Engineering Representations to Support Evidence-based Clinical Practice

John M. Flach, Peter Reynolds, Caroline Cao & Tiffany Saffell
Wright State University
Dayton, OH, USA

This paper provides an introduction to Cognitive Systems Engineering (CSE) and Ecological Interface Design (EID), as important complements to more conventional Human Factors Engineering approaches. These complementary perspectives are essential for supporting productive thinking in complex work domains, such as healthcare. We suggest that EHR systems provide a unique opportunity to take advantage of these approaches to support Evidence-Based Practice (EBP) in healthcare and we show examples of these approaches to three different healthcare problems: cardiovascular health, pain management, and anemia.

COGNITIVE SYSTEMS ENGINEERING

As a result of dramatic advances in the integration of information technologies into modern work domains during the 1980s and 90s, the primary role of humans in many work domains shifted from manual control (e.g., piloting) to supervisory control (e.g., monitoring and fault diagnosis). In this new role, humans were asked to do more than simply conform to pre-established procedures and standards. Additionally, they were expected to intervene when the automated systems failed or when situations arose, which were not anticipated in the design of the automated systems. Filling this new role required tapping into the creative problem solving capabilities of humans.

Cognitive Systems Engineering (CSE) developed as a complementary discipline to more traditional Human Factors (HF) in order to address the demands to design systems to support creative problem solving (Norman, 1986; Rasmussen, 1986). Whereas HF tended to emphasize the information processing limitations of humans, CSE began to focus attention on expertise and the capacity of humans for productive thinking (Wertheimer, 1959). This productive thinking was necessary in order to creatively adapt to the inevitable situation variability that could not have been anticipated in the design of standard procedures or automated control systems. The goal was to tap into human experts' abilities to organize information (e.g., chunking) and to leverage task constraints (e.g., heuristics) to effectively by-pass some of the normal constraints on information processing (Eriksson & Charness, 1994) and to create solutions to novel problems.

The distinction between *task* and *work analysis* is a second important distinction between classical HF and CSE. Classically, task analysis has focused on both physical and mental activities and has evaluated those activities relative to compliance to standard procedures or to normative models of rationality. Work analysis was introduced as an important complement to task analysis. Rather than focusing on activities, work analysis focuses on possibilities or, conversely, constraints (Vicente, 1999; Naikar, 2013). That is, the goal of work analysis is to provide insight into the *deep structure* of a work domain. This provides a context for considering alternatives to the standard procedures that might be required in order to adapt to situations that were not anticipated when the standard procedures were formulated. Such adaptation requires sensemaking or abductive reasoning (Flach & Voorhorst, 2016).

A third distinction between HF and CSE is associated with the design of interfaces. In classical HF, the emphasis was on ease of acquisition of *information* from the interface. Thus, the mantra was to *match the human's mental model,* since that would typically facilitate the pick-up of information. In CSE there was an additional goal of maximizing meaningful insight into the work domain problems or processes (Woods, 1991). Thus, the design goal was to shape the human's internal models, so that they would reflect the deep structure of the process dynamics. For example, in designing interfaces for process control it was not sufficient to display the states of a process (e.g., temperatures and volumes), it was additionally necessary to reflect the underlying thermodynamic constraints (e.g., mass and energy balances) (Vicente, Christofffersen & Pereklita, 1995). This additional requirement to create representations that revealed the underlying process dynamics was given the label: *Ecological Interface Design* (Rasmussen & Vicente, 1989). The goal was not simply to support procedural compliance, but additionally to support learning and abductive reasoning required for adapting to variability that could not be anticipated in the design of *a priori* procedures or rules.

Ecological Interface Design (EID)

The EID approach to the creation of representations was framed in the context of a triadic model of semiotics as illustrated in Figure 1. The triadic model shows the coupling between the human agents (e.g., physicians) and a work domain (e.g., patients) through a representation [e.g., an electronic healthcare record (EHR) system]. Classically, this coupling has been parsed between the engineers/computer scientists who focused on the link between the problem domain and the representation and the HF engineers who focused on the link between the representation and the human agents. The engineers/computer scientists tended to focus on the 'physics' of the underlying processes (i.e., the objective state variables) and sensor capabilities. The HF engineers tended to focus on human perceptual and cognitive limits. There was at least an implicit assumption that if each of these two components (engineering and HF) were both addressed properly, together they could be combined to create *user-centered* representations that would facilitate information processing and lead to successful interactions.

In contrast, the EID approach begins by focusing on the pragmatic or functional constraints reflected in the direct interactions between the agent and the work domain. That is, EID assumes that understanding the deep structure of the functional work (e.g., the consequences of potential actions relative risks and objectives specific to a work domain) is a prerequisite for choosing the appropriate state variables and for organizing representations to shape appropriate mental models. Thus, in EID the goal is to create *use-centered* representations that make the meaningful functional constraints salient (Flach & Dominguez, 1995). Such representations are intended to shape the agent's mental models so that they are well grounded with respect to the functional work constraints.

