2014 IEEE International Conference on Systems, Man, and Cybernetics
October 5-8, 2014, San Diego, CA, USA

A Strategy-Based Ecological(?) Display for Time-Series Structural Change Diagnosis

Antony Hilliard, Greg A. Jamieson
Department of Mechanical & Industrial Engineering
University of Toronto
Toronto, Canada
anthill@mie.utoronto.ca

***Abstract*—Ecological Interface Design (EID) is usually associated with applying Work Domain Analysis and developing customized control tools for specific, stable systems governed by causal work domain constraints. However, for domains with weaker, intentional, less specific constraints, EID based on analysis of cognitive strategies has been proposed, and demonstrated using metaphorical iconic forms. EID methods to find invariant strategies in intentional work domains are less developed. Such methods may be a way to generalize ecological interfaces beyond specific instances of complex systems. An example developed for an energy efficiency monitoring application is presented and discussed.**

Keywords—displays; ecological interface design; cognitive strategies; diagnosis; search; parameter estimation; recursive estimates; change-point-detection; energy; efficiency

I. Ecological Interface Design Principles

Ecological Interface Design (EID) was introduced 25 years ago as a set of design guidelines and principles to reduce human error in process control domains [1]. The core principle is that interfaces should help humans perceive and control complex work environments in terms of useful regularities at the least effortful mode of cognitive control necessary [2]. Also, display and control surfaces should be combined into interfaces that support epistemic (knowledge-seeking) action [3] and help workers abductively maintain their dynamic world model [4].

Over 25 years of development, adoption, and evaluation, EID has been codified into an engineering method [5]. Work Domain Analysis (WDA) serves as a psychologically relevant framework to represent system regularities in terms of structural means-ends, part-whole, and causal/topographical relationships [6]. This analysis then specifies information content requirements and informs the design of analogical graphic forms whose surface features behave consistently with deep regularities of system function [7].

This codification of EID was first and has been most frequently applied in research on systems with strong causal [6] constraints that are high risk and therefore highly engineered. However, WDA was not the only system analysis method proposed by Rasmussen, Pejtersen, and Goodstein.

Strategies Analysis was also proposed as a method to inform ecological information systems [6]. This paper will review the largely un-explored history of this proposal, and present an example of a display designed in this spirit. The example is drawn from the domain of Energy Monitoring & Targeting (M&T). M&T is a business process common in many industries that involves detecting and diagnosing utility (e.g. electric or gas) waste in order to inform corrective action [8].

II. Strategies Analysis for Interface Design

A. Principle

The standard formulation of EID ideally suits complex work domains with stable, causal, and specific constraints. These domains lend themselves to pre-analysis of the work domain as a source of behavior-shaping constraints because:

- Stable constraints mean that regularities can be pre-analyzed when designing an interface. In domains where the rules change, there is less value in supporting people in working within the rules.
- Causal constraints, 'laws of nature', tend to be stable and strong, and can often be idealized into few enough dimensions to be represented as analogical forms on a 2-D interface surface. Intentional constraints are higher-dimensional and difficult to specify.
- Specific constraints mean that design can target "an interface for a single, specific application; generality is not important" [2, p. 595]. Sub-features of an interface may still generalize across domains that share purposes, high-level abstractions, or more aggregated part-whole levels (e.g. one type of aircraft vs. another).

Work domains with all these properties are rare. Since work domains with flexible, semi-intentional constraints are more common, and generalizing across applications allows development costs to be more easily recouped, extending EID to such systems is an appealing prospect.

Foundational EID guidelines are relevant to this goal. Designers are encouraged to "present information embedded in a structure that can serve as an externalized mental model, effective for the kind of reasoning required by the task" [1, p. 525]. However, WDA was not the only recommended structure for an externalized model: "support should not aim at a particular process, but at an effective strategy, i.e. a category of processes." [1, p. 525]. Later, Rasmussen et al. explicitly recommended that "systems serving an autonomous user ... [where] the intentionality depends entirely on the users needs"

978-1-4799-3840-7/14/$31.00 ©2014 IEEE

[6, p. 187], should be designed with *cognitive strategies* as the basis for an ecological interface.

