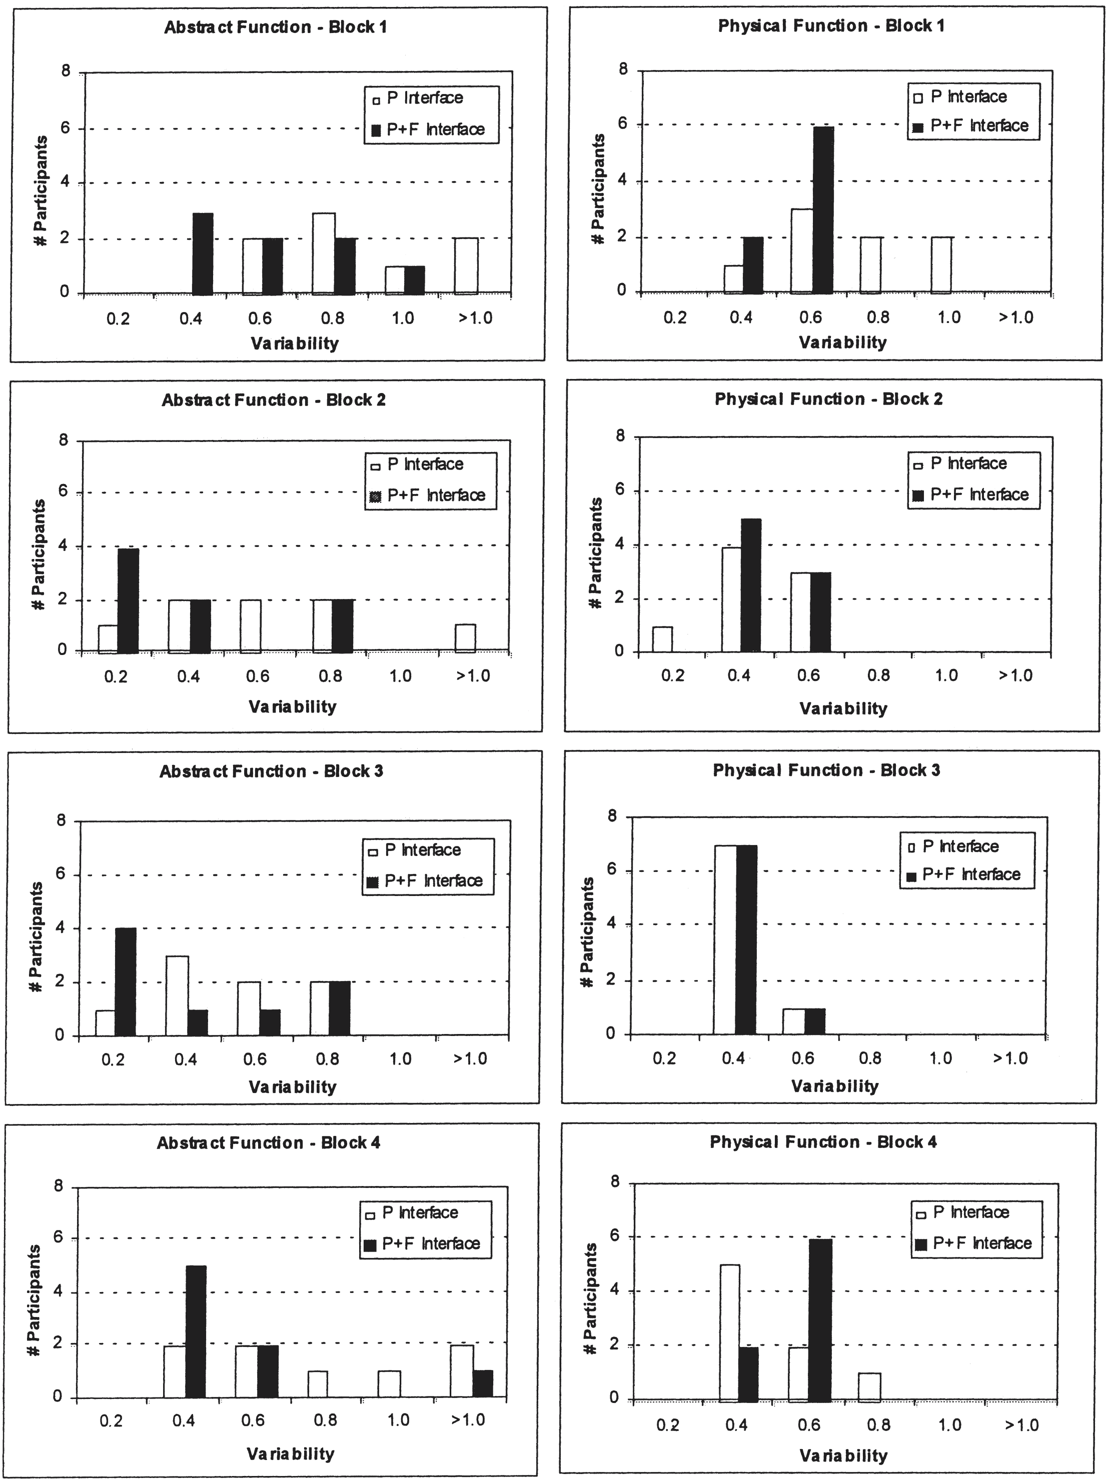

Figure 6. Trajectory variability distributions in Experiment 1 for the AH levels of abstract function (left column) and physical function (right column).

