


Cognitive Systems Engineering:

Connecting Situations and Awareness



John M. Flach

*Professor Emeritus
Wright State University
Dayton, OH*

©2026 All rights reserved by the author.

Table of Contents

PREFACE	4
HOW TO USE THIS BOOK	7
ACKNOWLEDGEMENTS.....	8
SETTING THE STAGE	9
CHAPTER 1: ENGINEERING ‘COGNITIVE SYSTEMS’	10
CHAPTER 2: FROM INFORMATION TO MEANING PROCESSING.....	15
CHAPTER 3: THE SYSTEM IMAGE	18
CHAPTER 4: CAPABILITIES AND LIMITATIONS.....	22
CHAPTER 5: TOO MUCH DATA OR TOO LITTLE MEANING?	25
CHAPTER 6: WORK VERSUS TASK ANALYSIS.....	28
CHAPTER 7: MATCHING VERSUS SHAPING EXPECTATIONS	30
EXPLORING DEEPER	33
CHAPTER 8: MULTIPLE PERSPECTIVES ON DESIGN	34
CHAPTER 9: DESCRIBING A WORK DOMAIN: FUNCTION, ABSTRACTION, AND MEANING	37
CHAPTER 10: GRASPING COMPLEXITY THROUGH DECOMPOSITION AND ABSTRACTION	43
CHAPTER 11: WORK DOMAIN ANALYSIS: A NOT ATYPICAL SCENARIO	46
CHAPTER 12: COGNITIVE SYSTEMS ENGINEERING: IS IT WORTH THE SQUEEZE?	50
CHAPTER 13: COGNITIVE SYSTEMS ENGINEERING: PUTTING IT TO USE	52
CHAPTER 14: SERIOUS PLAY: A DIFFERENT SLANT ON DESIGN THINKING	54
CHAPTER 15: WHAT DOES AN ‘ECOLOGICAL INTERFACE’ LOOK LIKE?	59
CHAPTER 16: SKILLED MUDDLING	63
CHAPTER 17: MUDDLING THROUGH WICKED PROBLEMS.....	66
CHAPTER 18: THE HEART OF MUDDLING	69
CHAPTER 19: RATIONAL VS INTUITIVE: ANOTHER FALSE DICHOTOMY	72
CHAPTER 20: THE GORDIAN KNOT OF COMPLEXITY.....	75
CHAPTER 21: REVISUALIZING THE DECISION LADDER IN TERMS OF FLOW	77
CHAPTER 22: HEURISTICS? THE GOOD, THE BAD, AND THE UGLY.	80
CHAPTER 23: CHANGE: GOING WITH THE FLOW.....	82
CHAPTER 24: PURPOSE OR FUNCTION.....	85
CHAPTER 25: BETTER MAPS OR BETTER MUDDLING?.....	89
CHAPTER 26: MULTI-LOOP CONTROL SYSTEMS.....	93
CHAPTER 27: POLYCENTRIC CONTROL, DEGREES OF FREEDOM & RESILIENCE	98
CHAPTER 28: DEGREES OF FREEDOM: A BLESSING AND A CURSE.....	101
CHAPTER 29: POLYCENTRIC CONTROL: ANOTHER VIEW	103
CHAPTER 30: RE-ENVISIONING BOYD’S OODA LOOP AS A LAYERED (POLYCENTRIC) ADAPTIVE CONTROL SYSTEM	105
CHAPTER 31: BOYD: AN ORGANIZATIONAL PERSPECTIVE.....	107
CHAPTER 32: A POLYCENTRIC CONTROL STORY	110
OVER THE EDGE.....	113
CHAPTER 33: TRUTH? A PRAGMATIC PERSPECTIVE ON MEANING.....	114
CHAPTER 34: GETTING PAST THE BOXES	116

CHAPTER 35: THE MYTH OF ‘CONTROL’	120
CHAPTER 36: EXPLAINING MECHANISMS VS. ORGANISMS	122
CHAPTER 37: PLAYING TWENTY QUESTIONS WITH NATURE	127
CHAPTER 38: WHAT TIME IS IT?	130
CHAPTER 39: ALL CIRCLES ARE NOT THE SAME	133
CHAPTER 40: EXPERIENCE: REALITY OR ILLUSION	135
CHAPTER 41: DO WE MAKE OR DISCOVER MEANING?	137
ABOUT THE AUTHOR	140

Preface

It is a bit of a jolt for me to realize that it has now been more than 30 years since the publication of Rasmussen, Pejtersen and Goodstein's (1994) important book "Cognitive Systems Engineering." It is partly a jolt because it is a reminder of how old I am getting. In 1984 I completed the Ph.D. in Human Experimental Psychology at The Ohio State University, and I started my first academic job as an Assistant Professor at the University of Illinois Urbana-Champaign. I had accepted a split appointment in three departments – Mechanical & Industrial Engineering, Psychology, and the Institute of Aviation. I would be working with Chris Wickens, who had just published a book - *Engineering Psychology and Human Performance*, and who was a leading authority in the field. I was slated to teach the graduate course in *Engineering Psychology*, and I was anticipating being able to use Chris's new text, which I thought was outstanding. However, I was disappointed to learn that Chris was using the text for another graduate course in *Human Performance*, and I would have to find another source.

I can't remember exactly how I got a hold of it, but somehow, I obtained the manuscript for a book that Jens Rasmussen was writing – *Information Processing and Human-Machine Interaction* – which became the primary text for the course. That was my introduction to Jens Rasmussen's work, and it led to a long-term relationship that greatly shaped my perspective and my career. As we went through the book, I asked the students to take notes on each chapter that I shared with Jens – although the book was in its final stage and it was too late for this feedback to be incorporated into the final product. However, after a few years we were able to do the same with manuscripts for Jens' second book (*Cognitive Systems Engineering*) and I believe that our feedback during the writing of that book was useful for Jens.

Jens' perspective was new for me. However, I found it to fit well with both the control theory that I had learned from Richard Jagacinski, and with James and Eleanor Gibson's ecological perspective on psychology that I had learned from Dean Owen and Rik Warren at Ohio State. Familiarity with the language and logic of closed-loop systems was of great value when interacting with Jens. I think Jens, who was trained in engineering control theory, was quite surprised to find a psychologist who had experience with control theory and its implications for understanding the coupling between perception and action.

As a teacher, I have always been a good student. That is, in choosing readings and topics I was typically motivated by curiosity, rather than any sense of authority. I used my courses

to force myself to explore new topics and approaches. One of the reasons I had originally been excited to use Chris's text was that I saw it as a great opportunity to learn from him. But his text was very familiar and was closely aligned with the information processing framework that was the dominant paradigm in which I had been trained at OSU. In contrast, Jens' work offered a completely different perspective. Thus, I was not teaching it to the students, but rather we were learning it together. This became obvious when Kim Vicente came to Illinois during my third or fourth year.

Kim Vicente was an exceptional student. He had gotten his undergraduate degree at the University of Toronto, where he worked with Neville Moray. He then obtained his MS degree at Virginia Tech. However, while at Virginia Tech he became interested in Rasmussen's work and also became interested in learning more about ecological psychology. Neville Moray, who had come to Illinois, suggested to Kim that I might be a good source for learning about ecological psychology. So, he applied and was accepted into the Ph.D. program at Illinois. However, he delayed coming to Illinois for a year to accept an opportunity to go to Denmark to spend a year working at RISO National Laboratory with Rasmussen.

When Kim arrived at Illinois after an intensive year working with Jens, we were using draft chapters of Jens' second book in the Engineering Psychology course. More than once after he arrived, I can remember speculating about an aspect of Jens' thinking – and then having Kim raise his hand and correct me – to explain what Jens actually meant. Each time, upon hearing Kim's explanation it was clear that his understanding was deeper than mine. So, it turns out, that Kim was doing a lot of the teaching, and I was happy to learn from him.

I have now been teaching and applying Rasmussen's Cognitive Systems Engineering (CSE) framework for more than 40 years. About 10 years ago, I started sharing my thoughts about CSE in a blog. However, I sometimes wonder if my efforts have had any impact at all.

In 2017 I joined a small software development company, where Alex Morrison (a student of Dave Woods) and I tried to integrate the CSE approach into an interdisciplinary approach to design, but after some initial success, enthusiasm for CSE waned and lost support due to economic constraints. It seemed that customers wanted quick incremental solutions, and they could not be convinced to make the investment in time and effort that it takes to do a thorough work domain analysis. This is not unreasonable. It does take significant time to learn about a domain – it simply can't be done in a six month or one-year contract. It may be that the investment in time and effort required by CSE is only feasible for large industries or government agencies that own work domains – such as nuclear power, aviation, or national defense. But maybe things are changing – maybe new opportunities are emerging.

John Allspaw recently made a LinkedIn post in which he wrote the following:

What if I told you that there exists an approach that is different than "AI does everything autonomously without human involvement" or "People now only supervise the work of AI"?

What if I told you also that perspectives like this come from a field almost entirely unknown to the mainstream tech industry? And that the entire rationale for this field and its perspectives was to fuel the actual design of real tools, real support systems, and real interfaces, that is different than what the industry is used to?

The field is known as Cognitive Systems Engineering.

It is also a jolt to realize that many people who are actively involved in the design of complex sociotechnical systems have very little understanding of what CSE is or why it is valuable. For example, they might ask how is CSE different than Human Factors or UX Design? It seems despite the efforts of me and many others, the message has not reached many of the communities who are involved in shaping the future of sociotechnical systems (e.g., computer scientists, human factors engineers, designers, chief technology officers).

It is also a jolt when I consider how technology has changed over the past 25 years. A common theme in Rasmussen's arguments for why CSE was necessary was the fast pace of change due to advances in modern information technologies. Boy did he get that right. Thirty years ago, the Internet was in its infancy. Google was a research project, not formally incorporated as a company till 1998. There was no Facebook, which wasn't launched until 2004. No YouTube - the first video was uploaded in 2005. No Twitter - the first tweet wasn't sent until 2006. And perhaps most significantly, there were no smart phones. The first iPhone wasn't released until 2007.

I'm not sure if Jens (or anyone else) fully anticipated the impact of the rapid advances of technology on work and life today. However, I do feel that most of his ideas about the implications of these advances for doing research on and for designing sociotechnical systems are still relevant. In fact, the relevance has only grown with the changes that information technologies have fostered. Despite the emergence of other terms for new ways to think about work, such as *UX Design* and *Resilience Engineering*, I contend that having an understanding of CSE remains essential to anyone who is interested in studying or designing sociotechnical systems.

After reading John Allspaw's post, I was motivated to collect some of my past blogs into this book. Each chapter is a lightly edit version of an original blog post. Maybe, there is still an opportunity for me to be a teacher! But if any of this puzzles you – you might look for your own Kim Vicente – who can help you make sense of it.

The subtitle of the book plays off the concept of Situation Awareness. I believe this term was historically used by the military in relation to the need to keep track of the enemy and to anticipate potential risks and opportunities. For example, combat pilots would talk about maintaining or losing situation awareness. The term was later introduced and popularized in human factors by Mica Endsley. Although I have been critical of some of the ways models of situation awareness has been framed, I think the term nicely captures the spirit of early functional psychologists (e.g., William James, John Dewey), pragmatist philosophers (C.S. Pierce) and ecological psychologists (Eleanor and James Gibson) who framed human experience in terms of the capacity of people to adapt to their ecology (or situation). As Herbert Simon illustrated with his metaphor of the ant on the beach, theories of adaptive behavior have to consider relations between the properties of the ant and of the beach. Similarly, James Gibson's concept of affordance (e.g., graspable) is a prime example of a relational construct that is a joint function of an agent (e.g., hand size) and an ecology (object size). Thus, I chose the subtitle to suggest that CSE is a natural extension of ecological theories of human performance to the challenge of designing systems that expand our adaptive capacity.

The book is organized into three sections. The first section is intended to set the stage by contrasting a CSE approach to 'meaning processing' from a more traditional approach to Human Factors that was based on an 'information processing' model. The second section goes into more detail about a CSE approach including discussions of work analysis, ecological interface design (EID), and sensemaking (both in terms of individuals and organizations). The final section is a bit more speculative, including some discussions of a more philosophical nature.

How to use this book

This book is intended as a supplement for an upper level undergraduate or graduate course or seminar on human factors, human-centered computing, human-machine systems, user-centered design, or cognitive systems engineering. The first section should be required reading, as an important check on the assumptions that students might have due to prior exposure to more conventional courses in psychology or human performance. The other two sections provide a range of topics that can be selected to reflect an instructor's specific focus for the course (e.g., whether the emphasis is on design, cognitive work analysis, human performance, or organizational sensemaking).

Acknowledgements

I won't even try to list all the teachers, students, and colleagues who have shaped my thinking about humans, life, and the universe. But I do need to recognize Fred Voorhorst's contributions as the artist who created many of the graphics that accompanied some of my posts and who was often a sounding board that helped shape the final thoughts expressed.

Also, I would like to dedicate this work to my wife, Linnea (Bits) who made sure that I didn't take the academic world of ideas too seriously and that I always stayed grounded in the day-to-day realities of what matters. Even though she has passed from this world, her influence endures and colors how I see myself and how I see the world. Finally, thanks to my sons, their wives, and their children who keep me moving forward with the hope that tomorrow will be just a little better than yesterday.