B. Cognitive Strategies Analysis Overview

In an EID framework, cognitive strategies are idealized "categories of cognitive processes" [9], that "will share important characteristics such as

- A particular kind of mental model,
- A certain mode of interpretation of the observed evidence, [and]
- A coherent set of tactical planning rules" [6, p. 70] (bullets added).

Strategies constructs are also considered by psychologists such as Payne [10] or Gigerenzer [11]. Both share a cognitive psychology emphasis on testable computational cognitive models and general-purpose heuristics. Payne emphasizes cognition-in-the-head and specific sequences of elementary information processes, while Gigerenzer's "Adaptive Toolbox" framework explicitly includes an ecological perspective on cognition as environment-actor interaction.

Every interface design method aims to support cognitive strategies, whether implicitly or explicitly. WDA-based EIDs implicitly try to support skill-, rule- and knowledge (SRK)-based strategies. Task Analysis-based interfaces implicitly support the strategy inherent in the specific observed task process. User-Centred-Design (UCD) interfaces will support users' preferred strategy, but will be limited to existing strategies familiar to users due to the task-artifact cycle [12].

How then do 'ecological' design principles [6, p. 133] apply to an explicitly strategies-based approach to EID? An ecological psychology perspective suggests inducing strategies that let actors efficiently exploit regularities in the environment to effectively act, using the least complex mental model or representation necessary [2], [3], [11]. At the same time, at least one of each SRK strategy should be supported to "allow a user to act with which ever level of cognitive interpretation/control they wish" [6, p. 189]. A formative strategy-based EID could introduce novel strategies or enable previously un-feasible strategies [12]. Finally, an EID should support naturalistic switching between multiple strategies to overcome local resource shortages [9]–[12].

C. Example: Bookhouse Strategies-based EID

The most often referenced (and perhaps only) strategies-based ecological interface is the BookHouse library information retrieval system [4], [6], [13]. Informed by field studies comprising 500 recorded user-librarian negotiations and 200 interviews, Pejtersen performed a full cognitive work analysis of a public library. This included a strategies analysis for both information retrieval and cultural mediation tasks, and a worker competencies / user characteristics analysis of children aged 4-15 and adult library patrons' skills and needs.

The subsequent EID process developed support for four distinct information retrieval strategies by specifying the novel "AMP" database scheme. The design process specified how the work domain (library fiction) should be structured (with a knowledge base) to better support not just bibliographic but also analytical, analogical, and browsing strategies.

The BookHouse interface relied heavily on iconography and symbolic metaphors. To practitioners in 2014, the realistic art style and 'virtual room' navigation may not have aged well. However, it is a theoretically grounded and pioneering information system design effort.

D. Conclusion: Matching domains and EID methods?

The canonical WDA-based EID was developed for systems with stable, causal, and specific work domain constraints, and encourages use of analogical graphic forms.

TABLE I. COMPARISON OF EID DOMAINS AND METEHODS

EID Example	Domain Properties			Analysis Method	Graphic Forms
	Stable?	*Causal?*	*Specific?*		
Nuclear Power	Yes	Yes	Yes	WDA	Analogic
Library Fiction	Mostly	No	Mostly	Strategies	Metaphoric
Energy Monitoring	No	Mostly	No	Strategies?	Analogic?

For intentional work domains with weaker, less specific constraints, Strategies-based EID was proposed [6, p. 187], and demonstrated using metaphorical iconic forms (see TABLE I.). However, EID has not been demonstrated for (more common) categories of domains with less stable, semi-intentional constraints. We present an example of a strategies-based approach for one such domain category.