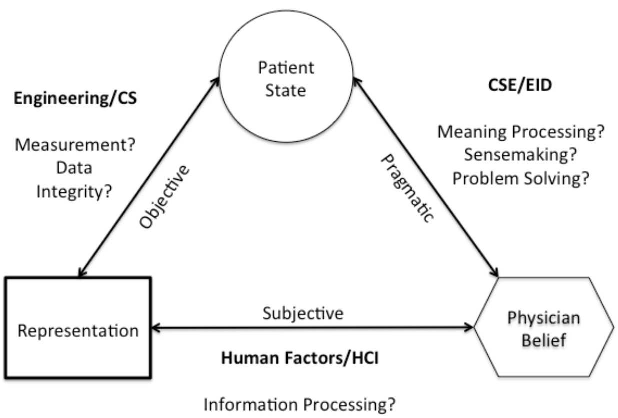

Figure 1. The triadic semiotic system underlying sensemaking or abductive reasoning.

EID is a holistic approach to the semiotic triad in Figure 1. That is, it recognizes that meaning processing cannot be understood as the sum of the component dyads. However, the EID approach gives precedence to the pragmatic or functional constraints as an overarching context for evaluating the other component dyads. The pragmatic constraints provide a context for choosing the appropriate coordinates for *objectively* representing the state space associated with a physical process in terms of the affordances offered. Additionally, the pragmatic constraints provide a context for choosing how to organize the data into chunks or patterns that reflect meaningful associations/constraints among the various state dimensions in order to shape the *subjective* impressions of an observer.

In terms of the form of representations, the general goal for EID is to create *configural* representations (Bennett & Flach, 2011). That is, it is necessary to provide easy access to both the particulars (e.g., raw test values) and the general (e.g., a global index of health). Further, the particulars must be presented in the appropriate context of the whole. This is because the meaning of a test result generally depends on its relation to other variables. A good representation should show the implication of each data point for the overall health of the complex system. In a configural representation the data are nested or layered within larger patterns that represent global properties of the work domain.

In essence, the goal of EID is to by-pass the capacity limits of human working memory by organizing (chunking) information and creating patterns that will trigger smart heuristics that are grounded in the functional constraints of the work domain. Such representations can enable more recognition-based (e.g., Klein, 1989) styles of processing (i.e., skill- and rule-based processing) and can reduce the need for more cognitively demanding knowledge-based computations. However, while reducing the need for knowledge-based processing, EID representations also include support for hypothesis testing and learning through direct interaction with the representation (i.e., direct manipulation) (Shneiderman, 1992; Shneiderman, Plaisant & Hesse, 2013). The combination of direct perception and direct manipulation in EID representations provides broad support for creative problems solving by leveraging both the opportunity of advanced interactive graphics and the power of human pattern detection and recognition capabilities.

EVIDENCE-BASED PRACTICE IN MEDICINE

There are many obvious parallels between the current roles of operators in high-technology work domains and the role of clinical physicians. That is, the role is essentially to monitor and diagnose problems associated with a complex system. In both situations, the work requires that the human decision makers do more than comply with standard procedures. The work requires that people adapt the procedures to fit with a wide variety of situations that cannot be fully specified in advance (e.g., unique patient situations).

The term *Evidence-Based Practice* (EBP) is a "formalization of the care practice that the best clinicians have practiced for generations" (McKibbon, 1998). EBP involves integrating the scientific knowledge about disease mechanisms and pathophysiology with knowledge about specific patients, situations, and preferences to make well-informed decisions about patient care. It is important to appreciate that EBP does not imply mindless, adherence to *a priori* procedures or practices. It explicitly recognizes the variance associated with patients and situations. The key to EBP is that in adapting treatments to individual patients, the adaptations should be informed by consideration to the established scientific knowledge and standards for care.

In many respects, EBP in healthcare reflects exactly the type of work that CSE and EID evolved to support. In this context, the goal for the work analysis portion of CSE is to identify the appropriate knowledge or evidence-base that constitutes the 'deep structure' of specific healthcare problems. The goal of the EID facet of CSE is then to organize information into representations that make the 'deep structure' (i.e., the evidence-base) salient to the healthcare professionals.

EID in Healthcare

Despite the apparent fit between the underlying logic of CSE/EID and the logic/dynamics of EBP, the medical field has been very resistant to the CSE/EID approach. For example, consider the state of current Electronic Healthcare Record (EHR) systems. There is very little evidence of any effort to create representations that reflect the 'deep structure' of healthcare. The interfaces tend to use standard HCI dialogue boxes, text, and data tables. There is little or no innovative use of graphics to create configural patterns or to make the deep structure of health salient in the interface representations. The result is that current EHR systems are clumsy with respect to finding the particular data and for seeing how they contribute to global health. For example, a major complaint about current EHR systems is difficulty with medication reconciliation and with assessing cumulative effects and interactions among medications.

