The Foundations and Pragmatics of Cognitive Work Analysis

A Systematic Approach to Design of Large-Scale Information Systems

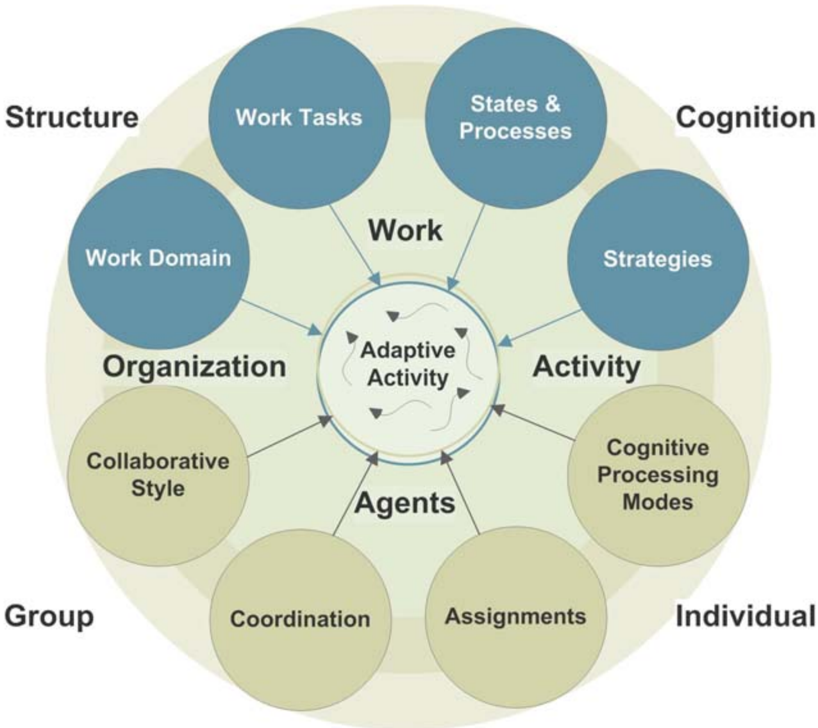

Gavan Lintern

Cognitive Systems Design

www.CognitiveSystemsDesign.net

Edition 1.0

Cover Graphic

My inspiration for the cover graphic evolved from a figure developed by Penny Sanderson and published in Sanderson, Naikar, Lintern and Goss (1999). Penny had been influenced by other figures from Rasmussen, Pejtersen and Goodstein (1994) and Vicente (1999). These figures capture the fundamental idea behind Cognitive Work Analysis; that we must orient socio-technical design towards supporting adaptive human activity. That orientation requires consideration of four sets of constraints; those provided by the structure of the work domain and work tasks, those provided by cognitive processing and cognitive strategies used in transition between cognitive states, those associated with individual work task assignments and cognitive processing modes, and those associated with group coordination and collaborative style. The inner ring of labels parses the space in terms of work constraints (upper half), agent constraints (lower half), organizational constraints (left half), and activity constraints (right half). Each stage of Cognitive Work Analysis deals with one or more of these sets of constraints and Cognitive Work Analysis, in its entirety, deals with all of them.

The three images below are the ones from the named publications that inspired development of the cover graphic.