III. EXAMPLE: ENERGY MONITORING & TARGETING

A. Domain Characteristics

Energy Monitoring & Targeting (M&T) is a business process done to "decipher [utility energy] information in order to make good energy and business decisions" [14, p. 2], and "detect avoidable energy waste that might otherwise remain hidden." [8, p. 1]. M&T was developed in the mid-1980s as a general purpose energy efficiency method [15], and has been incorporated into the contemporary ISO 50001 standard [16].

M&T domains are different from those in past applications of EID. In M&T, while some causal constraints such as laws of conservation of energy, thermodynamics, and heat transfer remain stable, the underlying business purposes, processes and functionality do not. Businesses adapt over time to changing economic priorities and their equipment degrades or is maintained. Furthermore, M&T information systems **must** generalize across businesses to be economically viable.

We characterized M&T as a detection, search and diagnosis task [17]. It is abductive in that the goal is *understanding* energy use. Workers can reconcile ambiguity by either changing targets (if socially desired), or by searching for additional observations with a practicable diagnosis strategy.

Sometimes, energy waste can be perceived directly (e.g. lights on in an empty room), or mediated by tools (an infrared picture directly indicates heat loss). In such cases 'good housekeeping' strategies [18] are effective. More often, un-productive energy use can't be directly perceived because cues are unmeasured (imperceptible or forgotten), indistinguishable

This work was supported by the Canadian Federal Economic Development Agency of Southern Ontario and Energent Inc.

(obscured by variability), or control-rejected (compensated for elsewhere).

When direct perception fails, workers need alternate strategies to infer energy performance from observable data.

B. Strategies for Monitoring & Targeting Energy

Through literature review [18] and a longitudinal field study of energy analysts and managers performing M&T in three work environments [17] we identified six strategies:

- 'Good Housekeeping' condition survey
- Equipment inventory, reconciling with consumption
- Consumption time-series profile pattern-detection
- Comparative analysis against normative/historic model
- Energy consumption and cost analysis
- Event – action time-series association

The most relied-on strategy was 'comparative analysis' [18]. In this strategy, the conceptual model of the business is as a symbolic 'black box', with

- Metered utility consumption as energy input y.
- Some measures x_n of productive output or controlled disturbances (e.g., production, hours of operation, sales, weather) as indicators of effective energy use.

The constraints between inputs and outputs are then inferred, most often empirically using historic data and regression [8]. This generates a symbolic energy performance model, usually of a standard linear form

$$\hat{y} = \alpha + \beta_1 x_1 + \ldots + \beta_n x_n \tag{1}$$

$$y - \hat{y} = \varepsilon \tag{2}$$

This model is then used as a normative comparator to monitor ongoing energy consumption y. Residuals ε, the difference between actual and model-predicted energy consumption, are a cue to infer changes in energy performance.

This strategy for detecting energy waste is an example of the time series structural change detection family [19]. Similar problems are addressed in the quality control domain using the classic Shewart Control Chart [8]. However, unlike product quality, variation in energy consumption is not inherently bad. M&T work is concerned with diagnosing persistent changes in energy performance that are financially significant.

C. Existing representations for Energy M&T

Control Charts are used as a representation in M&T. However, perceiving small changes in mean is difficult. If the accumulated energy over/under-consumption is summed before plotting (as in Figure 1), it becomes a Recursive Cumulative Sum of Residuals (CUSUM) chart [8], [15], [19], which has improved perceptual characteristics. A CUSUM representation transforms a shift in the position of an irregular control chart line into a shift in *slope* between smooth CUSUM line segments. These perceptual advantages may have contributed to CUSUM charts remaining the standard tool in M&T for over 25 years [15]. From our field study and analysis, however, we identified two weaknesses of CUSUM charts in supporting M&T work needs.

First, because CUSUM charts show only an aggregated residual, they are ambiguous. A CUSUM chart change may be due to changes in energy performance, but could also be due to quirks of model representation, missing data, changes in measurement practice, interference from past persistent changes or other factors [20]. Diagnosing this ambiguity is difficult without a correct understanding of the energy model, which practitioners in our field study did not demonstrate [17].