In many respects, the EHR interfaces are similar to early power plants that were designed based on a single-sensor-single-indicator philosophy. Such systems display the data – but do it in a way that invariably leads to data overload and does almost nothing to facilitate data integration and sensemaking. The only difference between these primitive nuclear power control rooms and the current EHR systems is that healthcare is far more complex than nuclear power. In the primitive nuclear power plant the data fills the walls of the control room – but in EHR systems the data fills the walls of many 'virtual' rooms (i.e., data sets) that can only be accessed through standard computer dialogue boxes. Thus, forcing physicians to navigate through a confusing maze of dialogue boxes to even find the data – much less to discover patterns in the data that reflect meaningful interactions or other global aspects of health.

In our efforts to promote CSE/EID in the healthcare domain, a common argument against generalizing from the successes in aviation (Borst, Flach & Ellerbroek, 2015) and process control (Vicente, Christoffersen & Pereklita, 1995) has been that, unlike healthcare, the processes in aviation and process control were 'engineered' based on well-defined physical laws (e.g., thermodynamics). We feel that such skepticism is based on a rather shallow understanding of EID.

First, it is true that the deep structure of healthcare is very different than the deep structure of aviation or process control. Every work domain is somewhat unique from all others, when it comes to the 'deep structure.' Thus, it would be naive to think that the particular representations that work in one domain would generalize to any other domain. In fact, that is fundamental to EID – the commitment to apply work domain analysis to learn about the particulars of a specific work domain. We suggest that one of the ultimate failures in the design of current EHR interfaces is the uncritical generalization of interface representations (e.g., dialogue boxes) that work for websites or other software packages to the domain of healthcare. These forms are often familiar to the healthcare users due to prior experiences with similar representations – so they are thought to be consistent with good human factors practice (i.e., match the user's expectations). However, while possibly facilitating interaction with the computer, these representations are obstacles to doing the work (i.e., interacting with patients and assessing health). Such representations provide absolutely no value with respect to supporting productive thinking.

It is also true that the evidence-base of healthcare may not be as well-defined and well-validated as the physical laws that underlie physical systems such as aircraft. As we already noted, healthcare is far more complex than many of the engineered domains. However, healthcare does have an evidence-base, and integrating this evidence base into the decision process is as important for healthcare as it is for any other complex domain. Where physical systems might have 'laws' derived from first principles, healthcare might only have 'models' derived from correlational studies. However, the point of EBP is that good physicians should take such models into account when making clinical decisions. For example, the Framingham model provides an important evidence-base for making judgments with respect to cardiovascular health. This model, based on longitudinal observations of a large population, provides important evidence about the relation between cardiovascular health and various test measures such as cholesterol levels and blood pressure.

It is ironic that the increased complexity of healthcare, reflected in the relatively tentative nature of the evidence-base is used as a reason not to pursue an EID approach to representation design, when the primary motivation behind EID is to help people to creatively manage the complexity and associated uncertainty. We believe the kind of support that an EID approach offers has even more potential value for healthcare than it has for the other domains, precisely because the complexity is so much greater. Because the evidence-base for healthcare is extremely complex, EBP is impossible without support in the form of smart representations! Classical human factors considerations remain important, but they are simply not good enough if the goal is EBP!

PROCESS AND EXAMPLES

Although there is a desire for systematic processes or recipes for doing work analysis and for creating EID interfaces and several sources have suggested such processes (Burns & Hajdukiewicz, 2014; Naikar, 2013), our approach has typically been a very opportunistic search in which we tap multiple sources (e.g., practitioners, researchers, incident reports, literature, etc.) and employ various methods (e.g., field observations, knowledge elicitation, table-top analysis, iterative design) to learn about the work domain. The unique aspect of this search relative to more classical HF approaches is that the focus is on the work domain or problem ecology, rather than on the behavior or activities of operators. This creates the possibility to look beyond the horizon of standard procedures to explore the full range of possibilities. Thus, creating opportunities to innovate.

Also, we treat work analysis as a co-requisite of interface design, not a prerequisite. That is, we employ an iterative design approach where the generation and evaluation of design hypotheses are integral to the work analysis process. This is consistent with Schrage's (1999) concept of 'serious play,' where design hypotheses in the form of concrete simulations are used to stimulate feedback from domain experts in a closed-loop, trial and error design process. We feel that such an approach is an effective way to explore complex problems and to converge on satisfying design solutions. In this process, we find formalisms such as the Abstraction Hierarchy to be useful tools for documenting what we learn in the process. However, because the learning process never ends, we believe it is a mistake to make completing such formalisms a prerequisite for design.