Setting the Stage

Chapter 1: Engineering ‘Cognitive Systems’

The term “Cognitive Systems Engineering” was first suggested by Erik Hollnagel and Dave Woods in a RISO technical report in 1982, that was titled “*Cognitive Systems Engineering: New wine in new bottles.*” Later, Don Norman used the term “Cognitive Engineering” as the title for a chapter in a book (User-Centered System Design - 1986) that he co-edited with Stephen Draper. Norman (1986) writes:

Cognitive Engineering, a term invented to reflect the enterprise I find myself engaged in: neither Cognitive Psychology, nor Cognitive Science, nor Human Factors. It is a type of applied Cognitive Science, trying to apply what is known from science to the design and construction of machines. It is a surprising business. On the one hand, there actually is a lot known in Cognitive Science that can be applied. On the other hand, our lack of knowledge is appalling. On the one hand, computers are ridiculously difficult to use. On the other hand, many devices are difficult to use - the problem is not restricted to computers, there are fundamental difficulties in understanding and using most complex devices. So the goal of Cognitive Engineering is to come to understand the issues, to show how to make better choices when they exist, and to show what the tradeoffs are when, as is the usual case, an improvement in one domain leads to deficits in another (p. 31).

Norman (1986) continues to specify two major goals that he had as a Cognitive Systems Engineer:

To understand the fundamental principles behind human action and performance that are relevant for the development of engineering principles in design.

To devise systems that are pleasant to use - the goal is neither efficiency nor ease nor power, although these are all to be desired, but rather systems that are pleasant, even fun: to produce what Laurel calls “pleasurable engagement” (p. 32).

Rasmussen (1986) was less interested in pleasurable engagement and more interested in safety - noting the accidents at Three Mile Island and Bhopal as important motivations for different ways to think about human performance and work. As a control engineer Rasmussen was concerned that the increased

utilization of centralized, automatic control systems in many industries (particularly nuclear power) was changing the role of humans in those systems. He noted that the increased use of automation was moving humans “from the immediate control of system operation to higher-level supervisory tasks and to long-term maintenance and planning tasks” (p. 1). Because of his background in control engineering, Rasmussen understood the limitations of the automated control systems, and he recognized that these systems would eventually face situations that their designers had not anticipated (i.e., situations for which the ‘rules’ or ‘programs’ embedded in these systems were inadequate). He knew that it would be up to the human supervisors to detect and diagnose the problems that would thus ensue and to intervene (e.g., creating new rules on the fly) to avert potential catastrophes.

The challenge that he saw for CSE was to improve the interfaces between the humans and the automated control systems to support supervisory control. Rasmussen wrote:

Use of computer-based information technology to support decision making in supervisory systems control necessarily implies an attempt to match the information processes of the computer to the mental decision processes of an operator. This approach does not imply that computers should process information in the same way as humans would. On the contrary, the processes used by computers and humans will have to match difference resource characteristics. However, to support human decision making and supervisory control, the results of computer processing must be communicated at appropriate steps of the decision sequence and in a form that is compatible with the human decision strategy. Therefore, the designer has to predict, in one way or another, which decision strategy an operator will choose. If the designer succeeds in this prediction, a very effective human-machine cooperation may result; if not, the operator may be worse off with the new support than he or she was in the traditional system ... (p. 2).

Note that the information technologies that were just beginning to change the nature of work in the nuclear power industry in the 1980's when Rasmussen made these observations have now become significant parts of almost every aspect of modern life - from preparing a meal (e.g., Chef Watson), to maintaining personal social networks (e.g., Facebook and Instagram), to healthcare (e.g., electronic health

record systems), to manufacturing (e.g., flexible, just-in-time systems), to shaping the political dialog (e.g., President Trump's use of Twitter). Today, most of us have more computing power in our pocket (our smart phones) than was available for even the most modern nuclear power plants in the 1980s. In particular, the display technology in 1980s was extremely primitive relative to the interactive graphics that are available today on smart phones and tablets.

A major source of confusion that has arisen in defining this relatively new field of CSE has been described in a blog post by Erik Hollnagel (2017):

The dilemma can be illustrated by considering two ways of parsing CSE. One parsing is as C(SE), meaning cognitive (systems engineering) or systems engineering from a cognitive point of view. The other is (CS)E, meaning the engineering of (cognitive systems), or the design and building of joint (cognitive) systems.

From the earliest beginnings of CSE, Hollnagel and Woods (1982; 1999) were very clear about what they thought was the appropriate parsing. Here is their description of a cognitive system:

A cognitive system produces "intelligent action," that is, its behavior is goal oriented, based on symbol manipulation and used knowledge of the world (heuristic knowledge) for guidance. Furthermore, a cognitive system is adaptive and able to view a problem in more than one way. A cognitive system operates using knowledge about itself and the environment, in the sense that it is able to plan and modify its actions on the basis of that knowledge. It is thus not only data driven, but also concept driven. Man is obviously a cognitive system. Machines are potentially if not actually, cognitive systems. An MMS [Man-Machine System] regarded as a whole is definitely a cognitive system (p. 345)

Unfortunately, there are still many who don't quite fully appreciate the significance of treating the whole sociotechnical system as a unified system where the cognitive functions are emergent properties that depend on coordination among the components. Many have been so entrained in classical reductionistic approaches that they can't resist the temptation to break the larger sociotechnical system into components. For these people, CSE is simply systems engineering techniques applied to cognitive components within the larger sociotechnical system. This approach fits with the classical disciplinary divisions in universities where the social

sciences and the physical sciences are separate domains of research and knowledge. People generally recognize the significant role of humans in many sociotechnical systems (most notably as a source of error); and they typically advocate that designers account for the 'human factor' in the design of technologies. However, they fail to appreciate the self-organizing dynamics that emerge when smart people and smart technologies work together. They fail to recognize that what matters with respect to successful cognitive functioning of this system are emergent properties that cannot be discovered in any of the components. Just as a sports team is not simply a collection of people, a sociotechnical system is not simply a collection of things (e.g., humans and automation).

The ultimate challenge of CSE as formulated by Rasmussen, Norman, Hollnagel, Woods and others is to develop a new framework where the 'cognitive system' is a fundamental unit of analysis. It is not a collection of people and machines - rather it is an adapting organism with a life of its own (i.e., dynamic properties that arise from relations among the components). CSE reflects a desire to understand these dynamic properties and to use that understanding to design systems that are increasingly safe, efficient, and pleasant to use.

References:

<https://erikhollnagel.com/ideas/cognitive-systems-engineering-1982>

Hollnagel, E. & Woods, D. D. (1982). [Cognitive systems engineering: New wine on new bottles](#) (Risø-M-2330). Roskilde, Denmark: Risø National Laboratory.

Hollnagel, E. & Woods, D.D. (1999). Cognitive Systems Engineering: New wine in new bottles. *In. J. Human-Computer Studies*, 51, 339-356.

Norman, D.A. (1986). Cognitive Engineering. In D.A. Norman & S.W. Draper (Eds). *User-Centered System Design*, 31-61, Hillsdale, NJ: Erlbaum.

Rasmussen, J. (1986). *Information Processing and Human-Machine Interaction: An Approach to Cognitive Engineering*. New York: North Holland.

Chapter 2: From Information to Meaning Processing

In the early 1960s, we realized from analyses of industrial accidents the need for an integrated approach to the design of human-machine systems. However, we very rapidly encountered great difficulties in our efforts to bridge the gap between the methodology and concepts of control engineering and those from various branches of psychology. Because of its kinship to classical experimental psychology and its behavioristic claim for exclusive use of objective data representing overt activity, the traditional human factors field had very little to offer (**Rasmussen**, p. ix, 1986).

Cognitive Engineering, a term invented to reflect the enterprise I find myself engaged in: neither Cognitive Psychology, nor Cognitive Science, nor Human Factors. (**Norman**, p. 31, 1986).

The growth of computer applications has radically changed the nature of the man-machine interface. First, through increased automation, the nature of the human's task has shifted from an emphasis on perceptual-motor skills to an emphasis on cognitive activities, e.g. problem solving and decision making.... Second, through the increasing sophistication of computer applications, the man-machine interface is gradually becoming the interaction of two cognitive systems. (**Hollnagel & Woods**, p. 340, 1999).

As reflected in the above quotes, through the 1980s and 90s there was a growing sense that the nature of human-machine systems was changing, and that this change was creating the demand for a new approach to analysis and design. This new approach was given the label of *Cognitive Engineering (CE)* or *Cognitive Systems Engineering (CSE)*. Table 1, illustrates one way to characterize the changing role of the human factor in sociotechnical systems. As technologies became more powerful with respect to information processing capabilities, the role of the humans increasingly involved supervision and fault detection. Thus, there was increasing demand for humans to relate the activities of the automated processes to functional goals and values (i.e., pragmatic meaning) and to intervene when circumstances arose (e.g., faults or unexpected situations) such that the automated processes were no longer serving the design goals for the system. Under these unexpected circumstances, it was often necessary for the humans to create or invent new procedures on the fly, to avoid potential hazards or to take advantage of potential opportunities. This demand to relate activities to functional goals and values and to improvise new procedures required meaning processing.

Information Processing	Meaning Processing
• System Image • Communication Channel • Dyadic Semiotic System (symbol processing) • Emphasis on Human Limitations • 7+ or – 2 Chunks • Heuristics = Biases • Workload - Too much data • Task Analysis – Procedures & Activities • Design Focus • Respect Limitations • Match expectations (mental models) <i>Important, but not sufficient!</i>	• System Image • Control System • Triadic Semiotic System (pragmatic assessment) • Emphasis on Human Capabilities • Recoding/Chunking • Heuristics = Tricks of the trade • Situation Awareness – Too little meaning • Work Analysis – Constraints & Possibilities • Design focus • Leverage Capabilities • Shape expectations (mental models) <i>Potential enhancements in perspective!</i>

In this booklet, I will elaborate the differences between information processing and meaning processing as outlined in Table 1. However, before I get into the specific contrasts, it is important to emphasize that relative to early human factors, CSE is an evolutionary change, NOT a revolutionary change. That is, the concerns about information processing that motivated earlier human factors efforts were not wrong, and they continue to be important with regards to design. The point of CSE is not that an information processing approach is wrong, but rather that it is insufficient.

The point of CSE is not that an information processing approach is wrong, but rather that it is insufficient.

The point of CSE is that design should not simply be about avoiding overloading the limited information capacities of humans, but it should also seek to leverage the unique meaning processing capabilities of humans. These meaning processing capabilities reflect peoples' ability to make sense of the world relative to their values and intentions, to adapt to surprises, and to improvise in order to take advantage of new opportunities or to avoid new threats. In following chapters I will make the case that the overall vision of a meaning processing approach is more expansive than an information processing approach. This broader context will sometimes change the significance of specific information limitations. It will also provide a wider range of options for by-passing these limitations and for innovating to improve the range and quality of performance of sociotechnical systems.

In other words, I will argue for a systems perspective - where the information processing limitations must be interpreted in light of the larger meaning processing context. I will argue

that constructs such as 'expertise' can only be fully understood in terms of qualities that emerge at the level of meaning processing.

References:

Hollnagel, E. & Woods, D.D. (1999). Cognitive Systems Engineering: New wine in new bottles. *In. J. Human-Computer Studies*, 51, 339-356.

Norman, D.A. (1986). Cognitive Engineering. In D.A. Norman & S.W. Draper (Eds). *User-Centered System Design*, 31-61, Hillsdale, NJ: Erlbaum.

Rasmussen, J. (1986). *Information Processing and Human-Machine Interaction: An Approach to Cognitive Engineering*. New York: North Holland.

Chapter 3: The System Image

A rather detailed account of the nineteenth-century history of the steam engine with governor may help the reader to understand both the circuits and the blindness of the inventors. Some sort of governor was added to the early steam engine, but the engineers ran into difficulties. They came to Clark Maxwell with the complaint that they could not draw a blueprint for an engine with a governor. They had no theoretical base from which to predict how the machine that they had drawn would behave when built and running.

There were several possible sorts of behavior: Some machines went into runaway, exponentially maximizing their speed until they broke or slowing down until they stopped. Others oscillated and seemed unable to settle to any mean Others - still worse - embarked on sequences of behavior in which the amplitude of their oscillation would itself oscillate or would become greater and greater.

Maxwell examined the problem. He wrote out formal equations for relations between the variables at each successive step around the circuit. He found, as the engineers had found, that combining this set of equations would not solve the problem. Finally, he found that the engineers were at fault in not considering time. Every given system embodied relations to time, that is was characterized by time constants determined by the given whole. These constants were not determined by the equations of relationship between successive parts but were emergent properties of the system.

... a subtle change has occurred in the subject of discourse It is a difference between talking in a language which a physicist might use to describe how one variable acts upon another and talking in another language about the circuit as a whole which reduces or increases difference. When we say that the system exhibits "steady state" (i.e., that in spite of variation, it retains a median value), we are talking about the circuit as a whole, not about the variations within it. Similarly the question which the engineers brought up to Clark Maxwell was about the circuit as a whole: How can we plan it to achieve a steady state? They expected the answer to be in terms of relations between the individual variables. What was needed and supplied by Maxwell was an answer in terms of time constants of the total circuit. This was the bridge between the two levels of discourse. Gregory Bateson (2002) p. 99-101.