Block 3 all participants reported only one way of doing the task because they had not yet experienced any perturbations. After the perturbation block, however, 4/8 P and 8/8 P+F participants reported adjusting their actions as a function of the simulation dynamics. An example of this context-conditioned variability is shown in Table 1 for a P+F participant.

The two groups seemed to differ in the degree to which they exhibited this pattern. Three categories of control recipes were reported overall: normal/regular, slow (higher time constant multiples), and fast (lowest time constant multiple). Of the P participants, 5/8 reported only one way of controlling the simulation after Block 4 (just as they had after Block 3) compared with 0/8 of the P+F participants. Of the P participants, 3/8 reported two conditions for controlling the simulation (e.g., normal and slow) compared with 4/8 of the P+F participants. Finally, 0/8 P participants reported three conditions for controlling the simulation (i.e., normal, slow, and fast) compared with 4/8 P+F participants. These results suggest that the P+F participants reported greater signs of context-conditioned variability and adjusted their actions to achieve the goals in a manner that was tailored to the magnitude of the perturbations. This finding provides subjective evidence for greater higher-level control for the P+F group compared with the P group.

Discussion

The results are in close agreement with the rationale behind EID. The P+F group showed signs of better adaptation to the unpredictable perturbations than did the P group, as evidenced by faster and more robust trial completion times as well as more stable control of the goal variables. Also, the P+F group showed more signs of higher-level control than did the P group, as evidenced by lower trajectory variability at the abstract function AH level and a greater number of participants reporting context-conditioned control recipes. Conversely, the P group showed more signs of lower-level control than did the P+F group, as evidenced by lower trajectory variability at the physical function AH level and a greater number of participants reporting context-independent control recipes. However, these differences were not statistically significant.

Interestingly, 2/8 P participants were also able to do relatively well in adapting to changes in component dynamics, relative to all participants. Their control recipes suggest that they developed analytical heuristics for controlling the process (e.g., Table 2). Because the goal temperatures remained the same for all trials, the steady-state heater setting could be derived based on the outflow demand goals by two simple equations (i.e., HTR1 = D1; HTR2 = 1/3D2). The two proficient P participants discovered

TABLE 1: Part of a Control Recipe for a P+F Participant after the Perturbation Phase

Fast reactions:

- Turn on valves in the same manner as discussed before: Mo ratio of the two tanks determines sum of VA1+VB1/VA2+VB2, and respective value of Mo determines VA, VB level.
- Turn on the heater: The larger T1 or T2 appears, the higher the respective HTR should be at; however, neither of them should be above 5 initially, avoiding T1 or T2 overheated.

Regular reactions:

- Turn on valves and heaters pretty much the same principle as above. But the heater of T1 (which requires larger T) should be set between 6–7, in order to increase the temperature faster.

Slow reactions:

- As soon as the slow reaction is observed after turning on the valves and pumps, turn every control (i.e., valves and HTRs) to its extreme value. This can speed up the magnitudes of change.
- After Mi's reach the desired value (i.e., the value above green region in Mo), turn down the valves switch to the readings of FVA, and FA's etc. so that Mi is stabilized. Then turn VO's to extreme value 'til MO reaches the green region.
- When T1/T2 reach green regions, completely shut down the heaters, until the slope of E1/E2 returns to a vertical line and then turn up the heaters to the heat flow reading below it.

TABLE 2: Part of a Control Recipe for a Proficient P Participant after the Perturbation Phase

- Calculate steady-state temperature: VO1 × (T1 – To) = VO1 × 30; VO2 × (T2 – To) = VO2 × 10 – of total heat required divide by 30 (i.e., HTR1 = VO1 × 30/30 = VO1, HTR2 = VO2 × 10/30 = VO2/3 is the final steady state setting of the heaters).
- When T2 = 15° turn T2 down to the HTR2 steady state setting to slow down the heating process. Similarly, when T1 ~ 33 turn down T1 to steady state setting.
- If T1 and T2 are responding slowly to change in HTR1 and HTR2, turn HTR1 and HTR2 to T1ss (HTR1 = T1goal), T2ss (HTR2 = T2goal/3). When T1 is at 33° and T2 is at 15° to let the system catch up to what your settings were previously.
- If T1 and T2 cease to rise, increase it by 1–2 units (depends on how close to goal T1 and T2 are).
- If T1 and T2 overshoot a lot, decrease HTR1 and HTR2 to zero and let it cool down till 1 unit above goal and increase it to Tss.
- Basically not much different from ordinary operation. Just remember that whatever we do, the result will appear a bit later. So turn down or up the HTR in advance before it reaches the goal state.

these heuristics. This result suggests that the absence of higher-level functional information may, to some extent, be compensated for by heuristics that map to higher-level functions. This compensatory strategy is possible only if the cognitive demands associated with analytical derivation are psychologically feasible, as they were in the case of the straightforward heuristics reported by the two proficient P participants.