In lieu of providing a recipe for designing representations, we will present several examples of concepts that are being developed as representations for EHR systems. Hopefully, these concrete examples will provide some insight into how work analysis informs the EID design process.

Cardiovascular Health

Our interest in developing an improved representation for guiding decisions with respect to cardiovascular health was piqued as a result of field studies to evaluate the management of test results in family medicine practices (Elder, McEwen, Flach & Gallimore, 2009). At the time of these observations most of the family practices were still using paper record systems and the results from blood tests were reported as a list of text and values sent from the labs via a fax machine. It was reported that when reviewing multiple charts, physicians would sometimes fail to recognize that some of the values from the blood tests suggested risks that should be attended. We saw the introduction of EHR systems as an opportunity to apply the EID approach to create representations that would make these data and the implications for cardiovascular health more salient.

Work Analysis. In our search to understand the meanings of the various test results with respect to cardiovascular health, we discovered three sources of evidence that suggested ways to intelligently configure the information. One source was the Seventh Report of the Joint National Committee on Prevention, Detection, Evaluation, and Treatment of High Blood Pressure (JNC-7) (Chobanian et al., 2003). This suggested four different conditions (healthy, pre-hypertension, stage 1 hypertension, and stage 2 hypertension) and associated treatments contingent on blood pressure measures.

A second important source of evidence was the Third Report of the Expert Panel on Detection, Evaluation and Treatment of Blood Cholesterol in Adults (ATP III) (National Heart, Lung and Blood Institute, 2001). This report suggested three categories of health as a function of cholesterol in relation to numerous other factors (e.g., gender, weight, smoking history, blood pressure, diabetes) that were associated with three treatment options: no treatment, therapeutic lifestyle change, or drug therapy.

The third important source of evidence was actually a component underlying the rationality of the ATP III guidelines. This was the Framingham model of cardiovascular risk (D'Agostino et al., 2008). The Framingham model is based on a longitudinal study of a large cohort of people (> 5,000) to assess the development of cardiovascular disease that was begun in 1948 and is still on going. The Framingham model provides a basis for computing a 'score' that indicates the likelihood of a cardiac incident (e.g., heart attack or stroke) within the next 10 years as a function of cholesterol and numerous other factors.

Configural Representation. Figures 2 & 3 illustrate the configural representation that was developed for assessing cardiovascular health. Note that there are four regions. Region 1 provides the values of particular variables (e.g., cholesterol levels, blood pressure, whether the patient smokes). For many of the continuous variables the particular value is positioned on a color-coded continuum that provides a normative context for interpreting the value (e.g., green means the value is in the normal healthy range).

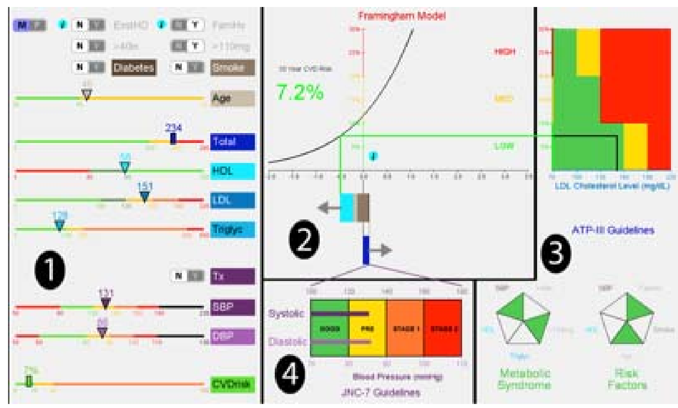

Figure 2. A configural representation that organizes data for a patient in the context of models of cardiac risk and published standards of treatment for cardiovascular disease.

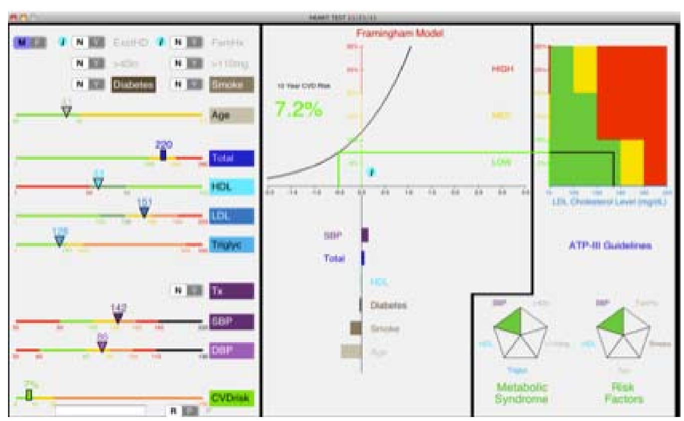

Figure 3. The exploded view makes it easier to see the weighted contribution of particular variables to the Framingham risk score.