This chapter will address the first line in the table below that is repeated from Chapter 1.

Information Processing	Meaning Processing
• System Image • Communication Channel • Dyadic Semiotic System (symbol processing) • Emphasis on Human Limitations • 7+ or – 2 Chunks • Heuristics = Biases • Workload - Too much data • Task Analysis – Procedures & Activities • Design Focus • Respect Limitations • Match expectations (mental models) <i>— Important, but not sufficient! —</i>	• System Image • Control System • Triadic Semiotic System (pragmatic assessment) • Emphasis on Human Capabilities • Recoding/Chunking • Heuristics = Tricks of the trade • Situation Awareness – Too little meaning • Work Analysis – Constraints & Possibilities • Design focus • Leverage Capabilities • Shape expectations (mental models) <i>— Potential enhancements in perspective! —</i>

Norbert Wiener's (1948) classic, introducing the Cybernetic Hypothesis, was subtitled "On Control and Communication in the Animal and Machine." The ideas in this book about control systems (machines that were designed to achieve and maintain a goal state) and communication (the idea that information can be quantified) had a significant impact on the framing of research in psychology. It helped shift the focus from *behavior* to *cognition* (e.g., Miller, Gallanter, Pribram, 1960).


However, though psychologists began to include *feedback* in their images of cognitive systems, the early program of research tended to be dominated by the image of an open-loop communication system and the research program tended to focus on identifying stimulus-response associations or transfer functions (e.g., bandwidth) for each component in a series of discrete information processing stages. A major thrust of this research program was to identify the limitations of each subsystem in terms of storage capacity (7 + or - 2 chunk capacity of working memory) and information processing rates (e.g., Hick-Hyman Law, Fitts' Law).

Thus, the Cybernetic Hypothesis inspired psychology to consider "intentions" and other internal aspects associated with thinking, however, it did not free psychology from using the language of the physicist in describing the causal interaction between one variable and another, rather than thinking in terms of properties of the circuit as a whole (i.e., appreciating the emergent properties that arise from the coupling of perception and action). The language of information processing psychology was framed in terms of a dyadic semiotic system for processing symbols as illustrated below.

Dyadic Semiotic System



Triadic Semiotic System



In contrast, CSE was framed from the start in the language of control theory. This reflected an interest in the role of humans in closing the loop as pilots of aircraft and supervisors of energy production processes. From a control theoretic perspective, it was natural to frame the problems of meaning processing as a triadic semiotic system, where the function of cognition was to achieve stable equilibrium with a problem ecology. Note that the triadic semiotic model emerged because of the work of functional psychologists (e.g., James & Dewey) and pragmatic philosophers (Peirce), who were most interested in 'mind' as a means for adapting to the pragmatic demands of everyday living.* Dewey's (1896) classic paper on the reflex arc examines the implications of Maxwell's insights (described in the opening quote from Bateson) for psychology (emphasis added):

The discussion up to this point may be summarized by saying that the reflex arc idea, as commonly employed, is defective in that it assumes sensory stimulus and motor response as distinct psychical existences, while in reality they are always inside a coordination and have their significance purely from the part played in maintaining or reconstituting the coordination; and (secondly) in assuming that the quale of experience which precedes the 'motor' phase and that which succeeds it are two different states, instead of the last being always the first reconstituted, the motor phase coming in only for the sake of such mediation. The result is that the reflex arc idea leaves us with a disjointed psychology, whether viewed from the standpoint of development in the

individual or in the race, or from that of the analysis of the mature consciousness. As to the former, in its failure to see that the arc of which it talks is virtually a circuit, a continual reconstitution, it breaks continuity and leaves us nothing but a series of jerks, the origin of each jerk to be sought outside the process of experience itself, in either an external pressure of 'environment,' or else in an unaccountable spontaneous variation from within the 'soul' or the 'organism.' As to the latter, failing to see unity of activity, no matter how much it may prate of unity, it still leaves us with sensation or peripheral stimulus; idea, or central process (the equivalent of attention); and motor response, or act, as three disconnected existences, having to be somehow adjusted to each other, whether through the intervention of an extra experimental soul, or by mechanical push and pull.

Many cognitive scientists and many human factors engineers continue to speak in the language associated with causal, stimulus-response interactions (i.e., jerks) without an appreciation for the larger system in which perception and action are coupled. In essence, they treat cognition as a property in the head or brain of an agent, that can be understood independent of any particular ecology within which it might be situated. They are still hoping to concatenate these isolated pieces into a more complete picture of human performance. In contrast, CSE starts with a view of the whole - of the coupling of perception and action through an ecology - as a necessary context from which to appreciate variations at more elemental levels. Thus, the CSE approach treats the human and ecology as a joint cognitive system – where performance is an emergent property of the functional coupling.

References

Dewey, J. (1896). The reflex arc concept in psychology. *Psychological Review*, 3(4), 357-370.

Bateson, G. (2002) [1979]. *Mind and Nature: A Necessary Unity*. Hampton Press.

Miller, G. A., Galanter, E., & Pribram, K. H. (1960). *Plans and the structure of behavior*. Henry Holt and Co.

Wiener, N. (1948). *Cybernetics: or control and communication in the animal and the machine*. Cambridge, MA: MIT Press.

*See Louis Menand's books as a good source for learning about CS Peirce, William James and the other pragmatists and their functional framework for understanding human cognition: *Metaphysical Club* (2001), *Pragmatism: A Reader* (edited by Menand)

Chapter 4: Capabilities AND Limitations

*First, the span of absolute judgment and the span of immediate memory impose severe limitations on the amount of information that we are able to receive, process, and remember. **By organizing the stimulus input simultaneously into several dimensions and successively into a sequence of chunks, we manage to break (or at least stretch) this informational bottleneck.***

*Second, **the process of recoding is a very important one in human psychology and deserves much more explicit attention than it has received.** In particular, the kind of linguistic recoding that people do seems to me to be the very lifeblood of the thought processes. Recoding procedures are a constant concern to clinicians, social psychologists, linguists, and anthropologists and yet, probably because recoding is less accessible to experimental manipulation than nonsense syllables or T mazes, the traditional experimental psychologist has contributed little or nothing to their analysis. Nevertheless, experimental techniques can be used, methods of recoding can be specified, behavioral indicants can be found. And I anticipate that we will find a very orderly set of relations describing what now seems an uncharted wilderness of individual differences. George Miller (1956) p. 96-97, (emphasis added).*

This chapter will focus on the second line in the table - shifting emphasis from human limitations to include greater consideration of human capabilities.

Information Processing	Meaning Processing
• System Image • Communication Channel • Dyadic Semiotic System (symbol processing) • Emphasis on Human Limitations • 7+ or – 2 Chunks • Heuristics = Biases • Workload - Too much data • Task Analysis – Procedures & Activities • Design Focus • Respect Limitations • Match expectations (mental models) <i>— Important, but not sufficient! —</i>	• System Image • Control System • Triadic Semiotic System (pragmatic assessment) • Emphasis on Human Capabilities • Recoding/Chunking • Heuristics = Tricks of the trade • Situation Awareness – Too little meaning • Work Analysis – Constraints & Possibilities • Design focus • Leverage Capabilities • Shape expectations (mental models) <i>— Potential enhancements in perspective! —</i>

Anybody who has taken an intro class in Psychology or Human Factors is familiar with Miller's famous number 7 plus or minus 2 chunks that specifies the capacity of working memory. This is one of the few numbers that we can confidently provide to system

designers as a spec of the human operator that needs to be respected. However, this number has little practical value for design unless you can also specify what constitutes a 'chunk.'

Although people know Miller's number, few appreciate the important points that Miller makes in the second half of the paper about the power of 'recoding' and the implications for the functional capacity of working memory. As noted in the opening quote to this blog - people have the ability to 'stretch' memory capacity through chunking. The intro texts emphasize the "limitation," but much less attention has been paid to the recoding "capability" that allows experts to extend their functional memory capacity to deal with large amounts of information (e.g., that allows an expert chess player to recall all the pieces on a chess board based on a very short glance).

Cataloging visual illusions has been a major thrust of research in perception. Similarly, cataloging biases has been a major thrust of research in decision making. However, people such as James Gibson have argued that these collections of illusions do not add up to a satisfying theory of how perception works to guide action (e.g., in the control of locomotion). In a similar vein, people such as Gerd Gigerenzer have made a similar case for decision making - that collections of biases (the dark side of heuristics) do not add up to a satisfying theory of decision making in everyday life and work. One reason is that in everyday life, there is often missing and ambiguous data and incommensurate variables that make it difficult or impossible to apply more normative algorithmic approaches.

One result of the work of Tversky and Kahneman in focusing on decision errors is that the term 'heuristic' is often treated as if it is synonymous with 'bias.' Thus, heuristics illustrate the 'weakness' of human cognition - the bounds of rationality. However, Herbert Simon's early work in artificial intelligence treated heuristics as the signature of human intelligence. A computer program was only considered intelligent if it took advantage of heuristics that creatively used problem constraints to find short cuts to solutions - as opposed to applying mathematical algorithms that solved problems by mechanical brute force.

This emphasis on illusions and biases tends to support the MYTH that humans are the weak link in any sociotechnical system and it leads many to seek ways to replace the human with supposedly more reliable 'automated' systems. For example, recent initiatives to introduce autonomous cars often cite 'reducing human errors and making the system safer' as a motivation for pursuing automatic control solutions.

Thus, the general concept of rationality tends to be idealized around the things that computers do well (use precise data to solve complex algorithms and logical puzzles) and

it tends to underplay those aspects of rationality where humans excel (detecting patterns and deviations and creatively adapting to surprise/novelty/ambiguity).

CSE recognizes the importance of respecting the bounds of rationality for humans when designing systems - but also appreciates the limitations of automation (and the fact that when the automation fails - it will typically fall to the human to fill the gap). Further, CSE starts with the premise that 'human errors' are best understood in relation to human abilities. On one hand, a catalogue of human errors will never add up to a satisfying understanding of human thinking. On the other hand, a deeper understanding may be possible if the 'errors' are seen as 'bounds' on abilities. In other words, CSE assumes that a theory of human rationality must start with a consideration of how thinking 'works,' in order to give coherence to the collection of errors that emerge at the margins.

The implication is that CSE tends to be framed in terms of designing to maximize human engagement and to fully leverage human capabilities, as opposed to more classical human factors that tends to emphasize the need to protect systems against human errors and limitations. This does not need to be framed as human versus machine, rather it should be framed in terms of human-machine collaboration. The ultimate design goal is to leverage the strengths of both humans and technologies to create sociotechnical systems that extend the bounds of rationality beyond the range of either of the components.

References

Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63(2), 81–97.

James Gibson

Gerd Gigerenzer

Tversky and Kahneman

Herbert Simon's

Chapter 5: Too much data or too little meaning?

This chapter will focus on the third contrast in the table - the shift from a focus on 'workload' to a focus on 'situation awareness.'

Information Processing	Meaning Processing
• System Image • Communication Channel • Dyadic Semiotic System (symbol processing) • Emphasis on Human Limitations • 7+ or – 2 Chunks • Heuristics = Biases • Workload - Too much data • Task Analysis – Procedures & Activities • Design Focus • Respect Limitations • Match expectations (mental models) — <i>Important, but not sufficient!</i> —	• System Image • Control System • Triadic Semiotic System (pragmatic assessment) • Emphasis on Human Capabilities • Recoding/Chunking • Heuristics = Tricks of the trade • Situation Awareness – Too little meaning • Work Analysis – Constraints & Possibilities • Design focus • Leverage Capabilities • Shape expectations (mental models) — <i>Potential enhancements in perspective!</i> —

The concept of information developed in this theory at first seems disappointing and bizarre - disappointing because it has nothing to do with meaning, and bizarre because it deals not with a single message but rather with the statistical character of a whole assemble of messages, bizarre also because in these statistical terms the two words information and uncertainty find themselves to be partners. Warren Weaver (1963, p. 27)


The construct of 'workload' is a natural focus for an approach that emphasizes describing and quantifying internal constraints of the human and that assumes that these constraints are independent of the particulars of any specific situation or work context. This fits well with the engineering perspective on quantifying information and for specifying the capacity of fixed information channels as developed by Shannon. However, the downside of this perspective is that in making the construct of workload independent of 'context,' it thus becomes independent of 'meaning' as suggested in the above quote from Warren Weaver.

Those interested in the impact of context on human cognition became dissatisfied with a framework that focused only on internal constraints (e.g., bandwidth, resources, modality) without consideration for how those constraints interacted with situations. Thus, the construct of *Situation Awareness (SA)* evolved as an alternative to *workload*. Unfortunately, many who have been steeped in the information processing tradition have framed SA in

terms of internal constraints (e.g., treating levels of SA as components internal to the processing system).