We addressed the first problem by developing a display of the energy performance model purpose, process, and product [21]. Our intent was to help use energy performance models as "a structure that can serve as an externalized mental model" [1, p. 525] to support knowledge based reasoning about aggregated CUSUM chart residuals [1, p. 524]. We will not discuss this design here, except to say that that it complements and enables solutions to the second problem.

A second weakness is that CUSUM charts, while an effective detection aid, provide little support for the more important search and diagnosis needs of M&T work. A change in CUSUM slope indicates the date and net magnitude of over/under-consumption, but CUSUM time-series shapes aren't necessarily diagnostic of underlying system changes. Without effective diagnosis, energy waste can't be rectified.

To address the second problem, we developed and prototyped a model-driven diagnosis aid display using a novel statistical strategy and inspired by EID principles. Our goal was to implement a "consistent information transformation concept for data integration" [1, p. 525] and to develop as close to "a consistent, one-to-one mapping between the invisible, abstract properties of the process and the cues or signs provided by the interface" [1, p. 530] as possible.

The remainder of this article will discuss adapting this novel statistical M&T strategy to behave as a better analogical display of energy performance.

IV. Recursive Estimates as a Diagnosis Strategy

A. Rationale for strategy support

We investigated several novel statistical strategies for M&T, and selected the Recursive Estimates (RE) fluctuation process [22] as a promising candidate, because it was:

- Capable of supporting energy diagnosis by estimating *structural change in model parameters*, β in (1),
- Compatible with the representations (linear regression models) already used in CUSUM-based M&T [8], and
- Implemented in the free and open-source R statistical computing environment [19].

Diagnosis support would specifically be for topographical search strategies [6, p. 72] within the functional structure represented by the energy performance model. Topographical search is already practiced in M&T through direct 'good housekeeping', equipment inventory-maintaining, and utility sub-metering [18]. However, sub-meters have disadvantages:

- They add capital and operational cost to calibrate and maintain data quality. Operational costs contradict the purpose of M&T work, which is to save money.
- They multiply the models required for comparative M&T, with associated maintenance and data availability costs.
- They support topographic search only at Physical Form (e.g. which physical areas are metered) or Physical Function (e.g. which equipment is metered) levels [6].

M&T energy models, by contrast, are often formulated in terms of Generalized Functions [6], as measures x_n (1) chosen to represent productive processes. This means that model parameters β can be interpreted in terms of energy intensity (e.g. kWh / Ton of Production). If a detected increase in energy consumption is associated with an increase in only one estimated model parameter, that cue can narrow the search space for conclusive diagnosis. For example, an increase to a building's weather sensitivity (i.e., fuel energy required per degree of outside cold weather) [8] would narrow the search to the process of "heating", which is associated with not only equipment (e.g., furnace efficiency) but functionally related business structure (e.g., building insulation, air-tightness, thermostat settings, internal heat generation) that would be difficult to isolate with a sub-metering approach.

B. Summary of RE charts

The algorithm for calculating RE charts is explained at length elsewhere [19], [22]. Briefly, for each new monitoring datum the RE algorithm:

- Re-fits a new monitoring model to include new data,
- Differences the estimated parameters *α, β* (1) between the new (monitoring) and old (trained) models, and
- Transforms the estimated parameter changes according to the training covariance matrix $Q^{1/2}$ and scales by (3).

$$\varepsilon_{\beta nRE} = (1/\hat{\sigma})(k/\sqrt{n})(\hat{\beta}_n - \beta_n) \quad (3)$$

The scaling in (3) is according to $\hat{\sigma}$, the model residual standard deviation over the training period, *k*, the time-series sequence of each data point and *n,* the length of the historical training set (in the same time units as *k*). This produces a set of time-series test statistic charts, one for each model parameter.