Why was the P+F group better able to rely on higher-level control and thus adapt to novelty and change? According to EID, this advantage has its basis in perceptual coupling to emergent feature graphics. The P+F group could see, rather than analytically derive, the state of the higher-level functions by relying on the emergent features shown in Figure 2. As a result, they were better able to engage in higher-level control and thereby to adapt their actions to the unexpected changes in process dynamics. This perceptually based explanation for the advantage of the P+F interface was tested in the following experiment.

EXPERIMENT 2

The purpose of this experiment was to investigate the impact of changes in interface form on success in adaptation and on level of control with the P+F interface only. The usual mapping between information content and visual form is illustrated in the middle column of Figure 7 ("invariant scale limits"). For the higher-level functional information the mass balance constraint is mapped onto an emergent feature graphic showing the relationship between mass input (MI1), mass output (MO1), and volume (V1). When input equals output, the vertical line connecting MI1 and MO1 is perpendicular, indicating that the volume should be stable. A similar type of graphic is used to represent the energy balance (see Figure 2). For the lower-level physical information the state of each component is mapped linearly onto the spatial position of a triangular pointer.

To investigate the relative contributions of the emergent feature graphics to successful adaptation and level of control, we introduced two types of form perturbations, one disrupting higher-level emergent features and another disrupting lower-level physical forms. According to EID, perturbations of higher-level forms should have a stronger negative impact on performance than should perturbations of lower-level forms because perceptual coupling to emergent feature graphics is said to be the basis for higher-level control and adaptation success.

Method

The method used was similar to that used in Experiment 1. Again, there were four blocks of 20 trials each, with the first three blocks being normal trials and the last block being perturbation trials. Each of the 16 participants in this experiment (not the same as those in Experiment 1) operated the simulation for approximately 25 hr in total, again resulting in approximately 400 participant hours of data collection.

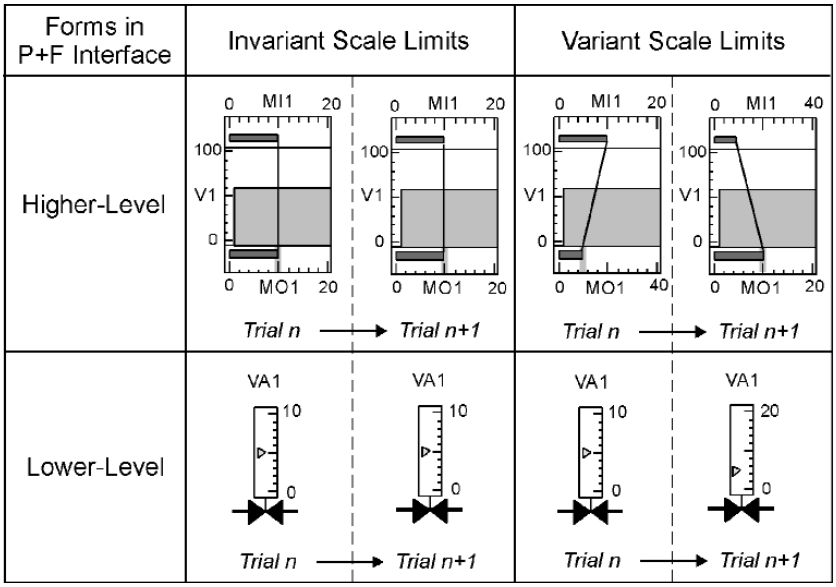

Figure 7. Examples of invariant and variant scale limit changes for higher-level (HL) and lower-level (LL) groups. All examples are based on the same work domain conditions in steady state.

As shown in Figure 8, the only difference in this study was the independent variables being manipulated. A repeated-measures mixed design was adopted, with interface form perturbations (lower level, or LL, and higher level, or HL) as a between-participants variable and trial type (normal and perturbation) as a within-participants variable. For the first three blocks of 20 trials the two groups used the same (unperturbed) P+F interface with invariant (i.e., consistent) LL and HL forms. During Block 4 for the LL group, the scale values of the low-level forms

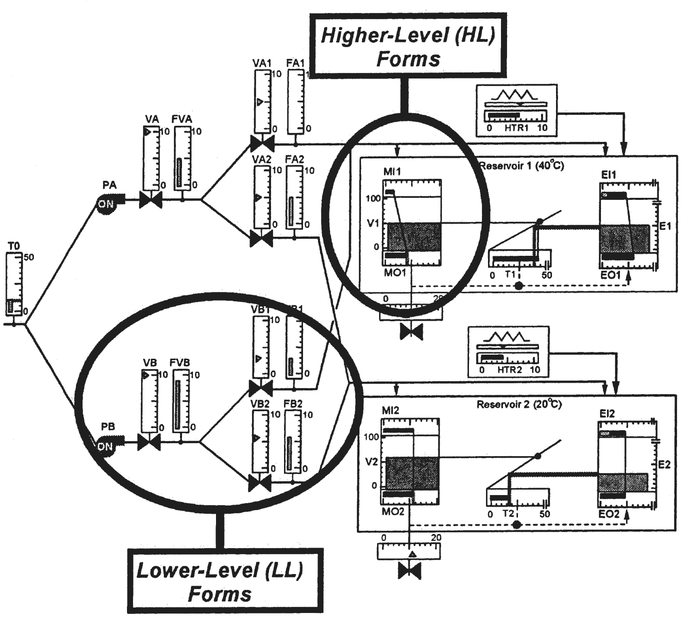

Group	
	HL - Perturbations are variant scales across trials for higher-level forms
	LL - Perturbations are variant scales across trials for lower-level forms