However, others have taken the construct of SA as an opportunity to consider the dynamic couplings of humans and work domains (or situations). For them, the construct of SA reflects a need to 'situate' cognition within a work ecology and to consider how constraints in that ecology create demands and opportunities for cognitive systems. In this framework, it is assumed that cognitive systems can intelligently adapt to the constraints of situations - utilizing structure in situations to 'chunk' information and as the basis for smart heuristics that reduce the computational burden, allowing people to deal effectively with situations that would overwhelm the channel capacity of a system not tuned to these structural constraints (see aiming off example).

There is no question that humans have limited working memory capacity as suggested by the workload construct. However, CSE recognizes the ability of people to discover and use situated constraints (e.g., patterns) in ways that allow them to do complex work (e.g., play chess, pilot aircraft, drive in dense traffic, play a musical instrument) despite these internal constraints. It is this capacity to *attune* to structure associated with specific work domains that leads to expert performance.



The design implication of an approach that focuses on workload is to protect the system against human limitations (e.g., bottlenecks) by either distributing the work among multiple

people or by replacing humans with automated systems with higher bandwidth. The key is to make sure that people are not overwhelmed by too much data!

The design implication of an approach that focuses on SA is to make the meaningful work domain constraints salient in order to facilitate attunement processes. This can be done through the design of interfaces or through training. The result is to heighten human engagement with the domain structure to facilitate skill and expertise. The key is to make sure that people are well-tuned to the meaningful aspects of work (e.g., constraints and patterns) that allow them to 'see' what needs to be done.

Situated Cognition

When we think of mind in relation to situations, do we think of mind as if it were the glass or as if it were the water?

The construct of **workload** treats the mind as if it were the glass. That is, it has a finite capacity for holding water and ice. The focus is on specifying that capacity - how much stuff can the mind hold at any time?

The construct of **situation awareness** treats the mind as if it were the water. That is, it has no shape of its own, but it adapts its shape to the constraints of the glass. The focus is on describing how well the water fills the glass - can the mind fully engage the situation - is there enough water? Are there gaps between the cubes?

Probably, neither minds nor situations are as rigid as the glass. Ultimately, however, the point is that minds are more like water than like glasses. Minds tend to be situated. They adapt to the constraints of situations.



References

Chapter 6: Work versus Task Analysis

“An ant, viewed as a behaving system, is quite simple. The apparent complexity of its behavior over time is largely a reflection of the complexity of the environment in which it finds itself.”—Herbert Simon

This chapter focuses on the fourth distinction in the table below.

Information Processing	Meaning Processing
• System Image • Communication Channel • Dyadic Semiotic System (symbol processing) • Emphasis on Human Limitations • 7+ or – 2 Chunks • Heuristics = Biases • Workload - Too much data • Task Analysis – Procedures & Activities • Design Focus • Respect Limitations • Match expectations (mental models) <i>— Important, but not sufficient! —</i>	• System Image • Control System • Triadic Semiotic System (pragmatic assessment) • Emphasis on Human Capabilities • Recoding/Chunking • Heuristics = Tricks of the trade • Situation Awareness – Too little meaning • Work Analysis – Constraints & Possibilities • Design focus • Leverage Capabilities • Shape expectations (mental models) <i>— Potential enhancements in perspective! —</i>

In some respect, it is debatable whether psychology completely escaped from the strictures of a Behaviorist perspective with the dawning of an information processing age. Although the computer and information processing metaphors legitimized the study of mental phenomena, these phenomena tended to be visualized as activities or mental behaviors (e.g., encoding, remembering, planning, deciding, etc.). Thus, for human factors, cognitive **task analysis** tended to focus on specifying the mental activities associated with work.

However, as Simon's parable of the ant suggests, this might lead to the appearance or inference that the cognition is very complex. When in reality, the ultimate source of that complexity may not be in the cognitive system, but in the work situations that provide the context for the activity (i.e., the beach). Thus, Simon's parable is the motivation for **work analysis** as a necessary complement to task analysis. The focus of work analysis is to describe the functional constraints (e.g., goals, affordances, regulatory constraints, social constraints, etc.) that are shaping the physical and mental activities of workers. The focus of work analysis is on describing work situations, as a necessary context for evaluating both awareness (rationality) and behavior.

While classical task analysis describes what people do, it provides little insight into **why** they do something. This is the primary value of work analysis, to provide insight into the functional context shaping work behaviors.

A common experience upon watching activities in an unfamiliar work domain is puzzlement at activities that seem somewhat irrational. Why did they do that? However, as one becomes more familiar with the domain, one often discovers that these puzzling activities are actually smart responses to constraints that the experienced workers were attuned to - but that were invisible to an outsider.

Thus, for those of us who see humans as incredibly adaptive systems - it is natural to look to the ecology that the humans are adapting to as the first source for hypotheses to explain that behavior. And for those of us who hope to design information technologies that enhance that adaptive capacity - it is critical that the tools not simply be tuned to human capabilities, but that these tools are also tuned to the demands of the work ecology. For example, a shovel must not only fit the human hands, but it must also fit well with the materials that are to be manipulated.

Thus, **work analysis** is essential for Cognitive Systems Engineering. It reflects the belief that understanding situations is a prerequisite for understanding awareness.

Some important resources for people interested in learning more about work analysis are Vicente (1999), Naikar (2013), and Jenkins, Stanton, Salmon & Walker (2009).

References

Chapter 7: Matching versus Shaping Expectations

In this chapter the focus will be on design implications of a Cognitive Systems Engineering framework.

Information Processing	Meaning Processing
• System Image • Communication Channel • Dyadic Semiotic System (symbol processing) • Emphasis on Human Limitations • 7+ or – 2 Chunks • Heuristics = Biases • Workload - Too much data • Task Analysis – Procedures & Activities • Design Focus • Respect Limitations • Match expectations (mental models) <i>— Important, but not sufficient! —</i>	• System Image • Control System • Triadic Semiotic System (pragmatic assessment) • Emphasis on Human Capabilities • Recoding/Chunking • Heuristics = Tricks of the trade • Situation Awareness – Too little meaning • Work Analysis – Constraints & Possibilities • Design focus • Leverage Capabilities • Shape expectations (mental models) <i>— Potential enhancements in perspective! —</i>

One of the well-known mantras of the Human Factors profession has been:

Know thy user.

This has typically meant that a fundamental role for human factors has been to make sure that system designers are aware of computational limitations (e.g., perceptual thresholds, working memory capacity, potential violations of classical logic due to reliance on heuristics) and expectations (e.g., population stereotypes, mental models) that bound human performance.

It is important to note that these limitations have generally been validated with a wealth of scientific research. Thus, it is important that these limitations be considered by designers. It is important to design information systems so that relevant information is perceptually salient, so that working memory is not over-taxed, and so that expectations and population stereotypes are not violated.

The emphasis on the bounds of human rationality, however, tends to put human factors at the back of the innovation parade. While others are touting the opportunities of emerging technologies, HF is apologizing for the weaknesses of the humans. This feeds into a narrative in which automation becomes the 'hero' and humans are pushed into the background as the weakest link - a source of error and an obstacle to innovation. From the

perspective of the technologists - the world would be so much better if we could simply engineer the humans out of the system (e.g., get human drivers off the road to increase highway safety).

But of course, we know that this is a false narrative. Bounded rationality is not unique to humans - all technical systems are bounded (e.g., by the assumptions of their designers or in the case of neural nets by the bounds of their training/experience). It is important to understand that the bounds of rationality are a function of the complexity or requisite variety of nature. It is the high dimensionality and interconnectedness of the natural world that creates the bounds on any information processing system (human or robot/automaton) that is challenged to cope in this natural world. In nature there are always potential sources of information that will be beyond the limits of any computational system.

The implication for designing sociotechnical systems is that designers need to take advantage of whatever resources are available to cope with this requisite variety of nature. For CSE, the creative problem-solving abilities of humans and human social systems is considered to be one of the resources that designers should be leveraging. Thus, the muddling of humans (i.e., incrementalism) described by Lindblom is NOT considered to be a weakness, but rather a strength of humans.

Most critics of incrementalism believe that doing better usually means turning away from incrementalism. Incrementalists believe that for complex problem solving it usually means practicing incrementalism more skillfully and turning away from it only rarely. (C.E. Lindblom, 1979, p. 517)

Thus, for designing systems for coping with complex natural problems (e.g., healthcare, economics, security) it is important to appreciate the information limitations of all systems involved (human and otherwise). However, this is not enough. It is also important to consider the capabilities of all systems involved. One of these capabilities is the creative, problem-solving capacity of smart humans and human social systems. A goal for design needs to be to support this creative capacity by helping humans to tune into the deep structure of natural problems so that they can skillfully muddle through with the potential of discovering smart solutions to problems that even the designers could not have anticipated.

In order to 'know thy user' it is not sufficient to simply catalog all the limitations. Knowing thy users also entails appreciating the capabilities that users offer with respect to coping with the complexities of nature.

This often involves constructing interface representations that shape human mental models or expectations toward smarter more productive ways of thinking. In other words, the goal of interface design is to provide insights into the deep structure or meaningful dimensions of a problem, so that humans can learn from mistakes and eventually discover clever strategies for coping with the unavoidable complexities of the natural world. When framed relative to Rasmussen's Abstraction-Decomposition space, the goal of interface design is to help people navigate through this work domain – to connect the *HOW* with the *WHY*.

References:

Lindblom, C.E. (1979). Still muddling, not yet through. *Public Administration Review*, 39(6), 517-526.

Exploring Deeper

Chapter 8: Multiple Perspectives on Design

Ten years ago, I left academia to join a small design startup, Mile Two. Ironically, after more than 30 years in academia teaching about applied cognitive psychology, human factors and cognitive systems engineering, it feels like my education on design has just begun. I discovered that essential people on our design teams were trained very differently than I was and brought skills and talents to the team that I was lacking. I recently collaborated with some of these teammates on a chapter to reflect on some of the different perspectives that contribute to the development of software to support cognitive work. The first figure was our attempt to summarize some of the different disciplinary perspectives and the second figure is a more elegant revision of that figure suggested by my friend and colleague Fred Voorhorst.




Figure 1. Multiple perspectives on design from Flach, Bennett, Butler & Heroux (In press Handbook of Human Factors).




Figure 2. Multiple perspectives on design by Fred Voorhorst.

These diagrams suggest four classes of constraint that must be considered in the design and development process:

Agency Constraints: These are properties of the "agents" in the system.

Traditionally, the agents have been humans, but as machines becoming increasingly autonomous - we now have to consider the internal constraints of these autonomous agents.

Implementation Constraints: These reflect properties of the computers that will host the software. They include things such as the types of displays and the control/input devices.

Functional Constraints: These are properties of the *work domains* or the *problem spaces* that are the targets for the cognitive work. These include the physical, regulatory, and organizational dynamics that limit the range of possible actions and consequences.

Aesthetic Constraints: These are properties of the psychological and cultural values that determine the degree of satisfaction or frustration associated with using the software products.

All of the four disciplines, at least implicitly, recognize the full range of constraints, however, each discipline tends to emphasize a subset of these constraints in training and practice.

Human Factors [HF] tends to emphasize the constraints associated with human agency. That is, the focus is on specifying the physical [e.g., reach envelopes], perceptual [e.g.,

visual acuity] and cognitive [e.g., working memory capacity] limitations of humans, so that these limitations are respected in the final design.

User Interface Design [UI] or Human Computer Interaction [HCI] tends to emphasize the impact of various computer capabilities on performance. For example considerations include types of displays [e.g., table-top, large-screen, hand-helds, or head mounted], types of controls [e.g., keyboard, mouse, game controller, touchscreen], and types of mediation [e.g., co-located or distributed via internet].

Cognitive Systems Engineering [CSE] focuses on the problem or work domain constraints. This involves describing the means-ends relations in terms of affordances and exploring alternative strategies or heuristics that might be used to efficiently navigate the functional space.

User-eXperience Design [UX] tends to consider the larger psychological and cultural context of use that impacts people's self-image and sense of satisfaction when using the design product.

I come from an HF perspective and was involved in the early development of the CSE perspective. Thus, much of my career has been spent arguing for the value of a CSE perspective for filling gaps in the design space that I felt were not adequately addressed by the HF perspective. The point was not that the HF perspective was wrong, but that it was not sufficient.

However, over the last few years spent developing software products I have come to realize that even the combination of HF and CSE leaves many gaps in the design space that require skills and perspectives associated with UI and UX design. I realize that my sense of the full demands of design were still quite naive and inadequate.

I now regret that I did not have and was not able to give my students a broader sense of the full breadth of the design challenge. Yet, I realize that it would be impractical to cover the full breadth of design in any reasonable course of study. Thus, much of the education must continue beyond the formal course work. And ultimately, I have come to the conclusion that design is a team sport that requires a diverse set of perspectives and people who are open to and who welcome alternative perspectives. To function on such a team I have had to dampen my tendencies to be an advocate for a particular perspective and I have had to heighten my ability to listen to and appreciate the value of alternative perspectives.

Chapter 9: Describing a Work Domain: Function, Abstraction, and Meaning

The concept of a work domain is a fundamental aspect of Cognitive Systems Engineering in relation to work analysis. It involves looking at cognition through a pragmatic lens as advocated by early American Functionalist Psychologists, such as **William James** and **John Dewey**. For them cognition involved making sense of the world in terms of its functional significance: What can be done? What will the consequences be? More recently, **James Gibson** (1979) introduced the word “**Affordance**” to reflect this functionalist perspective; where the term *affordance* is used to describe an object relative to the possible actions that can be performed on or with the object and the possible consequences of those actions. **Don Norman** (1988) has introduced the concept of *affordance* to designers who have found it to be a useful concept for thinking about how a design is experienced by people.