C. Modifications to RE strategy to improve analogical form

Stock test-statistic-based RE charts have several interpretation weaknesses, both statistical and perceptual [23]. Since CUSUM charts are more statistically robust, our design encouraged switching between CUSUM and RE strategies, and focused on reducing the perceptual ambiguities in RE charts.

To support rule-based M&T behavior, RE charts should ideally have a 1:1 mapping with (i.e. be analogical to) energy performance [2], [7]. However the stock RE algorithm produces charts scaled to represent a probabilistic test statistic for formal hypothesis-testing. This leads to some ambiguities when trying to interpret it as an analogical form representing underlying system change.

1) Response rate: The stock RE algorithm recursively grows the monitoring set with each new data point *k* (3), so an identical change produces a different chart shape depending on being early or late in the monitoring period. As the monitoring period expands, the charts respond slower, conflating large intermittent and small persistent changes.

2) Scaling: Stock RE charts scale up by $k/\sqrt{n}$ (3) as the monitoring set grows, so a constant energy performance deviation tends to produce an expanding $\sqrt{x}$ shaped RE chart.

3) Certainty: Straight-line stock RE charts are ambiguous - they can result either from consistent energy performance or from a driver variable holding constant for long periods (such as off-seasonal heating/cooling requirements). This means that if a physical system change occurs while a driver is 'stagnant', the effect won't be charted until the associated driver varies again. This ambiguity could reinforce over-confident estimations of change dates, for example not suspecting that poor fall heating performance could be due to building insulation damage caused during summertime renovations.

Relaxing the design constraint of statistical tests and reframing RE charts as analogical cues rather than degree-of-certainty time-series tests allows flexibility to achieve:

1) Consistent response: Discarding old observations allows the RE chart to respond consistently to changes regardless of date. Exponential decay weighting is a common choice, and can be metaphorically described as having "long-" or "short-" memory.

2) Proportionality: Scaling RE charts in a time-invariant way prevents early persistent changes from being amplified.

3) Informativeness mapping: De-emphasizing RE charts when underlying driver variables stop varying (conveying information) distinguishes "no information" vs. "no change".

These changes still under-specify design parameters such as the rate of exponential decay in weighting historic data, and how long driver variables must stagnate before estimates are considered uninformative. For lack of design guidance, we linked the stagnation criteria to the time constant of exponential decay, and superimposed two RE charts at pre-specified slow and rapid exponential decay rates. The RE chart with long memory (multi-year time constant) is shown as a black line in Figure 1, while the short memory (60-day exponential decay) is shown as a grey area chart.

While a design based on just an empirical linear model cannot achieve "a unique and consistent mapping between … the behavior of the process and the … cues that the interface displays" [1, p. 528], it is less ambiguous than CUSUM charts and improves on stock RE charts by behaving more consistently, proportionally, and informatively.