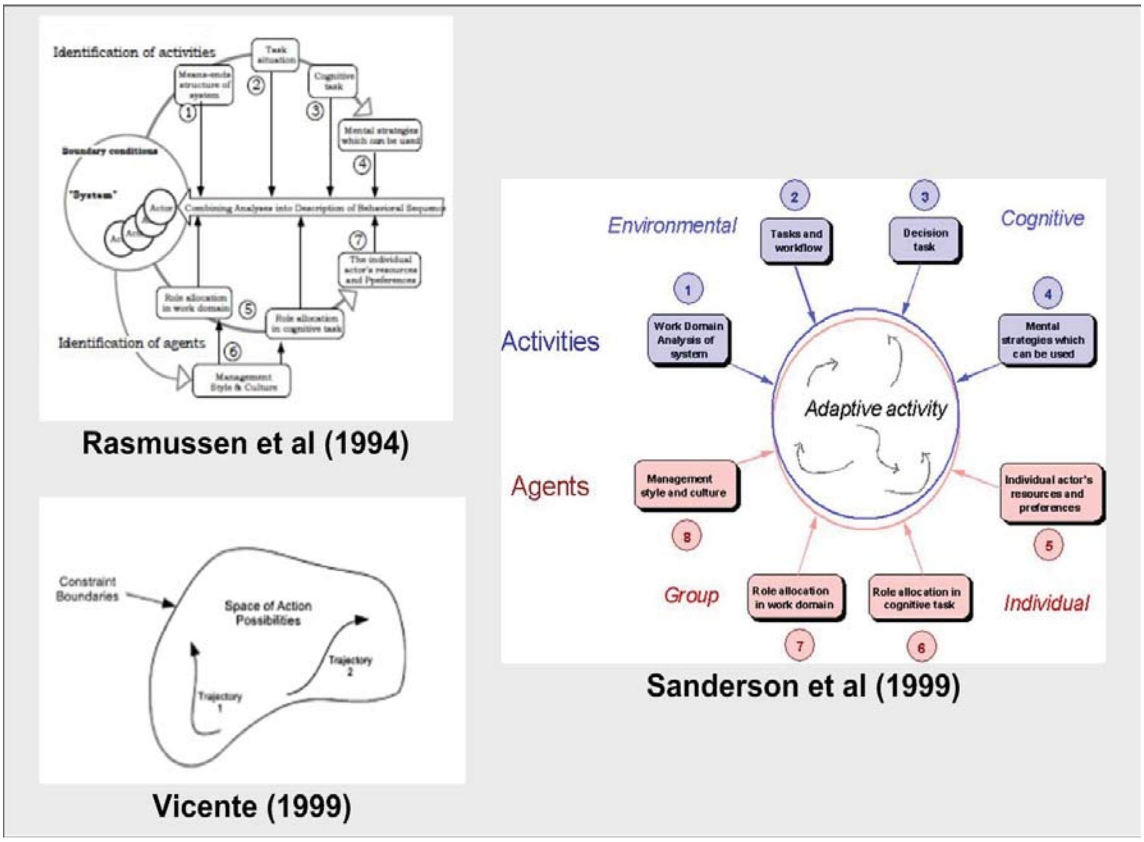

Table of Contents

Preface

I offer this book as a free download from my website at www.CognitiveSystemsDesign.net. I suppose, if I had thought it might be as popular as Dan Brown's "The Da Vinci Code ", I may have been tempted to market it commercially, but the fact is that this is a specialist topic and even if the book is wildly successful in this topic area, it will gain relatively few readers. I value readership above the rather paltry royalties that might accrue and so I hope to encourage readership by making this book as accessible as possible.

I request that you direct others to my website for their copy rather than giving it to them directly. If everyone adheres to that request, I will be able to track how many copies are in circulation and thereby assess whether this project has been successful and whether I should do it again for another topic. However, feel free to break this rule if your colleague would otherwise experience difficulty or delay in downloading a copy.

I was first attracted to the electronic self-publishing route because it is difficult and time-consuming enough to write a book without the further complications of negotiating with a publisher and the delays that accompany printing and marketing.

Electronic self-publishing has a further significant advantage; I can change the book at any time and at no particular cost. I will be adding chapters and appendices as time goes by and when I do, I will identify what I have changed so that you will be able to go straight to the new material.

One useful feature of electronic publishing, at least in the Portable Document Format, is that PC users can access the dictionary meaning of any word via http://dictionary.reference.com by right-clicking on that word and selecting “Look Up” from the drop-down menu. I assume Mac users can do this as well but I am not clear on how they can do it.

When you read this in Portable Document Format you may use the hyperlinks in the Table of Contents to jump directly to a desired chapter. Hyperlinks are indicated by a light-blue underline. The Portable Document Format also obviates the need for an index. Use the search function to find references within the book to any term. I have included page numbers in the Table of Contents primarily for the convenience of those who will want to read a paper copy.

Acknowledgments

I can never ponder issues related to socio-technical systems without reflecting on ideas I have discussed with Peter Kugler. Our conversation on these issues commenced in 1988 and continues today.

I struggled in my early encounters with Cognitive Work Analysis and was assisted considerably at that time through conversations with Cathy Burns and Kim Vicente.

I have had many long and detailed conversations with Neelam Naikar and Penny Sanderson, not always devoted to issues in Cognitive Work Analysis but always engaging and productive.

Wide-ranging discussions with Iya Whitely (nee Solodilova) and Anne Bruseberg continue to guide me into new areas.

I have been concerned for many years about the gulf between Cognitive Task Analysis and Cognitive Work Analysis and have sought to bring ideas from the first into the practice of the second. My ongoing conversations with Cindy Dominguez, Laura Militello and Gary Klein have helped enormously.

I have also become concerned in recent years about the seemingly insurmountable barriers between Cognitive Engineering and Systems Engineering. My ongoing conversation with Steve Deal has helped me frame ideas about how to attack those barriers.

This book would not have been possible without the support and forbearance of my wife, Miyuki Chikamatsu and my daughter, Anna Lintern.

Chapter 1

Introduction

Cognitive Work Analysis is notoriously difficult for those who encounter it for the first time. It is a complicated and expansive system of analysis, differing in scope and strategy from much of what currently goes on in cognitive engineering. There is little to do about this; the system is what it is for good reasons. Given that state of affairs, we need cohesive, pedagogical accounts of this analytic framework to guide beginners through their early efforts. Vicente (1999) has made good progress in this regard but much remains to be done. In this book I take a different but complementary approach to that taken by Vicente to introduce beginners to Cognitive Work Analysis.

I seek to resolve two issues. Cognitive work analysis remains difficult to understand and to execute because we have not made the foundational theory behind it sufficiently explicit and also because we have not been sufficiently tutorial in our approach to explaining it. In believing that these two things go together, I outline the theoretical basis for this framework of analysis and then offer a tutorial example that shows how the concepts can be applied. In future editions of the book, I will offer further tutorial examples as appendices.

Although I offer some refinements of Cognitive Work Analysis, there is nothing fundamentally new in this book. Rather, this is an effort to assemble the important ideas of Cognitive Work Analysis into a treatment that encourages solid understanding via a process of establishing specific concepts as knowledge anchors and then expanding that knowledge into a comprehensive system.

The Title of the Book

I have chosen the title of this book with deliberation. The book is centrally about Cognitive Work Analysis. The *foundations* are specifically the theoretical foundations. In chapter 2, I offer a brief account of several theoretical positions that establish a context for Cognitive Work Analysis. I do not necessarily want to claim that these theoretical positions have guided the development of Cognitive Work Analysis but that rather the assumptions on which those positions are based and the observations that have emerged from them are consistent with and offer support for the framework. In succeeding chapters, I outline the specific theoretical assumptions for each of the analyses that make up Cognitive Work Analysis.