Formalizing Functional Structure

This Functionalist view of the world has been formalized by the philosopher, **William Wimsatt** (1972), in terms of seven dimensions or attributes for characterizing any object; and by the cognitive systems engineer, **Jens Rasmussen** (1986), in terms of an **Abstraction-Decomposition Space**. Figure 7.1 illustrates some of the parallels between these two methods for characterizing functional properties of an object. The vertical dimension of the Abstraction-Decomposition Space reflects five levels of abstraction that are coupled in terms of a nesting of means-ends constraints. The top level, *Functional Purpose*, specifies the value constraints on the functional system – what is the ultimate value that is achievable or what is the intended goal or purpose? As you move to lower levels in this hierarchy the focus is successively narrowed down to the specific, physical properties of objects at the lowest *Physical Form* level of abstraction.




Figure 7.1. An illustration of how Wimsatt's functional attributes map into Rasmussen's Abstraction-Decomposition Space.

An important inspiration for creating the Abstraction-Decomposition Space was Rasmussen's observations of the reasoning processes of people doing trouble-shooting or fault diagnosis. He observed that the reasoning tended to move along the diagonal in this space. People tended to consider holistic properties of a system at high levels of abstraction (e.g., the primary function of an electronic device) in order to make sense of relations at lower levels of abstraction (e.g., the arrangements of parts). In essence, higher levels of abstraction tended to provide the context for understanding **WHY** the parts were configured in a certain way. People tended to consider lower levels of abstraction to understand how the arrangements or parts served the higher-level purposes. In essence, lower levels of abstraction provided clues to **HOW** a particular function would be achieved.

Rasmussen found that in the process of trouble shooting an electronic system, the reasoning tended to move up and down the diagonal of the Abstraction-Decomposition Space. Moving up in abstraction tended to broaden the perspective and to suggest dimensions for selecting properties at lower levels. In essence, the higher level was a kind of filter that determined significance at the lower levels. This filter effectively guided attention and determined how to chunk information and what attributes should be salient at the lower levels. Thus, in the process of diagnosing a fault, experts tended to shift attention across different levels of abstraction until eventually zeroing-in on the specific location of a fault (e.g., finding the break in the circuit or the failed part).

Wimsatt's formalism for characterizing an object or item in functional terms is summarized in the following statement:

According to theory T, a function of item i, in producing behaviour B, in system S in environment E relative to purpose P is to bring about consequence C.

Figure 7.1 suggests how Wimsatt's seven functional attributes of an object might fit within Rasmussen's Abstraction-Decomposition Space. The object or item (i) as a physical entity corresponds with the lowest level of abstraction and the most specific level of decomposition. The purpose (P) corresponds with the highest level of abstraction at a more global level of decomposition. The Theory (T) and System (S) attributes introduce additional constraints for making sense of the relation between the object and the purpose. Theory (T) provides the kind of holonomic constraints (e.g., physical laws) that Rasmussen considered at the Abstract Function Level. These constraints set limits on how a purpose might be achieved (e.g., the laws of aerodynamics set constraints on how airplanes or wings can serve the purpose of safe travel). The System (S) attributes provide the kind of organizational constraints that Rasmussen considered at the General Function level. These constraints describe the object's role in relation to other parts of a system in order to serve

the higher-level Purpose (P) (e.g., a general function of the wing is to generate lift). The Behavior (B) attribute fits with Rasmussen's Physical Function level that describes the physical constraints relative to the object's role as a part of the organization (e.g., here the distinction between fixed and rotary wings comes into play). The Environment (E) attribute crosses levels of abstraction as a way of providing the 'context of use' for the object. Finally, the Consequence (C) attribute provides the specific effect that the object produces relative to achieving the purpose (e.g., the specific lift coefficient for a wing of a certain size and shape).

While the details of the mapping in Figure 7.1 might be debated, there seems to be little doubt that the formalisms suggested by Wimsatt and Rasmussen are rooted in very similar intuitions about how the process of making sense of the world is rooted in a functionalist perspective in which 'meaning' is grounded in a network of means-ends relations that associates objects with the higher-level purposes and values that they might serve. This connection between 'meaning' and higher levels of abstraction has also been recognized by **S.I. Hayakawa** with his formalism of the **Abstraction Ladder**.

Hayakawa used the case of Bessie the Cow to illustrate how higher levels of abstraction provide a broader context for understanding the *meaning* of a specific object in relation to progressively broader systems of associations (See Figure 7.2).




Figure 7.2. An illustration of how Hayakawa's Abstraction Ladder maps into Rasmussen's Abstraction-Decomposition Space.

Figure 7.2 illustrates how the distinctions that Hayakawa introduced with his Abstraction Ladder might map to Rasmussen's Abstraction-Decomposition Space. It has been noted by Hayakawa and others that building engaging narratives involves moving up and down the

Abstraction Ladder (or equivalently moving along the diagonal in the Abstraction Decomposition Space). This is consistent with Rasmussen's observations about troubleshooting. Thus, the common intuition is that the process of sensemaking is intimately associated with unpacking the different layers of relations between an object and the larger functional networks or contexts in which it is nested.

The Nature of Expertise

The parallels between expert behaviors in trouble shooting and fault diagnosis by Rasmussen and observations about the implications of Hayakawa's Abstraction Ladder for constructing interesting narratives might help to explain why case-based learning (Bransford, Brown & Cocking, 2000) is particularly effective for communicating expertise and why narrative approaches for knowledge elicitation (e.g., Klein, 2003; Kurtz & Snowden, 2003) are so effective for uncovering expertise. Even more significantly, perhaps the 'intuitions' or 'gut feel' of experts may reflect a higher degree of attunement with constraints at higher levels of abstraction. That is, while journeymen may know *what* to do and *how* to do it, they may not have the deeper understanding of *why* one way is better than another (e.g., Sinek, 2009) that differentiates the true experts in a field. In other words, the 'gut feel' might reflect the ability of experts to appreciate the coupling between objects and actions with ultimate values and higher-level purposes. Further, this link to value and purpose may have an important emotional component (e.g., Damasio, 1999). This suggests that expertise is not simply a function of *knowing more*, it may also require *caring more*.

Conclusions

As Wieck (1995) noted, an important aspect of sensemaking is what Schön (1983) called *problem setting*. Weick wrote:

When we set the problem, we select what we will treat as "things" of the situation, we set the boundaries of our attention to it, and we impose upon it a coherence which allows us to say what is wrong and in what directions the situation needs to be changed. Problem setting is a process in which, interactively, we name the things to which we will attend and frame the context in which we will attend to them (Weick, 1995, p. 9).

The fundamental point is that the construct of function as reflected in the formalisms of Wimsatt, Rasmussen, and Hayakawa may provide important clues into the nature of how people *set the problem* as part of a sensemaking process. In particular, the diagonal in Rasmussen's Abstraction-Decomposition space may provide clues for how people parse the details of complex situations using filters at different layers of abstraction to ultimately make sense relative to higher functional values and purposes.

Thus, here are some important implications:

- A functionalist perspective provides important insights into the sensemaking process.
- This is the common intuition underlying the formalisms introduced by Gibson, Wimsatt, Rasmussen, and Hayakawa.
- Sensemaking involves navigating across levels of abstraction and levels of detail to identify functional or means-ends relations within a larger network of associations between objects (parts) and contexts (wholes).
- Links between higher levels of abstraction (values, purposes) and lower levels of abstraction (general functions, components and behaviors) may reflect the significance of couplings between emotions, knowledge, and skill.

The various formalisms described here provide important frameworks for understanding any sensemaking process (e.g., fault diagnosis, storytelling, or intel analysis) and have important implications for both eliciting knowledge from experts and for representing information to facilitate the development of expertise through training and interface design.

References

Bransford, J. D., Brown, A. L., and Cocking, R. (2000). *How People Learn*, National Academy Press, Washington, DC.

Damasio, A. (1999). *The Feeling of What Happens: Body and emotion in the making of consciousness*. Orlando, FL: Harcourt.

Flach, J.M. & Voorhorst, F.A. (2016). *What Matters: Putting common sense to work*. Dayton, OH: Wright State Library.

Gibson, J.J. (1979). *The Ecological Approach to Visual Perception*. New York: Houghton Mifflin.

Hayakawa, S.I. (1990). *Language in Thought and Action*. 5th New York: Houghton Mifflin Harcourt.

Klein, G. (2003). *Intuition at Work*. New York: Doubleday.

Kurtz, C.F. & Snowden, D.J. (2003). The new dynamics of strategy: Sense-making in a complex and complicated world. *IBM Systems Journal*, 42, 462-483.

Norman, D.A. (1988). *The Psychology of Everyday Things*. New York: Basic Books.

Rasmussen, J. (1986). *Information Processing and Human-Machine Interaction*. New York: North-Holland.


Schön, D.A. (1983). *The Reflective Practitioner*. New York: Basic Books.

Sinek, S. (2009). *Start with Why: How great leaders inspire everyone to take action*. New York: Penguin.

Weick, K.E. (1995). *Sensemaking in Organizations*. Thousand Oaks, CA: Sage.

Wimsatt, W.C. (1972). Teleology and the logical structure of function statements. *Hist. Phil. Sci.*, 3, no. 1, 1-80.

Chapter 10: Grasping Complexity through Decomposition AND Abstraction





As the classical story about the blind men and the elephant illustrates - the complexity of natural systems generally exceeds our grasp. That is, from any one perspective we can only grasp a part of the elephant. Thus, to get a sense of the whole elephant it is necessary to walk around to explore each of the parts and then to mentally stitch the parts together to get a sense of the whole elephant. The different possible positions of the blind men are analogous to the different disciplines in science. Each discipline owns a part of the elephant, and the challenge is to combine the various perspectives into a coherent understanding of the elephant.

However, for those who can see, there is another approach for getting a sense of the whole elephant. We can move away, increasing our distance from the elephant. As we move farther away, we can see less detail, but we can now see relations among the parts that were not visible from up close. Still, even when the whole elephant is visible, we can only see one side at a time. So, there is no single perspective that allows us to see the whole elephant. There remains the need for some cognitive work to stitch the different perspectives together into a complete understanding of the whole.

Thus, there are two ways that we can change perspectives in exploring the elephant. One way, illustrated by the blind men is through aggregation of the parts. The second way, available to those with eyes, is to change perspective through increasing distance from the elephant. This is a metaphor for abstraction.

One of the key insights of Jens Rasmussen is that to make sense of complex organizations or complex work domains it is necessary to explore through both


decomposition/aggregation and through abstraction. Further, he suggests, based on observations of experts trouble shooting faults in complex technologies, that in terms of thinking productively about complex systems some regions in the abstraction/aggregation space are privileged. In particular, he suggests that at high levels of abstraction the details become less important. For example, there is limited return from reducing a functional purpose into goals, and then further reducing them into sub-goals, which can be further divided into sub-sub-goals. On the other hand, at low levels of abstraction details become more and more important. For example, the shapes of the different components have to fit together in the space available. The threads of one part have to mesh with the threads of another part. The diagrams below are intended to illustrate an exploration space for exploring complex domains jointly by abstraction and decomposition/aggregation.



The key point is that to better understand the whole elephant we must use multiple senses and take multiple perspectives. John Boyd uses the analogy of building a snowmobile to illustrate the importance of combining an analysis mindset (e.g., exploring the components

of different vehicles - motors, steering mechanisms, traction), and a synthesis mindset (e.g., recombining those components to satisfy a different purpose - driving over snow). Decomposition/aggregation reflects an analysis mindset, and abstraction reflects a synthesis mindset. The hypothesis is that to think productively about complex problems one needs to be able to move along the diagonal in this abstraction/decomposition space.

Chapter 11: Work Domain Analysis: A not atypical scenario




Imagine a young human factors engineer, recently graduated from a psychology program at a small Midwestern university (in Dayton, Ohio). He goes to work for a large engineering corporation and is told that he will be participating in a project to design the next generation of decision support for Combined Air and Space Operations Centers (CAOCs). One of the primary goals for this decision support is to facilitate the ability to prosecute time-sensitive targets (TSTs). Where does he start? Perhaps he locates a resident domain expert and arranges a meeting. He introduces himself and begins, "Can you tell me everything that I need to know about TST?" Where does the domain expert begin? How does he capture and communicate a lifetime of experience in a 1-hour meeting? Perhaps, he provides a detailed work analysis that had previously been done, or perhaps he turns to his bookshelf and says, "Start by reading this shelf!"

Suppose that there had been a previous work analysis, how valuable do you think it would be? Would it reflect the new technologies and software options that will be implemented in the next generation CAOC? Would it accurately reflect the demands of future air operations? Even if the analysis did anticipate future opportunities and demands, would the data be in a form that would be useful to the young human factors engineer. That is, could he learn what he needed to learn about the domain from reading the work analysis report? Would the data in the report be represented in a way so that the important dimensions of the work problems would be salient?