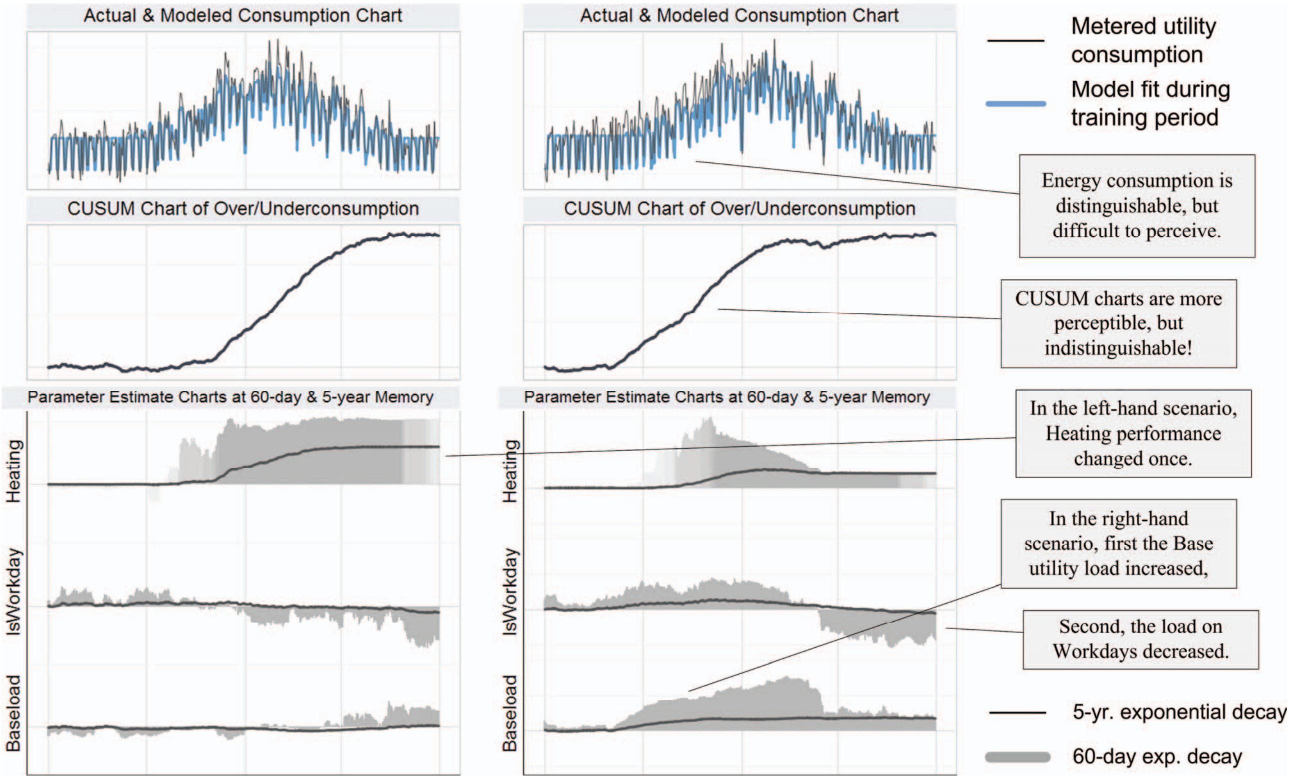

Fig. 1. Two year-long scenarios of energy performance changes in a hypothetical office building. Utility fuel is consumed by three categories of processes: winter space heating, workday hot water use, and always-on base load. An energy model represents these three processes as model parameters. The model is used to generate three time-series charts: predicted consumption (top), CUSUM of residual (middle) and a cluster of modified analogical RE (bottom).

D. *Design result*

An example demonstrating the behavior of the modified analogical RE charts is shown in Figure 1. The plots share a time X-axis with consumption, modeled, and CUSUM charts. Every relevant chart should be provided for knowledge-based reasoning, but CUSUM charts are integral to supporting rule-based interpretation. Two tactical rules that seem invariant are:

- If a change is apparent in the CUSUM chart, suspect the driver whose long-memory RE chart *most closely resembles* the CUSUM chart *at the time of the change*.
- For a change indicated by CUSUM or RE charts, inspect short-memory RE charts as indicating driver energy intensity. A steady RE chart indicates steady driver-related performance. An RE shift's earliest point indicates a suspected true onset time of the underlying structural change in the system's energy performance.

These two rules, one recognition-based, one analogy-based, are a starting point for developing situated expertise. Since each business application domain will have a different underlying set of causal and intentional constraints, and each change in performance will manifest as different changes to system functional measures, we cannot completely pre-specify specific interpretation procedures.

Knowledge-based reasoning from these graphic forms is possible if they are supplemented by a complimentary model summary report [21] which provides symbolic representations of the quality of data, processing, and models [20].

Finally, we hope that better initial diagnoses will be investigated by in-person topographical search, employing human perceptual skills to identify root causes of energy waste.