Interface Forms	Scale Mapping Across Trials	
	Block 3	Block 4
Higher-Level	HL: Invariant LL: Invariant	HL: **Variant** LL: Invariant
Lower-Level	HL: Invariant LL: Invariant	HL: Invariant LL: **Variant**

Figure 8. Design of Experiment 2.

varied in unpredictable ways from trial to trial. No changes were made to the higher-level forms. Conversely, during Block 4 for the HL group, the scale values of the higher-level forms were varied in unpredictable ways from trial to trial. No changes were made to the lower-level forms. These manipulations changed the form of the P+F interface but not the underlying simulation dynamics. Also, the magnitude of the changes in forms was physically the same for the LL and the HL group (see Hajdukiewicz, 2001, for details).

Figure 7 illustrates the nature of the two types of form perturbations in more detail. During the first three blocks the graphics behaved normally, as described earlier. During Block 4, however, the emergent feature was destroyed for the HL group. The mass (and energy) inflow and outflow scales were no longer the same within a trial, and they were varied unpredictably across trials. If P+F participants controlled the process primarily by perceptual coupling to emergent features, then this HL manipulation should have severely degraded performance.

As shown in Figure 7, the manipulation for the LL group was analogous. During the first three blocks the scale limits on the components were consistent across trials, thereby providing a predictable basis for being perceptually coupled to lower-level information. During Block 4, however, the scale limits on the individual components were varied unpredictably across trials. If P+F participants controlled the process primarily by perceptual coupling to emergent features, then the LL manipulation should not have degraded performance as much as the HL manipulation.

Results

Trial completion time. Figure 9a shows the results for normal trials in Block 3 and perturbation trials in Block 4, by group. For Block 3 there was no statistically significant difference in completion times between the HL and LL groups. This was as expected, given that the treatments were identical to this point. However, with the form perturbation trials in Block 4, the HL group exhibited significantly slower completion times than did the LL group. For both groups the perturbation trials in Block 4 were significantly slower than the normal trials in Block 3.

Figure 9b displays the differences between the two interface groups in terms of percentage change in trial completion times between matched trials from Block 3 and Block 4. The times for the HL group increased by 33% on average, whereas the completion times for the LL group increased by 14% on average, a statistically significant difference. These results show that higher-level form perturbations had a proportionately larger negative effect on com-

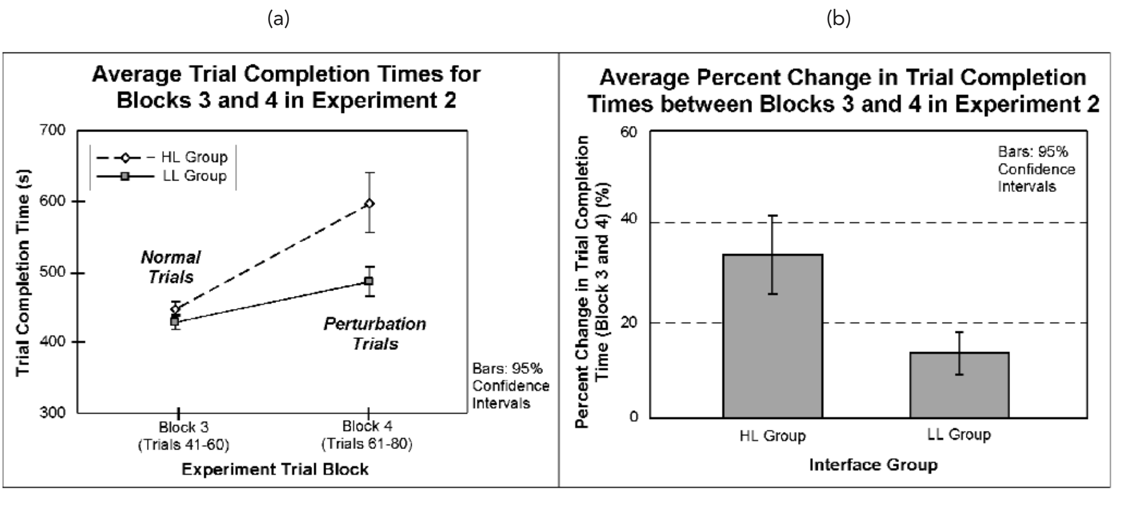

Figure 9. (a) Trial completion times for Trial Blocks 3 and 4 in Experiment 2; (b) percentage change from matched trials in Blocks 3 and 4.

pletion times than did lower-level form perturbations, suggesting that perceptual coupling to emergent feature graphics plays a key role in control with the P+F interface.

Goal variable measures. Goal variable measures provide another set of indicators for success in adapting to change, as well as an assessment of stability. Figure 10 shows the results comparing Block 3 (normal trials) and Block 4 (perturbation trials). The mean rise times in Block 3 were roughly the same for the two groups; in Block 4 there was a larger but not statistically significant difference between the groups. However, within each group the perturbation trials in Block 4 were significantly slower than the normal trials in Block 3 (41% slower for HL and 27% slower for LL participants, from matched trials in Blocks 3 and 4). This result suggests that the form perturbations negatively influenced rise times for both groups.