I have chosen to insert the word *pragmatics* into the title because, ultimately, Cognitive Work Analysis is a practical framework for developing a coherent and comprehensive description of the important properties of work. My dictionary, Houghton Mifflin (2000) defines pragmatism as a *practical, matter-of-fact approach to assessing situations or solving problems*.

Words

Words can be difficult. Many have multiple meanings and scientific and engineering usages often extend beyond the definitional boundaries contained in dictionaries. Vicente (1999) offered definitions for many of the troubling words he uses. I have gone back to those definitions time and time again. Many of the criticisms of Vicente's book have emerged from failure to understand how he was using specific terms; a failure that is inexcusable given that he defined the meanings of those terms explicitly. I also devote some effort in this book to defining my terms. Redefinitions of words first defined by Vicente will typically be clarifications rather than adjustments in meaning but I also add a few important terms to the lexicon of Cognitive Work Analysis.

In converging on a word meaning, I rely heavily on Wiktionary.org, Dictionary.com and Houghton Mifflin, 2000, but explain further when these references offer multiple meanings or where words have crept into common scientific usage with a meaning not implied by any dictionary. I avoid any usage that cannot be found in a pedigreed dictionary. I neither invent new meanings for words nor accept invented words or invented meanings unless there is a sound reason, as explicated by the inventor, for that term.

I abhor the current tendency in science and technology to sprinkle acronyms extensively throughout a narrative. Acronyms can make even simple ideas difficult to assimilate. Indeed, they require a reader who is new to the material to learn a new language. The practice of naming the acronym in its first use does not help very much. A reader will not necessarily remember it, especially if there are many other acronyms in an extended discussion. Many times, a reader will gloss over an acronym as something only half understood. I see no excuse for them and in this book I tolerate only those acronyms that are in general and widespread usage.

Chapter Titles

I have been troubled for some time by the misinterpretations of Cognitive Work Analysis that abound in the literature, for example, Lind (2003) on Work Domain Analysis, Cummings (2006) on the temporal implications of Work Domain Analysis, and Hollnagel and Woods (2005) on the Decision Ladder and the Abstraction Hierarchy. There is, it seems, no end of mischief that can be created by those who fail to understand.

As all readers of this book know, or at least will soon know, Cognitive Work Analysis has several analytic stages. It has occurred to me that the critiques of Cognitive Work Analysis noted above were encouraged by Vicente's strategy of using the names of the analyses for the titles of chapters in which he explained that analysis. I draw that conclusion because, without exception, those critics failed to show any understanding of the assumptions underlying that stage of the analysis they were critiquing and typically focused on analytic details. Not one of those critics offers even a glimmer of understanding of what is to be achieved in that analytic stage.

Although names of analytic stages seemed appropriate as chapter titles in 1999, I now think that a chapter title that identifies the purpose of the analytic stage will serve us better. I suggest that such a strategy will make it more difficult for critics to focus on peripheral issues. I map the correspondence between Vicente's chapters on analytic forms and the ones I use here in Table 1.1. As will be evident from Table 1.1, I have adjusted the stage sequence offered by Vicente (1991) and have added one stage. The additional stage, Stage 2 in my treatment, comes from a development by Naikar, Moylan, and Pearce (2006) of a Contextual Activity Matrix to depict the relationship of Work Problems to Work Situations.

The reasons behind my selection of each specific chapter title as a descriptor of what is to be achieved by the analysis described in that chapter will, I hope, become evident early in each of the respective chapters. I retain only two of the analysis titles used by Vicente (1991). The reasons that I have chosen the particular analysis titles that I have will also, I hope, become evident early in each of the respective chapters.

Vicente (1999) Chapter Titles	Lintern (2009) (this book) Chapter Titles	Lintern (2009) (this book) Name of Analysis
1. Work Domain Analysis	1. Functional Work Structure	1. Work Domain Analysis
	2. Partitioning & Organization of Work	2. Work Organization Analysis
2. Control Task Analysis	3. Cognitive States & Processes	3. Cognitive Transformations Analysis
3. Strategies Analysis	4. Cognitive Strategies	4. Strategies Analysis
5. Worker Competencies Analysis	5. Cognitive Processing Modes	5. Cognitive Processing Analysis
4. Social Organization & Cooperation Analysis	6. Work Coordination	6. Social Transactions Analysis

Table 1.1: Correspondence between Vicente' (1999) chapter titles and the chapter titles and names of the respective analyses as used in this book

A Niche for Cognitive Work Analysis

Some of the negativity towards Cognitive Work Analysis emanates, I believe, from a failure to understand what those of us who use Cognitive Work Analysis are trying to do. Most techniques of Cognitive Engineering are aimed at identifying and working on points of leverage, for example, on developing cognitive support tools in the form of such things as decision aids and planning support. In contrast, the framework of Cognitive Work Analysis was developed for a much larger problem; the design of large-scale socio-technical systems. Despite the value of other cognitive engineering strategies, they deal only with segments of the design problem for a complex socio-technical system. I do not intend that remark to be pejorative; many design assignments in cognitive engineering require precisely that form of intervention. My specific claim here is that Cognitive Work Analysis occupies a niche in the design world that is often not appreciated by those who focus on points of leverage or on the development of a cognitive support tools.