Over a 40+ year career exploring human performance in socio-technical systems I have observed or participated in numerous work domain analyses using various forms of representations such as concept maps, flow diagrams, journey maps, and abstraction hierarchies. It has been obvious that the processes involved in collecting the data and

constructing these representations have been enlightening for the people participating in the analyses. However, there is little evidence that anyone who did not participate in the research and generation processes gained much value from the products that were produced. Thus, future research teams rarely benefit from the work of prior teams. The new teams often must start with a blank slate. They must go back to the domain experts and collect their own data and create their own representations.

In 2008 we wrote a chapter in the book "Applications of Cognitive Work Analysis" (Bisantz & Burns) describing a process for creating a living database for archiving the products generated during Work Domain Analysis. We called this database - Integrated Constraint Evaluation (ICE). The general idea was to use the Abstraction Hierarchy to create an indexing system for linking and storing all the products generated during Work Domain Analysis into a relational database. These products could include raw data (e.g., transcriptions from interviews with domain experts, incident reports) and more abstracted data (e.g., concept maps, process diagrams, journey maps, prototypes). The design vision was that an indexing system that was organized around the functional relations across levels of abstraction would allow the products to be easily accessed, and that the index structure itself would help to situate or link the products in a way that would be meaningful, even to people who were not involved in generating the products.




I fear that without such a database a new work domain analysis will be required for every change in a domain. The result will be that cognitive systems and human factors engineers will always be following the design parade, sweeping up the litter from poorly designed systems. The hope is that this living database would allow new researchers and designers to build on the prior work of their peers and that this would allow cognitive systems engineers to be more efficient in responding to evolving challenges and would allow them to make early contributions to the design process.

Since writing the chapter we have been trying to find an organization willing to invest the resources to create something akin to ICE. However, the search has been in vain. It is hard to convince organizations that they own a work domain or that an upfront investment to build such a living database would pay off in the long run.

Perhaps, they're right. Or perhaps this is simply an idea that was ahead of its time. We wonder that in a world where people are investing in digital twins and Retrieval Augmented Generation (RAG) AI architectures, whether an idea of creating a living Work Domain database might be seen as an essential or at least complementary component. This archive could help to document the underlying assumptions and rationality behind a particular technology represented by a digital twin. It could also provide a structure for a database that becomes a component in a RAG architecture.


Despite the lack of support, we remain convinced that the products of Work Domain Analysis could be valuable to organizations. It is a shame that organizations lose the potential value of those products when the people who generated them walk out the door.



To read more: Flach, J.M., Schwartz, D., Bennett, A., Russell, S. & Hughes, T. (2008). Integrated constraint evaluation: A framework for continuous work analysis. In A.M. Bisantz

& C.M. Burns (Eds.) *Applications of Cognitive Work Analysis*. (p. 273 - 297). London: Taylor & Francis.

Chapter 12: Cognitive Systems Engineering: Is it worth the squeeze?



Unfortunately, we run into a variation of this question with every customer and every project. And sometimes, we also hear this from other disciplines who participate with us on design teams. There seems to be a general assumption that any time that a design team is not laying out interfaces (wire framing) or writing code is wasted time. There seems to be a general assumption that time invested to gain a deeper understanding of the nature of the work and to generate multiple hypotheses about alternative representations or innovative ways to employ new technologies is wasted. There is an implicit assumption that the customer can specify exactly what they need – they know the answers – they just need to have someone else write the code for them. Or if they are looking for innovation – customers believe it is possible to get instant solutions without any upfront investment in what Michael Schrage calls “Serious Play.”

Smart people think otherwise:

Give me six hours to chop down a tree and I will spend the first four sharpening the axe. Abraham Lincoln

If I had an hour to solve a problem, I'd spend 55 minutes thinking about the problem and five minutes thinking about solutions. Albert Einstein

I am not sure that the world is more complex today than ever before, but advanced technologies have opened possibilities for dealing with the complexities that have never before been available. And further, it seems clear that organizations who stick to old

strategies (e.g., large inventory buffers, or hierarchical central command structures), who fail to take advantage of new possibilities and new strategies will not succeed in competition with organizations who do take advantage of them.

Today, an auto manufacturer that offers customers “any color they want as long as it’s black” will not be competitive. Today, a military organization that fails to leverage advanced communication networks and advanced approaches to command and control (e.g., mission command) will not be successful.

Today, the value of taking some time to sharpen the axe, or to understand the problem, before hacking away at the tree with a dull axe or before building a solution to the wrong problem should be more apparent than ever. Today, the value of design teams that include a diverse collection of perspectives should be more apparent than ever. We need UI Designers, UX Designers, Programmers, Computer Scientists, Social Scientists, Engineers, and Systems Thinkers. We need multiple perspectives and we need to explore alternative ways of parsing problems and representing constraints.


People trained in CSE typically have experience with multiple disciplinary perspectives (e.g., human factors, industrial engineering, systems thinking). They typically have specialized skills related to work analysis: knowledge elicitation, field research methods, ethnographic methods, problem decomposition and representation methods, decision analysis methods and systems modeling methods.

CSE doesn’t have all the skills or all the answers when it comes to the challenge of designing competitive solutions that leverage the opportunities of advanced information technologies. However, they do bring important skills to the table. A design team that utilizes some of the skills that CSE offers is less likely to waste time hacking away with a dull axe or designing the perfect solution to the wrong problem.

Today – design takes a diverse team. And a team that includes a CSE perspective is more likely to cut sharply and to solve the right problem. *Yes, the juice produced from a CSE perspective is worth the up-front time to do a thorough work analysis, and to explore a variety of ways to represent information to decision makers.*

It's always better to measure twice and cut once!

Chapter 13: Cognitive Systems Engineering: Putting it to use



Yes - following on the previous post - I do believe that Cognitive Systems Engineering (CSE) generates juice that is well worth the squeeze. However, I think that it is important to distinguish between CSE as an academic enterprise exploring basic issues about the nature of work and the nature of human cognition versus a component of a design process.

When implemented in a design process, a CSE work analysis is sometimes mistakenly implemented as a prerequisite to other aspects of design (e.g., prototyping). The problem with this is that work analysis is never done. Work domains are not static - they are constantly changing due to new opportunities and new challenges associated with evolving technologies and operational contexts. Thus, there are always new depths to explore and often one question leads to even more questions. If you delay other aspects of design until the work analysis is complete - nothing will ever get built.

Thus, work analysis should be implemented as a co-requisite to other aspects of design. For example, customers or operators often have a difficult time articulating why they make certain choices, how new technologies might be helpful, or what they need to work more effectively without some concrete context. One way to provide that context is to create concrete scenarios (e.g., critical incidents). Another way is to provide them with a concrete model or prototype that they can manipulate. Even crude models (e.g., back of the napkin sketches or paper prototypes) can be very effective. In the process of reviewing a scenario or interacting with a prototype customers will sometimes be able to recognize and articulate new insights about the utility of the prototype or potential problems with it. This is reflected in Michael Schrage's concept of 'Serious Play.' in essence, prototypes can help to engage operators and allow them to participate in the idea generation process. This

can be a valuable source of knowledge about a work domain. Prototypes can greatly enhance knowledge elicitation and work analysis.

So, it is not a question of doing work analysis OR building design prototypes - success typically requires BOTH work analysis AND prototyping. And further, there is no fixed precedence. Ideally, the work analysis should be tightly coupled with more concrete aspects of design (e.g., wire framing, prototyping). In this coupling, work analysis can be both feedforward (generating hypotheses) and feedback (evaluating operator responses to concrete implementations).

With the modern explosion of technologies for managing complex information, work domains are rapidly changing. This requires a CSE perspective to assess the changing opportunities and risks and to generate alternative hypotheses for how to leverage these technologies more effectively to reduce risks and to stay competitive. This ongoing work analysis can be a resource for designing new interfaces and decision tools, for designing alternative concepts of operation, and for developing more effective training processes. However, design decisions can not wait for this work analysis to be complete, because it will never be complete.

In sum, CSE is both an academic enterprise and a field of practice. As an academic enterprise it focuses on understanding cognition situated in the context of the complexities of work environments. As such, it often challenges the conventional wisdom of a cognitive science based on reductive methods that utilize laboratory puzzles to decouple information processing stages from the dynamics of natural situations. As a field of practice, CSE has to function as a co-requisite of other components of design to probe the complexities of work domains. To be effective in practice, cognitive systems engineers have to learn to be team players and they must be able to coordinate and integrate the work analysis processes with the other design processes.

To be effective in practice cognitive systems engineers must function on interdisciplinary design teams as humble experts, rather than know-it all academics who want to lead the parade.

Chapter 14: Serious Play: A different slant on Design Thinking

Successful innovation demands more than a good strategic plan; it requires creative improvisation. Much of the “serious play” that leads to breakthrough innovations is increasingly linked to experiments with models, prototypes, and simulations. As digital technology makes prototyping more cost-effective, serious play will soon lie at the heart of all innovation strategies, influencing how businesses define themselves and their markets.”

“Serious play turns out to be not an ideal but a core competence. Boosting the bandwidth for improvisation turns out to be an invitation to innovation. Treating prototypes as conversation pieces turns out to be the best way to discover what you yourself are really trying to accomplish. Michael Schrage (1999)

“... generative design research [is] an approach to bring the people we serve through design directly into the design process in order to ensure that we can meet their needs and dreams for the future. Generative design research gives people a language with which they can imagine and express their ideas and dreams for future experience. The ideas and dreams can, in turn, inform and inspire stakeholders in the design and development process. (Sanders & Stappers, 2012, p. 8)

The concept of "Design Thinking" is very much in vogue these days and I share the associated optimism that there is much that everyone can learn from engaging with the experiences of designers. However, my own experiences with design suggest that the label 'thinking' is misleading. For me the ultimate lesson from design experiences is the importance of coupling perception (analysis and evaluation) with action (creation of various artifacts). For me Schrage's concept of "Serious Play" and Sanders and Stapper's concept of "Co-Creation" provide more accurate descriptions of the power of design experiences for helping people to be more productive and innovative in solving complex problems. The key idea is that thinking does not happen in a disembodied head or brain, but rather, through physically and socially engaging with the world.

A number of years ago, I was part of a brief chat with *Arthur Iberall* (who designed one of the first suits for astronauts) and he was asked how he approached design? His response was:


"I just build it. Then I throw it against the wall and build it again. Till eventually I can't see the wall. At that point I am beginning to get a good understanding of the problem."

The experience of building artifacts is where designers have an advantage on most of us. Building artifacts and interacting with the artifacts is an essential part of the learning and

discovery process. Literally grasping and interacting with concrete objects and situations is a prerequisite for mentally grasping them. Trying to build and use something provides an essential test of assumptions and design hypotheses. In fact, I would argue that the process of creating artifacts can be a stronger test of an idea, than more classical experiments. The reason is that the same wrong assumptions that led to the idea to be tested are often also informing the design of the experiment to test the idea.


Thus, an important step in assessing an idea is to get it out of your head and into some kind of physical or social artifact (e.g, a storyboard, a persona, a wireframe, a scenario, an MVP, a simulation).

As a Cognitive Systems Engineer, I am strongly convinced of the value of Cognitive Work Analysis (CWA) as described by Kim Vicente (1999) and others (Naikar, 2013; Stanton, et al., 2018). However, although not necessarily intended by Kim or others, people often treat CWA as a prerequisite for design. That is, there is an implication that a thorough CWA must be completed prior to building anything. However, I have found that it is impossible to do a thorough CWA without building things along the way. In my experience, it is best to think of CWA as a co-requisite with design in an iterative process, as illustrated in the figure below.




The figure illustrates my experiences with the development of the Cardiac Consultant App that is designed to help Family Practice Physicians to assess cardiovascular health. The first phase of this development was Tim McEwen's dissertation work at Wright State University. Tim and I did an extensive evaluation of the literature on cardiovascular health as part of our CWA. I particularly remember us trying to decompose the Framingham Risk equations. Discovering a graphical representation for the Cox Hazard function underlying this model was a key to our early representation concept. With Randy Green's help, we were able to code a Minimally Viable Product (MVP) that Tim could evaluate as part of his dissertation work. Note that MVP does not mean minimal functionality. Rather, the goal is to get sufficient functionality for empirical evaluation in a realistic context with MINIMAL

developmental effort. The point of the MVP is to create an artifact for testing design assumptions and for learning about the problem.




I was able to take the MVP that we generated in connection with Tim's dissertation to [Mile Two](#), where the UX designers and developers were able to refine the MVP into a fully functioning web App ([CVDi](#)). I have to admit that I was quite surprised by how much the concept was improved through working with the UX designers at Mile Two. This involved completely abandoning the central graphic of the MVP, that had provided us with important insights into the Framingham model. We lost the graphic, but carried the insights forward and the new design allowed us to incorporate additional risk models into the representation (e.g., Reynolds Risk Model).



Despite the improvements, there was a major obstacle to implementing the design in a healthcare setting. The stand alone App required a physician to manually enter data from

the patient's record (EHR system) into the App. This is where [Asymmetric](#) came into the picture. Asymmetric had extensive experience with the FHIR API and they offered to collaborate with Mile Two to link our interface directly to the EHR system - eliminating the need for physicians to manually enter data. In the course of building the FHIR backend, the UX group at Asymmetric offered suggestions for additional improvements to the interface representation, leading to the [Cardiac Consultant](#). Again, I was pleasantly surprised by the value added by these changes.