The mean oscillation times in Block 3 were roughly the same for the two groups. The difference between the groups was larger in Block 4 than in Block 3 and was statistically significant. Both groups experienced significantly slower oscillation times in Block 4, but the impact was far greater on the HL group (228% slower for HL and 110% slower for LL participants, on average). This result suggests that the form perturbations had a proportionately larger negative impact on stability for the HL group than for the LL group, in terms of oscillation times.

The mean number of oscillations per trial in Block 3 was similar for both groups. Both groups experienced a statistically significant increase between Blocks 3 and 4, but the impact was again far greater on the HL group (190% increase for HL and 51% increase for LL participants, on average). This result suggests that

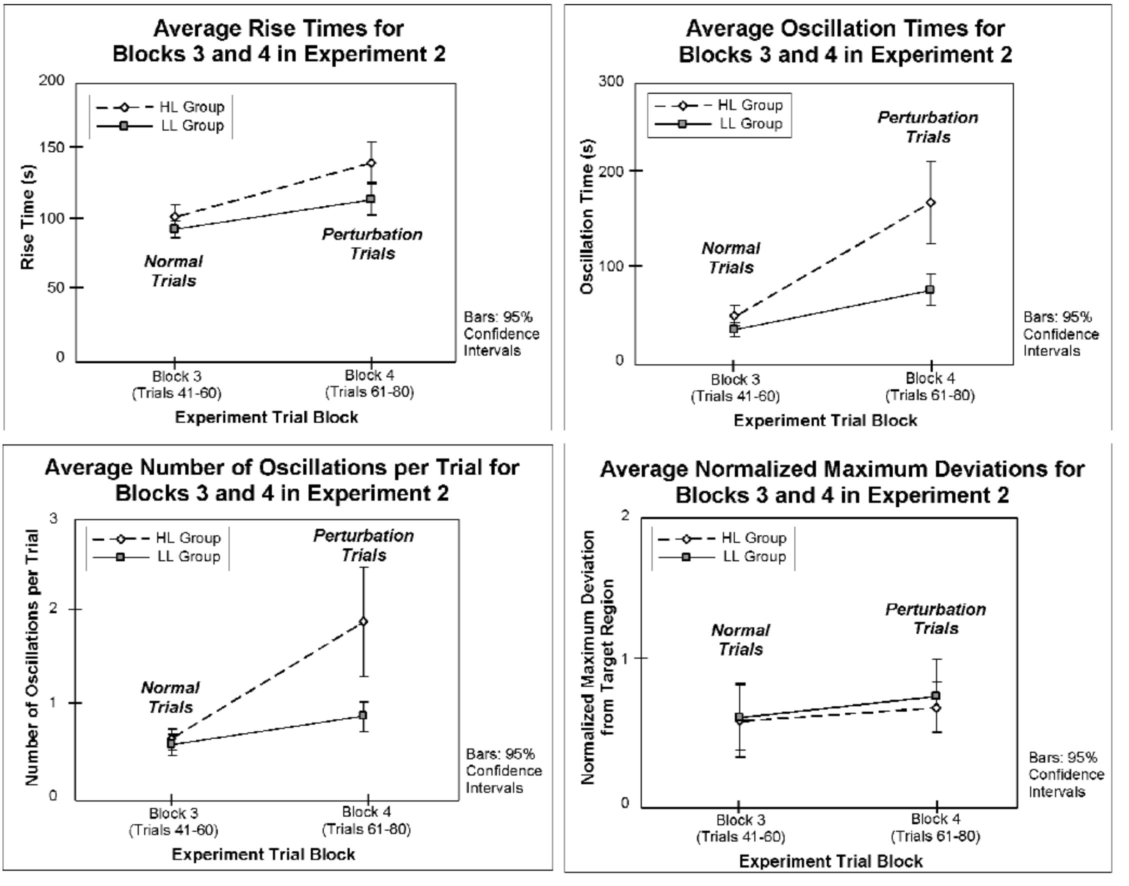

Figure 10. Goal variable measures for Experiment 2.

the form perturbations had a proportionately larger negative effect on stability for the HL group than for the LL group, in terms of number of oscillations per trial.

The mean normalized maximum deviations from the target regions were roughly the same for each group. There was a small, nonsignificant increase for both groups between Blocks 3 and 4 (15% increase for HL and 24% increase for LL participants, on average). This result suggests that the form perturbations had very little impact on stability for both groups, in terms of average normalized maximum deviations from the target regions.

In sum, the form perturbations had a significantly larger negative effect on stability for the HL group than for the LL group, in terms of both oscillations times and number of oscillations per trial. The perturbations had a negative impact on rise times for both groups. Finally, there was little impact of the form perturbations on the mean maximum deviations from the target regions.

AH trajectory variability. In Figure 11, left column, are shown the results for the abstract function AH level. As participants gained experience, both groups exhibited less variability at this level (from Blocks 1 to 3). In Block 4 the variability increased for both groups compared with Block 3; however, the variability seems proportionately higher for the HL group. The HL group had more participants with relatively higher variability than did the LL, although the difference was not statistically significant (i.e., 2/8 LL and 6/8 HL participants had values above the median for both groups; nonparametric Fisher test for a small sample size: p = .13 exact, two-tailed). This result suggests that form perturbations mapping to the higher-level variables of the AH may result in less use of higher-level control, whereas form perturbations mapping to the lower-level variables of the AH may result in continued use of higher-level control.