Region 2 represents the Framingham model as a graphic that maps the weighted sum of the particular health variables through the Framingham function to determine the Framingham score. A unique aspect of this mapping is that a contribution graphic on the x-axis (see Figure 2) shows how each of the component variables contributes to the score. Color was used to map the particular variables from Region 1 into the contribution graphic. Because this graphic could be difficult to read, there was an option to toggle to a second view that ranked the variables as a function of their weighted contribution to the risk score (see Figure 3). Variables that were to the right of center increased risk and the variables to the left of center reduced risk. The highest risk variables were at the top of the list.

Region 3 of the representation was based on the ATP III guidelines. A state space graphic shows three distinct regions associated with different treatment recommendations: green – no treatment required; yellow – lifestyle change; red – drug therapy. The boundaries of these regions was determined based on a combination of the particular risk factors shown in Region 1 and the position of the patient within this state space was a function of the Framingham score (projected from Region 2) and the LDL cholesterol level (specified in Region 1).

Region 4 was based on the JNC-7 guidelines. It showed four different treatment categories and the patient's state relative to those categories was determined as a function of blood pressure. Specifically, the highest projection of either the systolic or diastolic value determined the treatment recommendation.

Evidence-based Practice. It is difficult to fully appreciate how the interactive graphic supports reasoning about cardiovascular health and helps physicians to choose an appropriate treatment, based on the static images (Figures 2 & 3). Thus, we have created a video that illustrates our concept of use and the direct manipulation capabilities of the interface:

https://www.youtube.com/watch?v=dIZtMDo7im4&feature=youtu.be

In using the interface the significant impact of smoking on cardiovascular health became particularly salient. For example, the different between smoking and not smoking could result in a shift from the red region of the ATP-III state space to the green region. This is particularly relevant to EBP, because a physician who followed the ATP-III recommendations blindly would initiate drug therapy. However, a physician, who could "see" that smoking was a significant factor in the Framingham score, might recognize that this is actually a life-style problem. This physician could then make a more informed judgment about whether to address the life-style problem directly or to follow the ATP-III guidelines. To emphasize, EBP does NOT mean blind adherence to guidelines, rather it means intelligently adapting the guidelines to the particulars of a specific patient.

Pain Management

Our interest in pain management was the result of a conversation with Dr. Zach Hettinger at the National Center for Human Factors in Healthcare. Dr. Hettinger challenged us to apply the logic of EID to pain management. He noted that there was little support in current EHR systems to help physicians to assess the cumulative impact of multiple opioids on a patient's health (and pain). This was particularly difficult because of the different strengths of specific opioid medications. He also noted that in some cases patients might be getting multiple opioids, but that this information may be presented in different EHR windows as a function of when the orders were written (e.g., pre- or post-op) or the mode of delivery (IV or oral). Thus, a physician who is adjusting the oral dosing for a patient post surgery may not be aware that the patient is also getting additional opioids via IV.

Work Analysis. In trying to understand the cumulative effects of opioids, there were a number of obvious questions: Is there a common scale so that the effects of multiple different drugs can be combined to calculate the total level of opioid being administered? How do the various opioids dissipate over time? What is the standard dosing? How do you know if a person is getting too much? How do you know whether a person is getting enough?

We learned that medical researchers typically use a morphine equivalence scale to index the relative potency of different opioids. Note that this scale is not without some controversy, so physicians are warned to consider other factors and to be cautious when changing opioid medications (Schatman, Fudin, 2016).

We found that there are two ways to estimate the impact and dissipation of the opioids. In the literature we found estimates for the onset and duration of effectiveness for the various opioids and we also found estimates for the half-life of the opioids. Note that these did not seem to be independent, in that the duration of effectiveness tended to be equal to one half-life.

Subjective pain scales appeared to be the most common way to assess the patient's pain, and thus, the impact of the opioid relative to reducing pain and assessing whether the dosage is sufficient. The scales typically range from 1 (no pain) to 10 (intense pain), with some scales using face images to indicate different levels of pain.

A major concern relative to too much opioid is the danger of respiratory arrest. Thus, measures associated with respiration (e.g., respiratory rate in breaths per minute, or pulse oxygen levels) are useful indicators of whether a patient is getting too high a dose of opioid.

Configural Representation. Since the fundamental question with respect to managing opioids involved cumulative effects over time, our first hypothesis for how to configure the information was to create a time-history graph as illustrated in Figure 4. Color is used to indicate different opioids and the solid black line provides an indication of the total opioid level. A half-life model of dissipation was used for individual opioids and an effectiveness duration level was used for the total opioid level. With this format the initial dosing time of a specific opioid is shown as a peak and this code prevents the line indicating total level from occluding the lines for individual opioids. A dotted black line was included to indicate a 10% break-out level above the total opioid line. This was a common heuristic for indicating the minimal change need to have a noticeable impact on pain.