A long-standing complaint within Human Factors is that we (Human Factors practitioners in general) are brought into the design of large-scale systems only after human integration problems have become apparent. It is commonly argued that the expense of correcting these problems could be avoided if we were consulted earlier, possibly during concept development and then throughout the remainder of the design cycle. In the past, I have been skeptical. While I was confident that we could have avoided the sorts of common problems that were emerging, it was never clear to me that we would not have introduced other serious issues. We had no comprehensive analytic framework for addressing issues in concept development and then proceeding systematically through the human systems integration issues in the design of a complex socio-technical system. Only when I became acquainted with the framework of Cognitive Work Analysis did I begin to build confidence that I could, if invited early into the design process, contribute as an equal an effective partner.

I also wonder if the emphasis within Cognitive Work Analysis on representation introduces some negativity. Crandall, Klein and Hoffman (2006, p 107) note that knowledge elicitation has received more attention that knowledge representation within the general field of Cognitive Task Analysis. I suspect that formal education plays some role in determining enthusiasm for representation. The engineering disciplines employ representation extensively and systematically in many forms to impart understanding. In contrast, representation is employed less often within the behavioral sciences and then in an improvised and impromptu fashion.

From my own background in Psychology, my initial reaction to Cognitive Work Analysis was that it was *only* about representation (and therefore insubstantial). I adjusted that thought rather rapidly as I read further but continue to believe that Cognitive Work Analysis is largely about representation. However, I no longer use the pejorative *only* when I offer that view. I have come to believe firmly in the power of a theoretically motivated and well organized set of representations for assimilating, archiving and transferring knowledge.

Finally, Vicente (1999) emphasizes the activity-independent property of the Abstraction-Decomposition Space, the representational product of Work Domain Analysis. I suggest that this notion of activity independence troubles many people. Both Cognitive Psychology and Systems Engineering are process or activity oriented disciplines. Cognitive theories are typically framed as a series of processes or activities and Functional Analysis in Systems Engineering typically results in a representation of functional flow rather than a representation of functional structure.

Gibson (1979) is one in psychology who has taken this notion of activity independence seriously. I recall that his approach troubled me as I work through the first two chapters of his *Ecological Approach to Visual Perception*. My first thought was that this was *only* about the structure of the world and that there was no psychology in it. Again, note the pejorative *only*. However, one should not judge Gibson prematurely as I did then and, I suspect, as many others do. One has to get through the complete argument to appreciate its elegance.

By the time I encountered Rasmussen's work, I had assimilated Gibson's argument and did not for a moment cast the same aspersion. I do recall thinking that Rasmussen's distinction between structure and process was much like Gibson's. If we were to take the critics seriously, we would have to assume that Gibson and Rasmussen are alike in that they have built a flawed conceptual structure from a fundamental misunderstanding of the nature of the world. I happen to think otherwise, that each in their own way, immersed in a conceptually challenging and somewhat distracting intellectual culture, somehow came to remarkable insights about the way we need to conceptualize complex human environments.

Rasmussen was concerned with how to integrate multiple, diverse technical capabilities with human capability at many levels of organisation into a cohesive socio-technical system. As I note above, he was largely concerned with representation. Although he mentions the manner in which he collects data to populate his representations, those methods do not constitute an innovative contribution. In addition, while there is some discussion of how to use these representations for design, that too remains relatively undeveloped. Many others who employ the framework of Cognitive Work Analysis have made contributions in these other areas but my emphasis in this book is on the representational framework although I will, in a later editon of this book, devote a chapter to the design problem.

Chapter Summary

First and foremost, this book is a tutorial. It will have served its purpose if you, as reader, generate insights that help you understand what Cognitive Work Analysis is about. I seek to help you generate those insights by linking theory to illustration. Each of the chapters devoted to method outlines the theoretical basis for that method and then illustrates the method with an example that I hope will be easily understood by all. Additionally, I link the successive stages explicitly and illustrate how each stage not only provides information for design but also sets up the next stage.

Chapter 2

Theoretical Perspectives

The Nature of Theory

Opinions on what constitutes a theory are diverse. Sometimes I see summaries of structure, as derived through the application of taxonomic methods, characterized as theory. Sometimes those summaries incorporate relational statements as might be derived through the application of ontological methods. Indeed, an Abstraction-Decomposition Space, which is the representational product of Work Domain Analysis, is developed through use of ontological methods. The Abstraction-Decomposition Space is not, however, a theory, although I will later argue that the way we build one is guided by a pragmatic theory of reasoning.

In addition, I have occasionally encountered the opinion that Gibson's ecological approach is not a cognitive theory because it does not posit an internal cause-effect mechanism. I take issue with that opinion on two counts. I suspect that the author of a comment like this is demanding a linear action-reaction event such as a cue striking a billiard ball and that behavior shaping constraints will not serve. In addition, some appear to take Gibson's view that much cognitive activity unfolds beyond the central nervous system as a claim that there are no cognitive structures or processes within the central nervous system. That is, however, an incorrect reading of Gibson. His discussion of resonance to information is just one example of his concern with internal cognitive processes[2.1].

Given this sort of uncertainty, it is worth offering an opinion on what sort of properties a theory for Cognitive Work Analysis should capture. My dictionary (Houghton Mifflin, 2000) defines a theory as a set of statements or principles devised to explain a group of facts or phenomena or that can be used to make predictions about natural phenomena. This will do for the current purposes, although let's be cautious with the reference to *explanation;* there can be dissension about what constitutes an explanation versus a description. However, I conclude from this definition that by use of theory, scientists seek to make sense out of regularities they observe in natural phenomena.

2.1 The view that resources and processes external to the body can be characterized as cognitive has far more currency today (e.g., Hutchins, 1995; Hollan, Hutchins & Kirsh, 2000).