In the right column of Figure 11 are the results of the variability analyses for the physical function AH level. As participants gained experience, both groups exhibited less variability at this level (from Blocks 1 to 3). In Block 4 the variability distribution remained the same for the HL group but increased substantially for the LL group. All HL participants had lower variability than did the LL participants (nonparametric Fisher test for a small sample size: p = .0002 exact, two-tailed). This result suggests that form perturbations mapping to the lower-level variables of the abstraction hierarchy result in less use of lower-level control, whereas form perturbations mapping to the higher-level variables of the AH result in continued use of lower-level control.

Control recipes. After Block 3 all but one of the participants (8/8 HL and 7/8 LL) reported using some higher-level forms to control the process (e.g., balancing mass in and mass out with a vertical line). After the perturbation phase 7/8 LL participants still reported using some higher-level forms (e.g., "adjust the heater setting so energy in = energy out"), whereas only 3/8 HL participants reported doing so. Of the HL participants, 5/8 reported either using components in specific ways (e.g., heuristics) or using trial and error to achieve the goals. These results suggest that the higher-level form perturbations may have resulted in less use of the higher-level forms, making it more difficult for HL participants to achieve higher-level control.

Discussion

The experimental results are again in close agreement with the rationale behind EID. The form perturbations had a larger negative impact on completion times and stability for the HL group compared with the LL group, suggesting that perceptual coupling to emergent feature graphics plays a key role in controlling the process with the P+F interface. There was also some evidence linking coupling to emergent feature graphics with level of control. In the AH variability analyses and the control recipes, the LL group exhibited signs of relying more on higher-level control than did the HL group. Conversely, the HL group exhibited signs of relying more on lower-level control than did the LL group.

As in Experiment 1, there were some exceptions to the general pattern of results. Of the HL participants, 2/8 did not appear to be affected by the higher-level form perturbations. In the control recipes of these 2 participants, both reported using analytical means of deriving the