V. Discussion and Future Work

This paper has reviewed an under-explored proposal for an EID method based on Strategies Analysis rather than WDA. We have presented an example strategies-based display for Energy M&T, which differs from typical EID application domains in being unstable and generalized. However, because the M&T domain obeys some causal constraints (e.g. conservation of energy), visual analogical forms were feasible rather than relying on metaphors as in the canonical BookHouse strategy-based EID.

A. *Ecological Interface Criteria*

Since so few strategies-based EIDs have been described in the literature, it is worth considering how well the ecological interface criteria described earlier manifest in this example.

1) *Merged display/control surfaces?* The design presented here is admittedly only a display, not an interface.

2) *Consistent, one-to-one mappings of cues and properties?* Unlike canonical EIDs (though see [24]), this design only empirically estimates abstract properties through

coincident correlation, rather than *a priori* derivation of system constraints. The quality of the mappings depends on the quality of the normative model (which varies depending on whether it is derived from first principles or data mining). Using time-series charts limited the available dimensions for representation to the vertical axis. However, this still allowed emergent patterns to be plotted over time.

3) Can the same form be interpreted as signals, signs and symbols (in terms of SRK)? The output of the symbolic RE strategy is a set of time-series signals, which can be interpreted using recognition and analogical 'rules of thumb'. Organizing the display according to a symbolic energy model may encourage workers to reason about energy performance in terms of energy intensity of business processes *[17]*.

4) Does it allow naturalistic strategy switch behavior? The M&T tool shows CUSUM and RE charts simultaneously, allowing switching between CUSUM to robustly date and quantify change effects and RE to diagnose possible causes. A more ecological strategy switch would be to walk and physically search the site, guided by an initial diagnosis.

This strategy-driven display design arguably meets many of the principles of ecological interfaces, and we believe it was "explicitly designed on the basis of a detailed understanding of the work ecology" [4, p. 137].

B. Future Work

The design is an excerpt, and remains under development. It has been evaluated in a controlled laboratory experiment with representative students in their final year of a building energy management professional program. Evaluation in representative practical settings will be performed as it is introduced into practice. Future work to extend features for interactive thought experiments and hypothesis-testing will help extend it closer to an interface meeting all the design principles that inspired EID [1].

RE charts are not limited to energy monitoring. They are suited to time-series change detection in any system whose normative behavior can be described in terms of measured independent variables [19]. While originally developed for retrospective hypothesis-testing using statistical significance tests, when used on-line for system monitoring, workers can switch to more abductive, knowledge-seeking behaviors. In such cases, RE charts may be more useful in directing search than confirming hypotheses. Statistical-ecological design strategies offer many possibilities for lightweight yet sophisticated socio-technical system control schemes.

ACKNOWLEDGMENT

A.H thanks D. Jorjani and B. Hayes for design guidance and implementation, and W. Giang for EID discussions. This work was funded by Energent Inc., the Natural Sciences and Engineering Research Council of Canada, and the Federal Development Agency for Southern Ontario.

REFERENCES

[1] J. Rasmussen and K. J. Vicente, "Coping with human errors through system design: Implications for ecological interface design," *Int. J. Man-Mach. Stud.*, vol. 31, pp. 517–534, 1989.

[2] K. J. Vicente and J. Rasmussen, "Ecological interface design: Theoretical foundations," *IEEE Trans. Syst. Man Cybern.*, vol. 22, no. 4, pp. 589–606, 1992.

[3] A. Kirlik, "Abstracting Situated Action: Implications for Cognitive Modeling and Interface Design," in *Adaptive perspectives on human-technology interaction methods and models for cognitive engineering and human-computer interaction*, A. Kirlik, Ed. Oxford;; New York: Oxford University Press, 2006, pp. 212–229.

[4] K. B. Bennett and J. Flach, *Display and interface design: subtle science, exact art*. Boca Raton, Fla.: CRC Press, 2011.