Also note the use of the word *devised* in the definition. A theory is not a statement of fact but rather an imaginative construction. The test of a theory is not whether it is true versus false but whether it helps us understand the world in useful ways. Despite being an imaginative construction, theories can be powerful. The theory of gravity, for example, is a relatively simple statement that takes account of a diverse set of natural phenomena. Despite its simplicity, it has remarkable power. There is presumably no one reading this book who doubts that gravity will have its way on every location of our (and every other) planet. We believe that for locations we know, for places we have never been, for places we will never visit and even for places we have never even heard about.

There is a tendency, within behavioral science, to envy physicists. They study (or at least used to study) observable and a regular phenomena. Such envy is unnecessary: Cognitive Work Analysis is based on observable and regular behavioral phenomena that can impart considerable power to our analysis and design activities.

Foundational Perspectives

The foundational perspectives I outline in this section did not necessarily guide developments in Cognitive Work Analysis, but the concepts they have established represent core assumptions for an analysis and design strategy based on Cognitive Work Analysis.

Situated Cognition

The ethnographic research by Hutchins (1995), Jordan (1989), Lave (1988), Lave and Wenger (1991), Saxe (1991), Scribner and Fahrmeier, (1982) and Suchman (1987) offers profound insights. It reveals how adept workers can be at cognitively restructuring their work environment. Invariably, the work practices that evolve are cognitively economical and robust, typically more so than work practices prescribed by those who do not actively participate in the work.

I have reviewed a portion of this work for its relevance to aviation (Lintern, 1995). One lesson to be taken from it (for aviation and more generally) is that workers are both physically and cognitively active, reshaping how they think about their work environment as they develop their own work practices. The conceit of managers, and also of many designers, is that they know how the work should be accomplished and they need to instruct workers in the proper procedures. The ethnographic research on situated work practice reveals that conceit to be shallow.

In thinking about this issue, I reflect on developments in Artificial Intelligence. There are an enormous number of computationally-based support systems that would seem to offer huge advantages to current practices. Diagnostic systems for medical practitioners can serve to illustrate. These have been under development for decades but are still struggling to find their way into common usage within the medical profession. It is not unusual to hear the accusation that medical practitioners are too arrogant to embrace technology that might replace some of their skills. I suspect otherwise; that these systems do not mesh well with the cognitive strategies and work flow of medical practitioners. From that perspective it would seem that it is the designers of these systems who are overly arrogant.

Implications for analysis and design. The research in Situated Cognition indicates that we need to be very careful if we, as designers, specify cognitive strategies or work flow. The existing strategies and processes will have evolved over considerable time to be robust and effective. To change them without fully understanding the potential repercussions is to risk disaster. Furthermore, workers are adept at modifying strategies to accommodate the demands of new systems. Thus, we should ensure we understand how practitioners or experts go about their work (strategies, modes of cognitive processing) so that we can design supports for their work practices without imposing awkward strategies and we should resist the temptation to over-design systems (we should permit workers to finish the design, Vicente 1999).

Distributed Cognition

Within the work environment of ship navigation in confined waters, Hutchins (1995) reiterated many of insights to be drawn from Situated Cognition but added a particularly evocative and succinct description of distributed cognition. Up to that time, distributed cognition was a somewhat fuzzy concept that even experts in the field would debate.

Hutchins proposed that a ship navigation team, together with accompanying navigational artifacts and procedures, is a cognitive system that performs the computations underlying navigation. It is a distributed cognitive system because various elements of the computations are carried out over time and in different locations. The results of early computations are passed to another location and then integrated in further computation. Hutchins argued that this navigation system has cognitive properties that differ from the cognitive properties of the individuals within the system and that the cognitive potential of the navigation team depends as much on its social organization as on the cognitive potentials of its members. Thus the navigational system performs computations that need not necessarily be within the grasp of all (or even any) of its members.

The theory of distributed cognition forces a shift in how we think about the relationship between minds, social interactions and physical resources. Interactions between internal and external processes are complex and unfold over different spatial and temporal scales and neither internal nor external resources assume privileged status.

Implications for analysis and design. Most, if not all socio-technical systems we design will be distributed. As revealed in the illustration offered by Hutchins (1995), there is need for coordination between the distributed subsystems. We need to examine how people coordinate (share information, communicate, work collaboratively) and then we need to ensure that our designs support the essential modes of coordination.

Requisite Variety

Vicente (1999) appeals to the law of requisite variety in arguing that the complexity of a technological support needs to reflect the complexity of the work. Ashby (1957, p207), in framing this law, proposed that *only variety can destroy variety*, here taken to mean that only variety can control variety. In other words, a control system must incorporate as much variety as the system it controls. Alternatively, the functional scope and granularity of a work space must match the operational complexity of the work.

The law of requisite variety warns us against seeking to reduce control complexity by simplifying displayed information. Hollnagel and Woods (2005, p 85) also warn against this, but mis-characterize Ecological Interface Design as a strategy that reduces complexity. An ecological interface, when properly designed, will give selective information access at the level of complexity required for the anticipated control problem. An ecological interface does not continuously display all information at the most detailed levels as for example does a Single-Sensor/Single-Control strategy, but rather displays patterns that can be selectively interrogated to reveal information for the control problem at the essential level of detail.