The red line shows the patient's subjective rating of pain on a 10-point scale for the same time period. This can easily be compared to the opioid levels to help judgments about whether the levels of opioids are high enough to reduce the experience of pain. A history for respiration rate is also plotted in green over the same period. The green shaded region indicates the range of normal respiration rates (12 – 20 breathes per minute). This is useful for making judgments about whether the levels of opioids are too high. A drop in respiration levels would be a potential indication of unsafe levels of opioids.

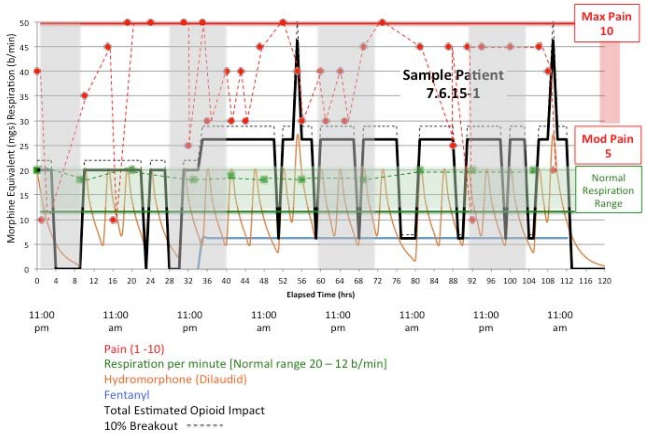

Figure 4. A time history graphic showing the cumulative morphine equivalent dosage levels resulting from multiple opioids. Respirations per minute and subjective pain ratings are included as indicators of the effectiveness and of potential side effects.

Initial reaction from physicians to the time-history graphic (Figure 4) was that it might be too much data in too complex a form. They challenged us to develop a simpler representation that would allow them to assess the current state of the patient in a single glance. Figure 5 shows our concept for a less complex representation that nevertheless shows important relations for managing pain.

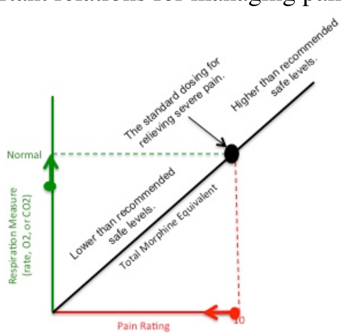

Figure 5. A configural representation that shows opioid levels relative to standard dosage recommendations and in relation to respiration rate and pain rating.

The configural graphic in Figure 5 uses a “box” metaphor in which the total morphine equivalent level is shown in relation to respiration rate and pain rating. The lower corner of the box is the origin and the lengths of the adjoining sides reflect the normal level for respiration and the highest pain level. Total morphine equivalent level is plotted along a diagonal of the box, such that the fourth corner of the box corresponds with the standard dosage of opioid (approximately 20 mg morphine equivalent). The last two readings of each of the three variables (i.e., respirations rate, pain rating, and total morphine) are plotted as vectors along these three axes. The head of the vector is the most recent value and the base of the vector is the prior value. If the two values are identical, then the values are shown as an enlarged point.

The vectors provide an indication of the direction of change. For example, in Figure 5 the morphine dosage stayed at the standard level, pain level was going down, and respiration level was in a healthy range and increasing. The box configuration provides a salient indication of the relations among the three variables and the standard dosage for opioid. If the dosing level is ‘inside-the-box” then it is lower than the standard and if it ‘outside-the-box” then it is higher than the suggested standard.

Figure 6 shows how the two graphics might be configured into a dashboard interface that also included a listing of the time, specific opioid dosages, and morphine equivalence. This is a very early concept that we are using as a stimulus to continue our work analysis in order to elicit a deeper understanding of the pain/opioid management problem.

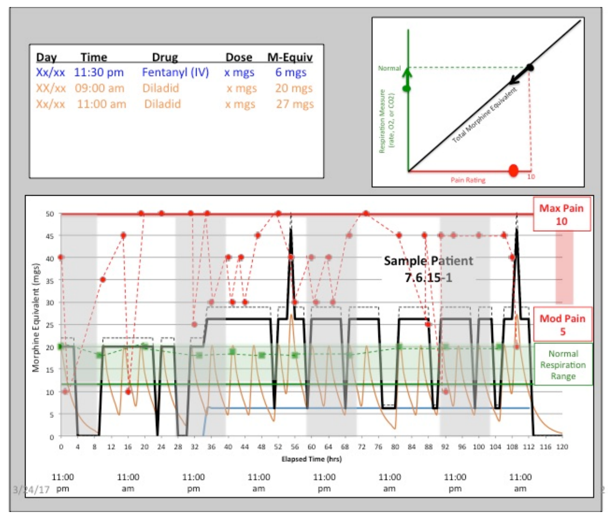

Figure 6. A dashboard to help physicians monitor opioid levels in order to manage a patient’s pain.