[5] C. M. Burns and J. R. Hajdukiewicz, *Ecological Interface Design*. Boca Raton, Florida: CRC Press, 2004.

[6] J. Rasmussen, A. M. Pejtersen, and L. P. Goodstein, *Cognitive systems engineering*. New York: Wiley, 1994.

[7] J. Rasmussen, "The human data processor as a system component: Bits and pieces of a model," Danish Atomic Energy Commission Riso, Roskilde, Denmark, Technical Report Riso-M-1722, 1974.

[8] Carbon Trust, "Monitoring and targeting: Techniques to help organisations control and manage their energy use," The Carbon Trust, CTG008, 2008.

[9] J. Rasmussen and A. Jensen, "A study of mental procedures in electronic trouble shooting," Riso National Laboratory Electronics Department, Roskilde, Denmark, Riso-M-1582, 1973.

[10] J. Payne, J. R. Bettman, and E. J. Johnson, *The adaptive decision maker*. Cambridge;; New York, NY: Cambridge University Press, 1993.

[11] Dahlem Workshop on Bounded Rationality and G. Gigerenzer, *Bounded rationality: the adaptive toolbox*, 1st MIT Press paperback ed. Cambridge Mass.: MIT Press, 2002.

[12] K. J. Vicente, *Cognitive work analysis: toward safe, productive, and healthy computer-based work*. Mahwah, N.J.: Lawrence Erlbaum Associates, 1999.

[13] A. M. Pejtersen, *The Bookhouse: modelling user' needs and search strategies a basis for system design*. Roskilde Denmark: Risø National Laboratory, 1989.

[14] J. H. Hooke, B. J. Landry, and D. Hart, "Energy management information systems - Achieving Improved Energy Efficiency: A handbook for managers, engineers and operational staff," 2003. [Online]. Available: http://oee.nrcan.gc.ca/publications/industrial/EMIS/index.cfm?attr=24.

[15] P. Harris, *Energy monitoring and target setting using CUSUM*. Cheriton Technology Publications, 1989.

[16] ISO Technical Committee 242, "ISO 50001:2011 - Energy Management." International Standards Organization, 2011.

[17] A. Hilliard and G. A. Jamieson, "Monitoring & Targeting Energy in Practice: A Field Study," in *Proceedings of the 2014 ECEEE Summer Study in Industry*, Arnhem, NL, 2014, pp. 591–601.

[18] CIPEC, "Energy Savings Toolbox - an Energy Audit Manual and Tool." Canadian Industry Program for Energy Conservation, 2010.

[19] A. Zeileis, F. Leisch, K. Hornik, and C. Kleiber, "strucchange: An R Package for Testing for Structural Change in Linear Regression Models," *J. Stat. Softw.*, vol. 7, no. 2, pp. 1–38, 2002.

[20] G. A. Jamieson and K. J. Vicente, "Designing effective human-automation-plant interfaces: A control-theoretic perspective," *Hum. Factors*, vol. 47, no. 1, pp. 12–34, 2005.

[21] A. Hilliard, G. A. Jamieson, and D. Jorjani, "Communicating a Model-based Energy Performance Indicator," in review.

[22] W. Ploberger, W. Krämer, and K. Kontrus, "A new test for structural stability in the linear regression model," *J. Econom.*, vol. 40, no. 2, pp. 307–318, Feb. 1989.

[23] A. Hilliard and G. A. Jamieson, "Recursive Estimates as an Extension to CUSUM-based Energy Monitoring & Targeting," in *Proceedings of the 2013 ACEEE Summer Study on Energy Efficiency in Industry*, Niagara Falls, NY, 2013, pp. 4–1..4–13.

[24] N. Lau, "Numerical Models in Representation Design: Computing Seawater Properties in an Ecological Interface," in *Proceedings of the Human Factors and Ergonomics Society 50th Annual Meeting*, 2006, pp. 245–249.