A pentagon display for a social system governed by human intentions (Figure 2.1) offers a simple illustration. Five dimensions that contribute to the global construct are represented by individual spokes of the pentagon. The measures of those dimensions are normalized to show a symmetric figure under normal or desirable conditions. Where a particular measure reflects an abnormal or undesirable condition, the spoke for that measure generates a distortion in the figure. That distortion will be noticed readily and the offending dimension identified. The relevant spoke can then be interrogated (via mouse click, for example) to foreground more diagnostic detail about the issue. By this means, requisite complexity is avoided until it becomes relevant and only that portion of the requisite complexity needed for the current situation is displayed. Nevertheless, the entire requisite complexity for the system is available.

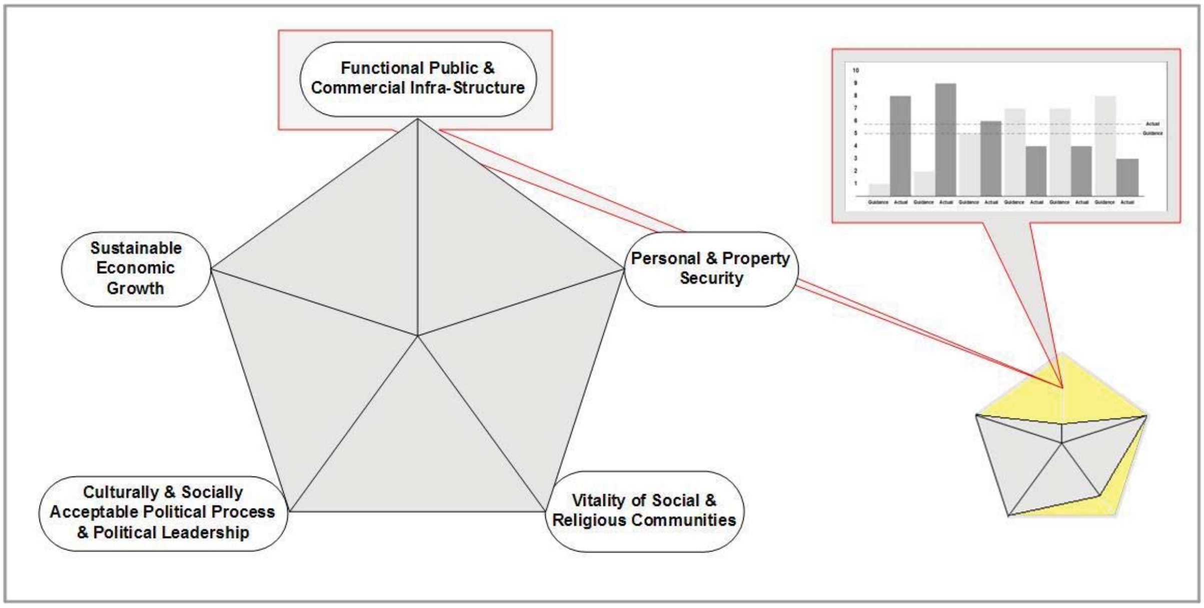

Figure 2.1: A pentagon representation for a social system governed by human intentions (assessment of progress in building a civil, democratic society): the outset of the pentagon at bottom right shows a distortion for Functional Public & Commercial Infra-Structure, the details of which can be displayed in a histogram that compares values for desirable and actual supply of Electricity, Gas, Telephone, Water, Gasoline and Public Transport.

Implications for analysis and design. In development of complex socio-technical systems, we need to ensure that the information potentially available to a worker matches the requisite variety of the work to be undertaken and we need to ensure that workers can find and assemble the constellation of information needed for the problem at hand.

Ecological Psychology

The foundational insight of Ecological Psychology is that cognition is tied up in the reciprocity between an organism and its environment. As with Situation Cognition and as stressed particularly by Hutchins (1995) and his colleagues (Hollan, Hutchins & Kirsh, 2000), much of cognition occurs in the world rather than in the head. As noted above, some critics of Ecological Psychology have taken this to be a claim that nothing relevant to cognition happens in the head, although I remain mystified how any reasonably competent reading of Gibson can lead to that conclusion.

The theory of affordances is a major contribution from Gibson that is relevant to our work. An affordance is a relationship between properties of an organism and matching properties of its environment. It is a relationship between capability and opportunity. In an explanation of the relevance of affordances to interface design (Lintern, 2000), I drew on work by Warren and Whang (1987) who discussed the relationship between shoulder width and aperture width as a passing-through affordance. Indeed, an affordance is always a relationship. Where the dimensions of that relationship can be quantified, it can be expressed as a dimensionless ratio.

Thus, an affordance-based fuel gauge compares the distance that needs to be traveled to the distance that can be traveled with currently available fuel. Depending on which way the ratio is constructed, a value of more or less than one will signify that you can (or cannot) get to your destination. This strategy removes from the operator the computation that is required when fuel and distance are presented separately. Gibson's affordance claim is that this is analogous to the way we operate in the world.

I suspect that many take an ecological display to be one that is pictorial or richly graphical. However, for a display to be ecological, it needs to be more than that. It needs to incorporate within its graphics a depiction of structure of the work environment in terms of the affordances essential to the work.

Implications for analysis and design. An affordance establishes meaning by revealing the reciprocity between information and action. We need to ensure that the information potentially available to a worker is meaningful and we can do that by examining how workers use information or, in other words, by analyzing their affordance structure. Also, recall that the law of *requisite variety* implies that the functional scope and granularity of a work space must match the operational complexity of the work. The law emphasizes that the degree of complexity must be equivalent but does not emphasize that the semantic structure of the workspace must match the semantic structure of the work. The appeal to the concept of affordances corrects that neglect.