So, the ultimate point of this story is to illustrate a Serious Play process that involves iteratively creating artifacts and then using the artifacts to elicit feedback in the analysis and discovery process. The artifacts are critical for pragmatically grounding assumptions and hypotheses. Further, the artifacts provide a concrete context for engaging a wide range

of participants (e.g., domain experts and technologists) in the discovery process (participatory design).

I have found that it is impossible to do a thorough CWA without building things along the way. In my experience, it is best to think of CWA as a co-requisite with design in an iterative process, rather than a prerequisite.

At the end of the day, Design Thinking is more about doing (creating artifacts), than about 'thinking' in the conventional sense.

Works Cited

Naikar, N. 2013. *Work Domain Analysis*. Boca Raton, FL: CRC Press.

Sanders, E.B. -N, Stappers, P.J. (2012). *Convivial Toolbox*. Amsterdam, BIS Publishers.

Schrage, M. (1999). *Serious Play: How the World's Best Companies Simulate to Innovate*. Cambridge, MA: Harvard Business School Press.

Stanton, N.A., Salmon, P.M., Walker, G.H. & Jenkins, D.P. (2018). *Cognitive Work Analysis*. Boca Raton, FL: CRC Press.

Vicente, K.J. (1999). *Cognitive Work Analysis*. Mahwah, NJ: Erlbaum.


Chapter 15: What does an 'Ecological Interface' look like?

What does an Ecological Interface Design (or an EID) look like?

As one of the people who has contributed to the development of the EID approach to interface design, I often get a variation of this question (Bennett & Flach, 2011).

Unfortunately, the question is impossible to answer because it is based on a misconception of what EID is all about. EID does not refer to a particular form of interface or representation, rather it refers to a process for exploring work domains in the hope of discovering representations to support productive thinking about complex problems.

Consider the four interfaces displayed below. Do you see a common form? All four of these interfaces were developed using an EID approach. Yet, the forms of representation appear to be very different.



What makes these interfaces "ecological?"

The most important aspect of the EID approach is a commitment to doing a thorough Cognitive Work Analysis (CWA) with the goal of uncovering the deep structures of the work

domain (i.e., the significant ecological constraints) and to designing representations in which these constraints provide a background context for evaluating complex situations.

In the DURESS interface, Vicente (e.g., 1999) organized the information to reflect fundamental properties of thermodynamic processes related to mass and energy balances.

The TERP interface, designed by Amelink and colleagues (2005) was inspired by innovations in the design of autopilots based on energy parameters (potential and kinetic energy). The addition of energy parameters helped to disambiguate the relative role of the throttle and stick for regulating the landing path.

In the CVD interface (McEwen et al., 2014) published models of cardiovascular risk (e.g. Framingham and Reynolds Risk Models) became the background for evaluating combinations of clinical values (e.g., cholesterol levels, blood pressure) and for making treatment recommendations.

In the RAPTOR interface, Bennett and colleagues (2008) included a Force Ratio graphic to provide a high-level view of the overall state of a conflict (e.g., who is winning).

Although interviews of operators can be a valuable part of any CWA, these are typically not sufficient. With EID the goal is not to match the operators' mental models, but rather to shape the mental models. For example, the Energy Path in the TERP interface was not inspired by interviews with pilots. In fact, most pilots were very skeptical about whether the TERP would help. The TERP was inspired by insights from Aeronautical Engineers who discovered that control systems that used energy states as feedback resulted in more stable automatic control solutions.

With EID the goal is not to match the operators' mental models, but rather to shape the mental models toward more productive ways of thinking.

A second common aspect of representations designed from an EID perspective is a configural organization. Earlier research on interface design was often framed in terms of an either/or contrast between integral versus separable representations. This suggested that you could EITHER support high-level relational perspectives (integral representations) OR provide low-level details (separable representations), but not both. The EID process is committed to a BOTH/AND perspective, where it is assumed that it is desirable (even necessary) to provide BOTH access to detailed data AND to higher order relations among the variables. In a configural representation the detailed data is organized in a way to make BOTH the detailed data AND more global, relational constraints salient.

For example, in the CVD interface, all of the clinical values that contribute to the cardiovascular risk models are displayed and in addition to presenting a risk estimate (that is an integral function of multiple variables) the relative contribution of each variable is also shown. This allows physicians to not only see the total level of risk, but also to see how much each of the different values is contributing to the risk level.

In configural representations a goal is to leverage the powerful ability of humans to recognize patterns that reflect high-order relations while simultaneously allowing access to specific data as salient details nested within the larger patterns.

The EID process is committed to a both/and perspective, where it is assumed that it is desirable (even necessary) to provide both access to detailed data (the particular) and to higher order relations among the variables (the general).

A third feature of the EID process is the emphasis on supporting adaptive problem solving. The EID approach is based on the belief that there is no single, best way or universal procedure that will lead to a satisfying solution in all cases. Thus, rather than designing for procedural compliance, EID representations are designed to help people to explore a full range of options so that it is possible for them to creatively adapt to situations that (in some cases) could not have been anticipated in advance. Thus, representations designed from an EID perspective typically function as analog simulations that support direct manipulation. By visualizing global constraints (e.g., thermodynamic models or medical risk models) the representations help people to anticipate the consequences of actions. These representations typically allow people to test and evaluate hypotheses by manipulating features of the representation before committing to a particular course of action or, at least, before going too far down an unproductive or dangerous path of action.

Rather than designing for procedural compliance, EID representations are designed to help people to explore a full range of options so that it is possible for them to creatively adapt to situations that (in some cases) could not have been anticipated in advance.

It should not be too surprising that the forms of representations designed from the EID perspective may be very different. This is because the domains that they are representing can be very different. The EID approach does not reflect a commitment to any particular form of representation. Rather it is a commitment to providing representations that reflect the deep structure of work domains, including both detailed data and more global, relational constraints. The goal is to provide the kind of insights (i.e., situation awareness)

that will allow people to creatively respond to the surprising situations that inevitably arise in complex work domains.

Works Cited

Amelink, H.J.M., Mulder, M., van Paassen, M.M., Flach, J.M. (2005). Theoretical foundations for total energy-based perspective flight-path displays for aircraft guidance. *International Journal of Aviation Psychology*, 15, 205 – 231.

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

Bennett, K.B., Posey, S.M. & Shattuck, L.G. (2008). Ecological interface design for military command and control. *Journal of Cognitive Engineering and Decision Making*, 2(4), 349-385.

McEwen, T., Flach, J.M. & Elder, N. (2014). Interfaces to medical information systems: Supporting evidence-based practice. IEEE: Systems, Man, & Cybernetics Annual Meeting, 341-346. San Diego, CA. (Oct 5-8).

Vicente, K.J. (1999). *Cognitive Work Analysis*. Mahwah, NJ: Erlbaum.

Chapter 16: Skilled Muddling

Lindblom (1979) wrote:

Perhaps at this stage in the study and practice of policy making the most common view... is that indeed no more than small or incremental steps - no more than muddling - is ordinarily possible. But most people, including many policy makers, want to separate the 'ought' from the 'is.' They think we should try to do better. So do I. What remains at issue, then? It can be clearly put. Many critics of incrementalism believe that doing better usually means turning away from incrementalism. Incrementalists believe that for complex problem solving it usually means practicing incrementalism more skillfully and turning away from it only rarely.


I am not sure that today is all that different. I think many people still see *muddling* (*incrementalism*) as a flawed approach to problem solving and decision making. It still seems to have a negative connotation and there is an implication that there is a better way to manage organizations or a better way to make decisions.

I think people equate *muddling* with a confused, aimless flailing. They have no concept of *skilled muddling*. However, for me skilled muddling is an apt description of how experts deal with complexity and uncertainty. It is not a confused, aimless process, but a carefully controlled combination of probing the world and utilizing the resulting feedback to make tentative steps in potentially satisfying directions.

The image I have is of skilled rock climbers - who are well tuned to the opportunities that different holds offer and who have a good understanding of the risks. They have developed strength and agility, so that they have a wide range of affordances in terms of what is reachable and graspable. They are careful about insuring that they have appropriate protection to guard against catastrophic falls. They are constantly establishing a firm base of support before reaching for the next hand or foot hold. And they are well-tuned to the information that allows them to identify the best supports and the viable paths toward their goal.

I see many parallels between *muddling* and Peirce's concept of *abduction* - in which experts use their experience to make tentative hypotheses about the consequences of choices and actions; then test their hypotheses through acting on them; and then utilize the feedback to adjust their assumptions and framing to reduce surprise (error). Further I

see this in terms of a closed-loop triadic semiotic process with three sources of constraint. Each of these sources reflect relations between agents and ecologies.



The first of these relations is the now familiar construct of affordance, which represents the potential for action. These reflect the relation between the effectivities of an agent (or organization) and the properties of objects in the ecology. In control systems terms this reflects the constraints of the plant dynamics in the forward loops.

The other two constructs are less familiar. The second construct, *specifying*, reflects the potential for perception or information pick-up. It reflects the relation between the sensor systems (e.g., eyes) of agents and structure in the related ecological mediums (e.g., optical array, or graphical interface). In control systems terms this reflects the constraints on the sensors in the feedback loops.

The third construct, *satisfying*, reflects the functional constraints on performance. This reflects the relation between the intentions, values, or preferences of agents and the actual consequences of particular choices or actions. In control terms this reflects the comparator where intentions are compared to feedback - resulting in a surprise or error signal.

These three components are ontologically independent. That is an affordance exists independently from the information to specify it and from the intentions of agents. Thus, a cup is graspable even though it is out of eyesight and you don't have a current use for it.

However, these three constructs are tightly coupled with respect to epistemology. Thus, for example, our experiences of affordances depend on how well they are specified, and on whether they help us to satisfy our intentions. You cannot actually grasp the cup unless you can sense it (e.g., it is in eyesight), and you are unlikely to grasp it if you don't have a use for it (e.g., to hold a liquid).

Thus, skilled muddling involves tuning the coupling of agents with their ecologies or situations - increasing the range of affordances, tuning the perceptual system and aligning your intentions with the activities that lead to satisfying consequences. The opening figure illustrates how these dimensions map into Rasmussen's intuitions about the design of cognitive systems that are capable of skilled muddling.

These ideas are developed more extensively in our book "A Meaning Processing Approach to Cognition" (Flach & Voorhorst, 2020).

Chapter 17: Muddling through Wicked Problems

Although I have used the term "wicked problems" in my writing, I only recently read Rittel & Webber's (1973) original description of this concept along with an editorial by Churchman (1967) commenting on his hearing Rittel talk about this construct.

I have little to add to the original formulation and encourage others to access and read both papers.

[Rittel, H.W. & Webber, W.M. \(1973\). *Policy Sciences*, 4, 155-169.](#)

[Churchman, C. W. \(1967\) Wicked Problems. *Management Science*, 14\(4\), B141-B142,](#)

Rittel and Webber list 10 attributes of wicked problem, that I will list here, but encourage readers to go to the original source for further explication.

1. There is no definitive formulation of a wicked problem.
2. Wicked problems have no stopping rule.
3. Solutions to wicked problems are not true-or-false, but good-or-bad.
4. There is no immediate and no ultimate test of a solution to a wicked problem.
5. Every solution to a wicked problem is a "one-shot operation"; because there is no opportunity to learn by trial-and-error, every attempt counts significantly.
6. Wicked problems do not have an enumerable (or an exhaustively describable) set of potential solutions, nor is there a well-described set of permissible operations that may be incorporated into the plan.
7. Every wicked problem is essentially unique.
8. Every wicked problem can be considered to be a symptom of another problem.
9. The existence of a discrepancy representing a wicked problem can be explained numerous ways. The choice of explanation determines the nature of the problem's solution.
10. The planner has no right to be wrong.

From the Churchman article:

... the term "wicked problem" refers to that class of social system problems which are ill-formulated, where the information is confusing, where there are many clients and decision makers with conflicting values, and where the ramifications in the whole system are thoroughly confusing. The adjective


"wicked" is supposed to describe the mischievous and even evil quality of these problems, where proposed "solutions" often turn out to be worse than the symptoms. p. B141

Churchman raises some ethical issues in the context of OR associated with approaching wicked problems piecemeal, that I think applies far more broadly than to just OR:

A better way of describing the OR solution might be to say that it tames the growl of the wicked problem: the wicked problem no longer shows its teeth before it bites.

Such a remark naturally hints at deception: the taming of the growl may deceive the innocent into believing that the wicked problem is completely tamed. Deception, in turn, suggests morality: the morality of deceiving people into thinking something is so when it is not. Deception becomes an especially strong moral issue when one deceives people into thinking that something is safe when it is highly dangerous.

The moral principle is this: whoever attempts to tame a part of a wicked problem, but not the whole, is morally wrong.