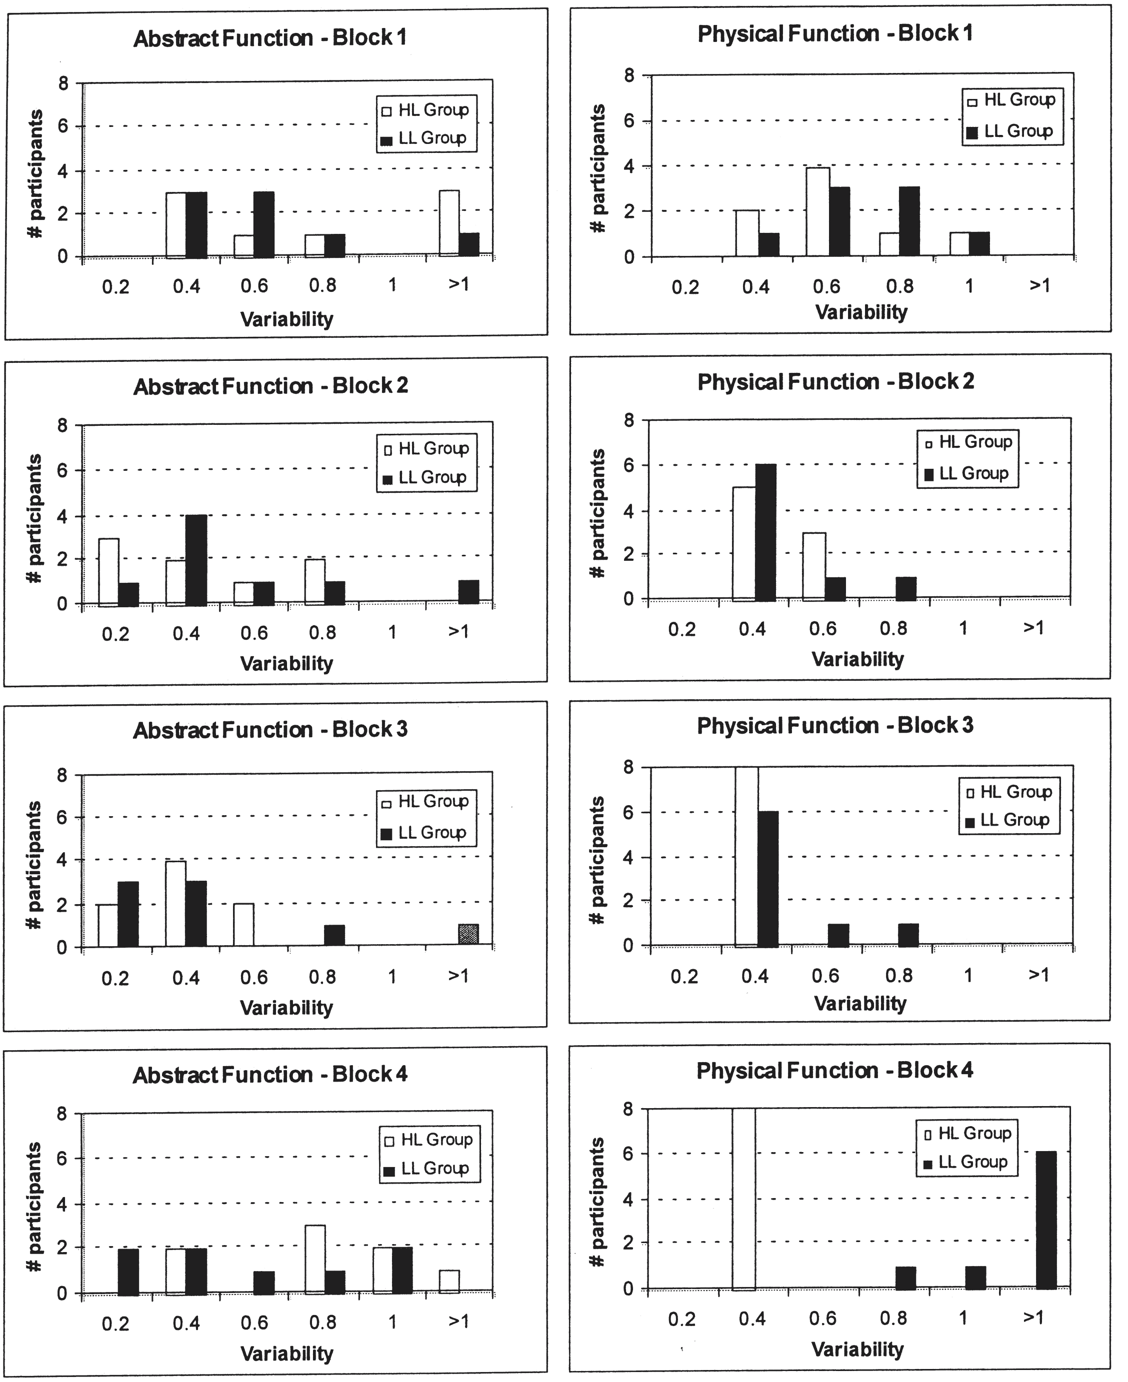

Figure 11. Trajectory variability distributions in Experiment 2 for the AH levels of abstract function (left column) and physical function (right column).

functional information that had been perturbed by the form manipulations. Table 3 provides one example that demonstrates the reported use of heuristics (in italics). This result suggests that the absence of higher-level form invariants may, to some extent, be compensated for by heuristics that map to higher-level functions. This compensatory strategy is possible only if

TABLE 3: Part of a Control Recipe for a Proficient HL Participant after the Perturbation Phase

- *Increase HTR2 by one unit for every 3 units of MO2, keeping in mind point 1 (this will be your set point for HTR2).*
- *Increase HTR1 by one unit for every 1 unit of MO1, keeping in mind point 1 (this will be your set point for HTR1).*
- Set HTR1 and HTR2 to max until T1 and T2 approach the desired temperature band.
- When T1 and T2 are within 1 unit of the desired temperature band, set HTR1 and HTR2 to their respective settings.
- Set VO1 equal to MO1, keeping 1 in mind.
- Set VO2 equal to MO2, keeping 1 in mind.
- Adjust HTR1 and HTR2 around their respective set points in order to stabilize T1 and T2, within their respective desired output temperature bands.

the cognitive demands associated with analytical derivation are psychologically feasible, as they were in the case of the simple heuristics reported by the 2 proficient HL participants.

CONCLUSIONS

Designing to support human adaptation to novelty and change is necessary and possible. EID provides a principled way of achieving this goal. Displaying functional information can result in greater adaptation because participants can become perceptually coupled to emergent feature graphics, thereby exhibiting higher-level control. In some cases it is possible to achieve higher-level control in the absence of these preconditions, although this requires knowledge and cognitive resources. This path represents the more difficult road to successful adaptation.

Several issues could be pursued to address the limitations of this research. First, it would be useful to replicate these findings in other application domains and in larger-scale test beds that are more representative of those found in industry. Second, it would be interesting to investigate the impact of other types of perturbations to see if the same results hold. Third, the scope of EID should be extended, both theoretically and empirically, to address adaptation at the team and social-organizational levels.

Finally, the generalizability of the findings obtained here should be investigated. The results in this article are surprisingly similar to those obtained in the motor control literature. For example, Kelso, Tuller, Vatikiotis-Bateson, and Fowler (1984) investigated a type of problem very different from the one studied here – namely, the motion of jaw muscles during speech. There seems to be very little reason, a priori, to believe that human-computer interaction with a thermal-hydraulic process simulation would have anything at all to do with unmediated human speech production. However, Kelso et al. found that people were able to demonstrate robust control in the face of novelty and change and that this control was achieved by keeping higher-level goals constant and by modifying lower-level processes in a way that was specific to both the task goal and the nature of the perturbation.

This strong similarity between seemingly disparate phenomena suggests that mechanical, biological, human-machine, and perhaps even social adaptation may be different manifestations of the same underlying phenomenon. A unified theory encompassing these phenomena would represent a significant contribution to applied research in cognitive engineering because it would provide a theoretical basis for shaping adaptation in complex systems. At the same time, the theory would also represent a significant contribution to multidisciplinary basic research in the technical, human, and social sciences because it would provide a common basis for understanding adaptation in its many disparate guises.