Self Organization

Self-organization is a process in which the internal organization of a system, normally an open system, increases in complexity without being guided or managed by an outside source. Self-organizing systems typically (though not always) display emergent properties. http://en.wikipedia.org/wiki/Self-organization [Accessed Nov 16 2007]

Theories that posit a mental image, a mental model or a mental schema as a formative cause of cognition eschew self-organization (e.g., Johnson-Laird, 1983). In contrast, others argue that an understanding of self-organization is central to understanding cognition. Peter Kugler and I have summarized the basic concepts for this latter view in Lintern and Kugler (1991) and Kugler and Lintern (1995).

Self-organizing systems can transition through a non-linear region into a different (and sometimes, unexpected) organizational form. As explained by Prigogine and Stengers (1984), an adjustment of a control parameter can generate critical fluctuations that cast the system into a new energetic mode. In the case, for example, that the change in the control parameter is serving to inject more energy into the system, a point is reached at which the system must reorganize to dissipate that energy. The term *dissipative structure* is often used to characterize the new organizational form.

The patterns of locomotion for a horse offer an illustration. As the rider nudges the horse into increasing its speed, the horse will initially increase the rates of limb motion but at some critical point will transition into a new, more efficient mode (e.g., canter to gallop). As is true of all nonlinear systems, equine locomotion is mostly linear. However, it is at the nonlinear transitions that interesting things occur. It is invariably true that nonlinear systems can appear linear if one takes a restricted view.

Those who promote self-organization as an explanation of cognition typically emphasize the role of local interactions in the development of patterns and might offer self-organization as a bottom-up, emergent view in contrast to the top-down view of mental imagery as the shaping influence on cognition. Some caution is needed here. While local constraints play an important role in self-organization, it is the interplay between local and global constraints that generate the new patterns.

This means that cognitive emergence owes as much to the functional layout of the environment as it does to the local interactions of individuals with each other and with artifacts. The cognitive architecture determines the way information flows through the system. This architecture encompasses the functional structure of the physical environment, the social organization of the work place and the functional structure of individual minds. New cognitive capabilities emerge from activity undertaken within the constraints imposed by the cognitive architecture and are shaped by those architectural constraints.

Note also that it can be difficult to anticipate the form of organization that will emerge after transition through a nonlinearity. It is never possible to do that from an understanding of the underlying mechanism. However, those experienced with a particular system can typically anticipate the forms of organization that will emerge at least within the range of their previous experience. Experienced equestrians can certainly anticipate the form of organization that will emerge as a horse increases or decreases its pace.

Engineers abhor nonlinearities but biology cannot survive without them. In cognitive engineering, we have a subtle problem. We need to conjoin system components that have been designed with linearity as a design goal to other system components (i.e., the human operators) for which nonlinearity is fundamental. The techno-centric approach is to force linearity on our nonlinear human work force; to suppress the self-organizing tendencies of human systems. These self-organizing tendencies are, however, critical to the system effectiveness (Lintern, 2003). The human-centric approach we seek in cognitive engineering is to work with (even to celebrate and to leverage from) processes of self-organization to make our systems more effective.

Implications for analysis and design. In the design of a large-scale information system we must remain concerned with the functional properties that constrain (or-shape) the possibilities for courses of action and the informational interactions that stimulate emergent patterns of action. We need to examine how both intentional as well as technological properties establish a functional structure that can shape cognition. In addition, we need to understand how interactions between people and interactions between people and technological subsystems generate emergent patterns of behavior as the basis for social organization and teamwork within the workplace. Note that for new designs, anticipation of emergent behavioral patterns demands careful analysis of similar systems if they exist but may otherwise require systematic evaluation with a full range of realistic operational scenarios.

Cognitive Systems

A cognitive system is a one that performs the cognitive work of knowing, understanding, planning, deciding, problem solving, analyzing, synthesizing, assessing, and judging as those activities are fully integrated with perceiving and acting. A complex socio-technical system is an entity that does cognitive work and is therefore a cognitive system.

The claim that a complex socio-technical system does cognitive work expands the view of what is *cognitive* beyond the individual mind to encompass coordination between people and their use of resources and materials. This view is aligned with the theory of distributed cognition enunciated by Hutchins (1995) and further described by Hollan, et al (2000). A foremost claim of this theory is that distributed cognition is not a special type of cognition but is rather a characterization of fundamental cognitive structures and processes (Hollan et al, 2000). Thus, all cognition is distributed.

Traditionally, we are used to thinking of cognition as an activity of individual minds but from the perspective of distributed cognition it is a joint activity that is distributed across the members of a work or social group and across the technological artifacts available for support of work[2.2]. Cognition is distributed spatially so that diverse artifacts shape cognitive processes. It is also distributed temporally so that products of earlier cognitive processes shape later cognitive processes. Most significantly, cognitive processes of different workers interact so that synergistic cognitive capabilities emerge via the mutual and dynamic interplay resulting from both spatial and temporal coordination among distributed human agents.

A distributed cognitive system is one that dynamically reconfigures itself within its functional constraints to bring subsystems into functional coordination. Many of the subsystems lie outside individual minds; in distributed cognition, interactions between people as they work with external resources are as important as the processes of individual cognition. Both internal mental activity and external interactions play important roles as do physical resources that reveal relationships and act as reminders. A distributed system that

2.2 Research in Situated Cognition and Ecological psychology focuses primarily on individual interactions with environment or artifacts but their foundational ideas are consistent with this distributed view.

involves many people and diverse artifacts in the performance of cognitive work is therefore properly viewed as a cognitive system.