Evidence-based practice. The pain/opioid management problem is a particularly good example of why EBP cannot be reduced to simple standards of practice or procedure following. There can be wide variation in terms of people’s tolerance for pain, the impact of the various opioids on that pain, and the potential for negative side effects. Thus, it is typically recommended that physicians start with a standard recommended dosage and then titrate the dosage to the levels appropriate for a particular patient/situation. However, current EHRs provide very little support for this titration process. The goal of our design is to provide an integrated representation to help the physician to easily see the levels of opioid relative to a patient’s pain (to judge whether the current dosage is sufficient) and relative to a patient’s respiration level (to judge whether the current dosage is too high).

Anemia

We have just recently begun to consider the problem of anemia. At this point we are in the very earliest stages of our work analysis in order to begin to understand the different types of anemia and the information that physicians use to discriminate the different types and to choose appropriate treatments. We include this in the paper to emphasize the critical role of work analysis in developing representations. The representations have to be constructed in ways that reflect the deep structure of the problem. Thus, the power of the representation is not in the particular form per se, but in how well the structure of the configural representation maps to the deep structure of the work or problem being represented.

Figure 7 shows an initial chart summarizing some of the properties that might be associated with the 'deep structure' of the anemia problem. Anemia typically results from either blood loss, destruction of red blood cells, or decreased or faulty red blood cell production. Two important sources of information for identifying the nature of the anemia problem is the number of reticulocytes and the mean corpuscle volume. Reticulocytes are immature blood cells and a high count typically indicates either a loss of blood or premature destruction of red blood cells. A low count typically indicates a problem with red blood cell production. Mean corpuscular volume (MCV) is the average volume of red blood cells. The normal range is 80-100 fL. Different diseases are associated with different health problems. For example, liver disease is typically associated with an MCV that is larger than normal (macrocytic). Renal disease is associated with normal MCV (normacytic). Iron deficiencies are typically associated with smaller than normal MCV (microcytic).

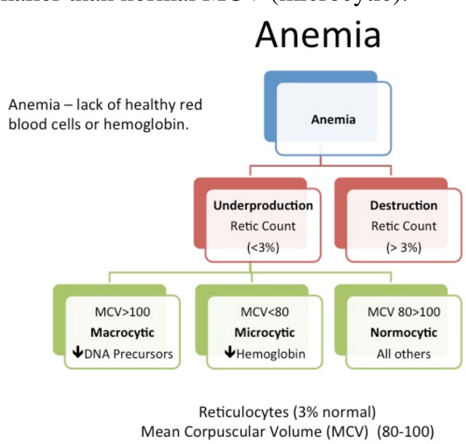

CONCLUSION

The goal for this paper was to introduce CSE and EID to the healthcare community. We believe that CSE is an essential complement to more classical human factors approaches for domains where the role of humans cannot be reduced to procedural compliance or simple rule following. In such domains it is critical for smart humans and smart technologies to work together to adapt procedures to highly variable situations and contingencies that could not be fully anticipated in advance. There is little doubt that healthcare is such a domain. Yet, there is little evidence that CSE is informing the development of healthcare technologies such as EHR systems.

EID is a challenge for interface designers to look beyond human-computer interaction to consider the pragmatics of human-work interaction. The goal of EID is not to 'match' the user's mental model or to simply provide access to data. Rather, the goal of EID is to design representations that make the deep structure of the work domain 'tangible' in order to shape mental models in ways that provide insights into complex problems and that support productive thinking (Wertheimer, 1959; Woods, 1991; Bennett & Flach, 2011).

We have provided three different examples, in the hope that the different perspectives will help the readers to see beyond the surface features of particular interfaces or representations and to better appreciate that it is not the specific form of the representations, but the mapping of the form onto the deep structure of a problem that makes a representation useful. Thus, for example, the representations developed for CVD and managing pain are very different. However, the principle of 'structural mapping' is informing both designs.

The development and integration of EHR systems into the healthcare domain provides an excellent opportunity for healthcare to build on the discoveries that CSE professionals have made in other similar domains, where we depend on people to make smart choices in the face of complex situations. The EHR systems offer the computational power and the display capabilities to create innovative representations that can allow healthcare professionals to 'see' and 'explore' problems in ways that were not possible before (Shneiderman, et al. 2013). The ultimate question is whether healthcare will take advantage of the lessons learned in these other domains, or will it be necessary for healthcare to make all the mistakes associated with clumsy automation (Weiner, 1988), in order to discover for themselves the value of a CSE perspective and the power of representations that are designed to make the deep structure of problems tangible (EID).