ACKNOWLEDGMENTS

The writing of this paper was sponsored by the Jerome Clarke Hunsaker Distinguished Visiting Professorship at MIT and by a research grant and a fellowship from the Natural Sciences and Engineering Research Council of Canada. We thank Pascal van Lieshout, Susan McCahan, Tom Stoffregen, the editor, and two reviewers for their comments.

REFERENCES

Blum, B. I. (1996). *Beyond programming: To a new era of design.* New York: Oxford University Press.

Buttigieg, M. A., & Sanderson, P. M. (1991). Emergent features in visual display design for two types of failure detection tasks. *Human Factors, 33,* 631–651.

Christoffersen, K., Hunter, C. N., & Vicente, K. J. (1997). A longitudinal study of the effects of ecological interface design on fault management performance. *International Journal of Cognitive Ergonomics, 1,* 1–24.

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 doctoral dissertation, University of Toronto, Department of Mechanical and Industrial Engineering.

Ham, D.-H., & Yoon, W. C. (2001). The effects of presenting functionally abstracted information in fault diagnosis tasks. *Reliability Engineering and System Safety, 73,* 103–119.

Hirschhorn, L. (1984). *Beyond mechanization: Work and technology in a postindustrial age.* Cambridge: MIT Press.

Irmer, C., & Reason, J. T. (1991). Early learning in simulated forest fire-fighting. In *Simulations, evaluations, and models: Proceedings of the 4th Models of Human Activities in Work Context Workshop* (pp. 1–20). Roskilde, Denmark: Risø National Laboratory.

Kelso, J. A. S., Tuller, B., Vatikiotis-Bateson, E., & Fowler, C. A. (1984). Functionally specific articulatory cooperation following jaw perturbations during speech: Evidence for coordinative structures. *Journal of Experimental Psychology: Human Perception and Performance, 10,* 812–832.

Loftus, G. R. (1993). A picture is worth a thousand *p* values: On the irrelevance of hypothesis testing in the microcomputer age. *Behavior Research Methods, Instruments, and Computers, 25,* 250–256.

Nickerson, R. S. (1995). *Emerging needs and opportunities for human factors research.* Washington, DC: National Academy Press.

Pask, G., & Scott, B. C. (1972). Learning styles and individual competence. *International Journal of Man-Machine Studies, 4,* 217–253.

Pawlak, W. S., & Vicente, K. J. (1996). Inducing effective operator control through ecological interface design. *International Journal of Human-Computer Studies, 44,* 653–688.

Rasmussen, J. (1985). The role of hierarchical knowledge representation in decisionmaking and system management. *IEEE Transactions on Systems, Man and Cybernetics, SMC-15,* 234–243.

Rasmussen, J., Pejtersen, A. M., & Goodstein, L. P. (1994). *Cognitive systems engineering.* New York: Wiley.

Reising, D. V. C., & Sanderson, P. M. (2000). Testing the impact of instrumentation location and reliability on ecological interface design: Fault diagnosis performance. In *Proceedings of the XIVth Triennial Congress of the International Ergonomics Association and 44th Annual Meeting of the Human Factors and Ergonomics Society* (Vol. 3, pp. 591–594). Santa Monica, CA: Human Factors and Ergonomics Society.

Shepherd, A. (1993). An approach to information requirements specification for process control tasks. *Ergonomics, 36,* 1425–1437.

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.

Turvey, M. T., Fitch, H. L., & Tuller, B. (1982). The Bernstein perspective: I. The problems of degrees of freedom and context-conditioned variability. In J. A. S. Kelso (Ed.), *Human motor behavior: An introduction* (pp. 239–252). Hillsdale, NJ: Erlbaum.

Vicente, K. J. (1999). *Cognitive work analysis: Toward safe, productive, and healthy computer-based work.* Mahwah, NJ: Erlbaum.

Vicente, K. J. (2000). Toward Jeffersonian research programmes in ergonomics science. *Theoretical Issues in Ergonomics Science, 1,* 93–113.

Vicente, K. J. (2002). Ecological interface design: Progress and challenges. *Human Factors, 44,* 62–78.

Vicente, K. J., & Rasmussen, J. (1990). The ecology of human-machine systems: II. Mediating “direct perception” in complex work domains. *Ecological Psychology, 2,* 207–250.

Vicente, K. J., & Rasmussen, J. (1992). Ecological interface design: Theoretical foundations. *IEEE Transactions on Systems, Man and Cybernetics, SMC-22,* 589–606.

Yu, X., Lau, E., Vicente, K. J., & Carter, M. W. (2002). Toward theory-driven, quantitative performance measurement in ergonomics science: The abstraction hierarchy as a framework for data analysis. *Theoretical Issues in Ergonomics Science, 3,* 124–142.

John R. Hajdukiewicz received a Ph.D. in mechanical and industrial engineering in 2001 at the University of Toronto. He is a senior research scientist at Honeywell Laboratories in Minneapolis, Minnesota.

Kim J. Vicente received a Ph.D. in mechanical engineering in 1991 from the University of Illinois at Urbana-Champaign. He is the Jerome Clarke Hunsaker Distinguished Visiting Professor of Aerospace Information Engineering at MIT and professor of mechanical and industrial engineering, biomaterials and biomedical engineering, computer science, and electrical and computer engineering at the University of Toronto.

Date received: July 17, 2001
Date accepted: October 25, 2002