A cognitive system is a thinking (or intelligent) information system. However, the enhanced intelligence is not generated by the activity of intelligent technological functions as many in the discipline of Artificial Intelligence want to claim, but emerges from the coordinated collaboration of distributed human agents via their interactions with each other and with functionally heterogeneous technological artifacts. In the sense that collaboration between human agents and their use of technological artifacts is coordinated, effective, robust and meaningful, the cognitive system is intelligent.

It is sometimes argued that computer-based agents can be employed to reason about the beliefs of human participants in teams. However, computer-based agents follow programmed rules, they do not reason. More generally, people reason but technological devices do not. Two people in coordination can possibly reason more effectively than either in isolation, and if they (as a coordinated dyad) avail themselves of the opportunities presented by technological devices that can compute logical relationships, find and organize information, and probably offer a number of as yet unimagined supporting functions, these entities (the two people together with their technological devices) constitute a reasoning system. Heterogeneity (people with different capabilities, the availability of diverse functional resources) will enhance the potential of the system to perform complex cognitive work.

Implications for analysis and design. Note that heterogeneity does not ensure more effective performance of a cognitive system. It is our job as designers to promote effective cognitive performance by assembling and configuring the requisite functional resources and the requisite collaborative supports. The recommendations offered within the implications for analysis and design for each of the preceding discussions of Situated Cognition, Distributed Cognition, Requisite Variety, Ecological Psychology and Self Organization suggest the means for accomplishing this.

Work Centering: Whence the Images

Most, if not all scientific developments emerge from an image that is acquired informally through natural interaction in the world. Theorizing in behavioral science has traditionally derived formative images from some sort of well-known mechanism. Most recently, the digital computer has played a central role, but appeals to formal logic and mathematical relationships have also been influential.

The rational choice strategy for decision making is one such logical-mathematical strategy. The theory of rational choice posits that decision makers first identify options for action and then define dimensions of evaluation, weight each dimension, rate each option on each dimension, multiply the dimensional weightings, total up the scores and select the option with the highest score. Klein (1998) notes that he entered his early decision research committed to the assumption that he would find evidence of option comparison. Only after confrontation with evidence that suggested otherwise could he divert his attention from that idea and develop the concept of recognition-primed decision making.

Klein's work on recognition-primed decision making, now held in high regard, involves a radical move that attracts little comment. Klein rejected decades, possibly even centuries, of reliance on logical and technological images in favor of a work-centered image, one drawn from the way that experienced operators conceptualize their work. Quite independently, it seems, Rasmussen had already made this move and researchers in Situated Cognition were actively working through it.

I have heard it said that Cognitive Engineering is no more than good Human Factors or good applied cognitive science. I reject that observation and do so primarily because of this move. Human Factors is guided predominately by theoretical images derived from technology and logic. In contrast, Cognitive Engineering is work-centered not only in practice but also in theory. We are no longer deriving formative images from mechanism (e.g., the computer) but from ethnographic descriptions or analyses of cognitive work.

Implications for analysis and design. To be work-centered means to be concerned, first and foremost, with what must be accomplished. A work-centered approach rejects ideas that have the human agent subservient to the technology (*man is best when doing least*, Birmingham and Taylor, 1954) or imply a parity between the human agent and the technology (the *team-player* analogy for interaction of humans with automation, Dekker and Woods, 2002). Cognitive Systems Engineers must first understand the nature of the work (what must be accomplished, how it is accomplished, how it might be accomplished in the future) and then set about designing technological supports and organizational configurations to enhance the conduct of that work.

A Theory of Work Practice

Developments in Cognitive Work Analysis have been guided by a largely unstated theory of work practice. While each of the theoretical perspectives outlined above can be considered a theory in its own right, and each contributes to how we might understand a theory of work practice, none constitutes a comprehensive theory of work practice. That theory of work practice needs to characterize the structure within which work is accomplished and the processes with which it is accomplished. The theory that underlies Cognitive Work Analysis does that and I will outline it throughout the early sections of the next six chapters in this book.

Chapter Summary

I sometimes hear the claim that we cannot predict human behavior. While that is true for specific details, there are certain aspects of human behavior that are predictable. For example, if you and I meet for dinner at a restaurant, I will not be able to predict what you order from the menu, but I can predict with good reliability that you will order something and also the upper and lower boundaries of how much you will eat. To illustrate with another example, an architect can design a family home without knowing the specifics of what will be done in that home. S/he knows enough about human behavior at the level of description required for architectural design to do what is necessary. One of the critical but unstated assumptions of cognitive engineering is that we can predict human behavior at that level of description required for design of cognitive support tools and cognitive systems.

The foundational perspective I outline in this section draws on certain observable and regular behavioral phenomena that can inform cognitive design. Theories or conceptualizations of Situated Cognition, Distributed Cognition, Requisite Variety, Ecological Psychology, Self Organization, Cognitive Systems, Work Centering and Work Practice strengthen our conceptualization of the regularities of human behavior. In this chapter, I suggest that each has specific implications for analysis and design.

The practice of Cognitive Work Analysis is also based on observable and regular behavioral phenomena of the sort that can inform cognitive design. It would be useful to connect the theories and conceptualizations I have discussed in this chapter to the practice of Cognitive

Work Analysis but that would make for a scholarly and detailed treatment, which is not my purpose here. Rather, I introduce these ideas to set them as context for the later discussion and I hope, if you are puzzled by any particular elements of my approach, you will be able to reflect on the ideas presented in this chapter in order to understand the rationale.