REFERENCES

Bennett, K.B. & Flach, J.M. (2011). *Display and Interface Design: Subtle Science, Exact Art.* London: Taylor & Francis. ISBN-13: 978-1420064384

Borst, C., Flach, J.M. & Ellerbroek, J. (2015). Beyond ecological interface design: Lessons from concerns and misconceptions. *IEEE: Transactions on Systems, Man, and Cybernetics*, 45(2), 164-175.

Burns, C.M. & Hajdukiewicz, J. (2004). *Ecological interface design.* Boca Raton, FL: CRC Press.

Chobanion, A. V., Bakris, G. L., Black, H. R., Cushman, W. C., Green, L. A., et al. (2003). Seventh report of the Joint National Committee on prevention, detection, evaluation, and treatment of high blood pressure. *Hypertension*, 42, 1206-1252.

D'Agostino, Sr, R. B., Vasan, R. S., Pencina, M. J., Wolf, P. A., Cobain, M., Massaro, J. M., & Kannel, W. B. (2008). General cardiovascular risk profile for use in primary care: The Framingham heart study. *Circulation, 117,* 743-753. DOI: 10.1161/CIRCULATIONAHA.107.699579

Elder, N., *McEwen, T.R.*, Flach, J.M. & Gallimore, J.J. (2009). The management of test results in family medicine offices: Complexity and quality. *The Annals of Family Medicine,* 7(4), 343 - 351.

Eriksson, K.A. & Charness, N. (1994). Expert Performance. *American Psychologist,* 49(8), 725-747.

Flach, J.M. & *Dominguez, C.O.* (1995). Use-centered design. *Ergonomics in Design*, July, 19 - 24.

Flach, J.M. & Voorhorst, F.A. (2016). *What matters?* Dayton, OH: Wright State University Library.

Klein, G. (1989). Recognition-primed decisions. In W.B. Rouse (ed.) *Advances in Man-Machine Systems Research.* Greenwich, CT: JAI Press.
McEwen, T.R., Flach, J.M. & Elder, N.C. (2014). Interfaces to medical information systems: Supporting evidence-based practice. *IEEE: Systems, Man, & Cybernetics Annual Meeting,* 341-346. San Diego, CA. (Oct 5-8).
McEwen, T.R., Flach, J.M., and Elder, N.C. (2012). Ecological interface for assessing cardiac disease. *Proceedings of the ASME 2012 11th Biennial Conference on Engineering Systems Design and Analysis*, ESDA2012, July 2-4, 881-888.Nantes, France. ASME ESDA2012-82974.
McKibbon, K.A. (1998). Evidence-based practice. *Bulletin of the Medical Library Association,* 86(3), 396-401.
Naikar, N. 2013. *Work domain analysis.* Boca Raton, FL: CRC Press.
National Heart, Lung, and Blood Institute. (2001). Executive Summary of the Third Report of the National Cholesterol Education Program (NCEP) Expert Panel on Detection, Evaluation, and Treatment of High Blood Cholesterol in Adults (ATP III). Expert panel on detection, evaluation, and treatment of high blood cholesterol in adults. *Journal of the American Medical Association, 285,* 2486-2497.
Norman, D.A. (1986). Cognitive Engineering. In *User Centered System Design*, ed. D.A. Norman and S.W. Draper, 31–61. Hillsdale, NJ: Erlbaum.
Rasmussen, J. (1986) *Information processing and human-machine interaction: An approach to cognitive engineering.* New York: Elsevier.
Rasmussen, J. and Vicente, K. J. (1989). Coping with human errors through system design: Implications for ecological interface design. *International Journal of Man-Machine Studies* 31:517–534.
Schatman, M.E. & Fudin, J. (2016). The myth of morphine equivalent daily dosage. *Medscape,* May 24.
Schrage, M. (1999). *Serious Play.* Brighton, MA: Harvard Business Publishing.
Shneiderman, B. 1992. *Designing the user interface: Strategies for effective human computer interaction.* Reading, MA: Addison-Wesley.
Shneiderman, B., Plaisant, C. & Hesse, B.W. (2013). Improving healthcare with interactive visualization. *Computer,* May, 58-68.
Vicente, K.J., Christoffersen, K. and Pereklita, A. (1995). Supporting operator problem solving through ecological interface design. *IEEE Transactions on Systems, Man, and Cybernetics, SMC-*25: 589-606.
Vicente, K. J. (1999). *Cognitive work analysis.* Mahwah, NJ: Erlbaum.
Wertheimer, M. (1959). *Productive thinking*. New York, NY: Harper and Row.
Wiener (1988). Cockpit automation. In E.L. Weiner & D.C. Nagel (eds.) *Human Factors in Aviation.* San Diego, CA: Academic Press.
Woods, D. D. (1991). The cognitive engineering of problem representations. In *Human-computer interaction and complex systems*, ed. G. R. S. Weir and J. L. Alty, 169-188. London, UK: Academic Press.