The construct of *wicked problems* reflects situations where intuitions based on conventional logic will often not be adequate. Wicked problems are chaotic. Thus,

conventional ways to decompose problems based on the intuitions of linear analytic techniques or traditional causal reasoning will fail. It is here where ‘experience’ and ‘wisdom’ are the best guides. It is here where Captain Kirk (a bias toward action), Spock (logic), and Dr. McCoy (emotions) must work together to keep the boat stable amid the waves of epistemic uncertainty.


In dealing with wicked problems the head and heart must trust each other and work together to muddle through. For these situations pragmatics take priority – the right choice is the one that works! And often the only reason it works will be that you did what was required to make it work!

This involves more than classical logic. It often requires going forward and following your intuitions, even when conventional logic says to turn back. It involves passion, persistence, and discipline. For these situations, it is not about making the ‘right choice.’ It is about making your choices work! And typically, this will require the efforts of both head (mind) and heart (body).

Again, it is tempting to argue about whether the head or heart should lead. But this reflects a dualistic trap based on either/or reasoning. It is ultimately a matter of coordination between heart and head. It is not about abandoning analytic aspects of cognition, but rather recognizing that the heart is a necessary partner. Success depends on effective coordination between heart and head in order to make the choices work out right! And still there are no guarantees. Serendipity and luck also play a role.

Without Captain Kirk, Spock and Dr. McCoy might get caught in infinite analytical loops and never pull the trigger, but without Spock and Dr. McCoy, Captain Kirk may not be able to make his choices work! And despite their combined efforts the waves of uncertainty may still get the final vote!

Chapter 18: The Heart of Muddling



Muddling literally involves a trial-and-error process of "feeling" our way through a complex ecology to satisfying outcomes. This involves applying our experiences and utilizing domain constraints (e.g., landmarks) in a smart way. However, muddling involves more than just "knowledge" or "intelligence." It also involves feelings. It takes passion to persist in the face

of inevitable errors. It takes a cool head to face the inevitable risks. It takes discipline to persist when progress is slow. And it takes a well-tuned value system to set appropriate goals and to gauge progress toward those goals.

Information processing models of cognition, based on a computer metaphor have tended to consider emotions to be a primitive source of 'noise' that interferes with successful adaptation and performance. In this context, the difficulties that Phinneas Gage had in adjusting to life after his accident have been attributed to damage to an internal executive that normally suppresses more primitive instincts associated with emotions. However, recent research by Damasio suggests that the difficulties that Gage experienced may be due to a lack of emotional constraints - that is, a kind of autism in which the emotions are cut-off from the executive processing, leading to a cold logical process - that is poorly tuned to the realities of everyday life, that lacks the common sense necessary for consistently achieving satisfaction.

Damasio's work suggests that rather than thinking about emotions as 'primitive' instincts that are a threat to satisfactory functioning, it might be better to think about emotions as evolutionarily tuned instincts that help to ground our newer cognitive capabilities in the realities of everyday situations. Perhaps, the tuning of emotional and aesthetic sensitivities are fundamental to human expertise. This might be part of why Don Norman's claim that

Beautiful things work better!

is true. Perhaps, expertise ultimately depends on our ability to coordinate our hearts and minds. Perhaps, designing technologies to support human expertise requires developing representations that are both well tuned to the domain constraints and to the aesthetic sensibilities of people. Perhaps, it is the intimate intermingling of the pragmatic with the aesthetic that make Pirsig's construct of *Quality* so difficult to explain in conventional terms.

Any philosophic explanation of Quality is going to be both false and true precisely because it is a philosophic explanation. The process of philosophic explanation is an analytic process, a process of breaking something down into subjects and predicates. What I mean (and everybody else means) by the word 'quality' cannot be broken down into subjects and predicates. This is not because Quality is so mysterious but because Quality is so simple, immediate and direct.

References

Damasio, A. (1994). *Descartes' Error: Emotion, reason, and the human brain*. New York: Penguin Books.

Norman, D.A. (2005). *Emotional design*. New York: Basic Books.

Pirsig, R.M. (1974). *Zen and the art of motorcycle maintenance*. New York: Perennial Classics. (p. 254).

Chapter 19: Rational vs Intuitive: Another false dichotomy



When discussing decision making people often make a contrast between rational decision making (Spock) and intuitive decision making (Kirk). Rational decision making is typically

described as deliberate, analytical thinking that can be modeled using classical logic and mathematics (e.g., probability theory). Although humans don't often live up to the ideals of logic or mathematical models - they generally accept the prescriptions of these types of models as the ideals that we should aspire to.

Intuitive decision making is a bit more mystical and is not easily described using classical logic or mathematical models. In fact, the experts who often seem to make decisions intuitively can have difficulty describing how they reached their decision, often attributing their choices to a 'gut feeling.' In many naturalistic situations, experts sometimes don't even recognize that they are making decisions. Rather, they seem to just 'see' or 'recognize' what actions are demanded by a situation with little (if any) conscious deliberation.

The qualitative distinctions lead many to conclude that there are two, distinct types of decision 'mechanisms.' And there is much discussion about which is better and/or which is more or less appropriate for specific situations. For example, many see analytical thinking as the ideal, if the situation allows it; and in contrast intuitive decision is implicitly seen to be 'irrational.' But they might see intuitive decisions as a necessary compromise when there is time pressure that does not permit more analytical thinking.

In contrast, others see intuitive decision making as the hall mark of expertise allowing them to be both faster and more accurate than novices. With extensive experience in a domain, experts seem to be able to automatically tune into solutions with little deliberation. This contrasts with novices in the same domain, who require extensive deliberation and still often fail to perform as well as the experts.

Rather than seeing these as two distinct mechanisms - perhaps decision making is ideally a both/and situation, involving coordination between deliberation (rational analysis) and recognition (intuition/feel). As an analogy - consider that most people utilize two eyes to 'see' the world. Each eye has a different field of view. Does it make sense to consider which view is more accurate? Does the combination allow people to see the world more accurately than either eye alone (e.g., to see depth)? Note that there are conditions known as Binocular Vision Dysfunction in which the eyes are not able to coordinate and fuse the separate perspectives to see the world as a single image and/or to see depth.

Perhaps, we should not be framing decision making as either rational or intuitive. Maybe, we should be talking about 'binocular decision making,' where the quality of the process is a function of the coordination and integration of rational and intuitive processes to get a deeper understanding of a situation. While each of the two perspectives provides different strengths and weaknesses, the combination could be more than the sum of the different

strengths and may actually yield emergent insights (e.g., more depth), needed for success in a complex world.

Chapter 20: The Gordian Knot of Complexity

The literature on leadership suggests that one of the characteristics of effective leaders is a bias toward action. In contrast, poor leadership is often associated with a ‘paralysis of analysis’ in which leaders are unable to act while they weigh all the variables and possibilities that will potentially impact the outcome of a choice. This paralysis often results in loss of opportunities in a dynamic environment where the windows of opportunity are opening and closing.

The penchant for action is illustrated in the legend of Alexander the Great and the Gordian Knot. The Gordian Knot is a metaphor for a complex problem that is intractable when approached by conventional analytical means (i.e., by disentangling it). However, the problem may be easily solved through decisive action (i.e., cutting the knot with a sword).

Many decisions that we make may be effectively Gordian Knots (e.g., buying a house, choosing a career, choosing a mate, deciding to have a family, managing a humanitarian crisis). There are so many factors to consider, the options are so disparate that making comparisons is extremely difficult, and there are often an uncountable number of options to consider (e.g., the perfect house may go on the market tomorrow, interest rates might change).

In addition, there are finite windows of opportunity for effective action (e.g., someone else might buy the house you wanted). Recognizing the challenge of these Gordian Knots, Walker Percy wrote

“Lucky is the man who does not secretly believe that every possibility is open to him.”

Percy knew that this secret belief could result in a ‘paralysis of analysis’ that would ensure a sad and difficult life.

In fact, this paralysis of analysis is exactly what Damasio observed with people with damage similar to Phineas Gage. For example, here is Damasio’s description of one of his patients with ventromedial prefrontal lobe damage:

I was discussing with the ... patient when his next visit to the laboratory should take place. I suggested two alternative dates, both in the coming month and just a few days apart from each other. The patient pulled out his appointment book and began consulting his calendar. The behavior that ensued, which was witnessed by several investigators, was remarkable. For the better part of a half-hour, the patient enumerated reasons for and against each of the two dates:

previous engagements, proximity to other engagements, possible meteorological conditions, virtually anything that one could reasonably think about concerning a simple date.... He was now walking us through a tiresome cost-benefit analysis, an endless and fruitless comparison of options and possible consequences. It took enormous discipline to listen to all of this without pounding on the table and telling him to stop, but we finally did tell him, quietly, that he should come on the second of the alternative dates. His response was equally calm and prompt. He simply said: "That's fine." Back the appointment book went into his pocket, and then he was off.

This clearly illustrates the problem of conventional analytical models of rationality based in logic or normative economic models when faced with the complexities of everyday life. They lack a 'stopping rule.'

The analytic models provide a means for doing the computations, for processing the data, for making comparisons, but the computations will continue blindly as long as data is being fed in. Typically, there are no intrinsic criteria for terminating the computation in order to act. In contrast, heuristics such as those described by Gigerenzer typically have explicit stopping rules. Thus, one of the advantages of heuristics relative to more normative approaches is that heuristics are typically recipes for action, rather than processes for doing computations.

We hypothesize that these intuitive heuristics are the foundations for common sense! They are the swords that allow us to solve the Gordian Knots of everyday life.

In many situations, the quality of the outcome may depend on acting to make the choice right, rather than on waiting to act until the right choice is certain.

References

Percy, W. (1966). The last gentleman. New York: Picador.

Damasio, A. (1994). Descartes' Error: Emotion, reason, and the human brain. New York: Penguin Books. (p. 193-194)

Gigerenzer, G. (2007). Gut feelings. The intelligence of the unconscious. New York: Penguin Books.

Chapter 21: Revisualizing the Decision Ladder in Terms of Flow

To say that Rasmussen's Decision Ladder or SRK Model has had a strong influence on how I think about human performance would be a gross understatement. This model resulted from Rasmussen's observations of how experts did trouble-shooting to diagnose problems in electronic equipment. He observed that often these experts took short cuts that bypassed some of the stages that are typically included in models of human information processing. Figure 20.1 illustrates how Rasmussen represented these short-cuts in the Decision Ladder. He suggested that these short cuts resulted in three different styles of processing: knowledge-based, rule-based, and skill-based.




Figure 20.1

Knowledge-based processing was a slow, deliberate style of analysis that involved all the typical processing stages. This is how people typically visualize rational, analytical

reasoning. *Rule-based processing* utilizes heuristics or tricks of the trade to by-passed the need for deep analysis. Essentially, an observation could trigger a simple rule - if you see a certain pattern or 'sign' then activate a previously learned procedure or action without the need for further analysis. These heuristic short cuts increase efficiency, allowing quicker less effortful responses to familiar situations. Though as Reason later observed, by short cutting the analysis, processing was vulnerable to specific types of errors. Finally, *skill-based processing* was an even more direct short-cut in which perceptions automatically elicited a response completely bypassing the need for information processing. With skill-based processes actions are directly guided by perceptual signals as might occur in a tightly coupled feedback loop. Skill-based processing is even more efficient than rule-based processing, and it also has unique vulnerabilities for error.

Note that many people before Rasmussen had noted similar qualitative differences in performance (e.g., Shiffrin & Shneider's distinction between automatic and controlled processing). However, in my opinion, the SRK model suggested a particularly elegant architecture to account for these qualitative differences.

Recently, I have been reading about Bejan's *Constructal Law* that suggests that all natural living systems from rivers to societies are constantly evolving or self-organizing into structures that maximize flow. For example, the Constructal Law predicts the branching structures of rivers, trees, and lungs. As I considered the implications for thinking about cognitive systems - it seemed obvious to me that Rasmussen's Decision Ladder was describing a flow architecture that seemed to be consistent with the Constructal Law. As a result, I thought it might be interesting to create a representation that made this more apparent. Figure 20.2 is what I created.

I imagined the SRK model as a system that like a river evolves from flowing through a field of possibilities. In this representation the major river is the flow of general knowledge that is global and abstract but can't completely realize many of the performance possibilities that require more detailed direct coupling with the field (knowledge-based flow). Branching off the major river are branches that reflect specialized areas of knowledge (e.g., music theory). These branches have secondary branches that provide heuristics that reflect local constraints within the field of possibilities (e.g., blues patterns) (rule-based flow). These branches have smaller branches that are couple directly to local regions in the field of possibility to trigger skilled actions (e.g., chords and picking patterns).





Figure 20.2

What do you think? Does the alternative representation make sense? I don't intend it to replace Rasmussen's representation, which is brilliant. But I am hoping that having an alternative representation might inspire further discussion and perhaps lead to a deeper understanding of how information and skill flows through cognitive systems

Chapter 22: Heuristics? The good, the bad, and the ugly.



Last week, in the middle of the pandemic, the highly publicized police killings of black men, and the resulting protests and demonstrations I learned of the death of Professor Anders Ericsson. Professor Ericsson was a preeminent psychologist who studied the development of expertise. He was interested in the development of high levels of skill that allow 'experts' to do things that are beyond the capacity of most humans. In particular, his work was instrumental in illustrating that deliberate practice was critical in developing the heuristics that allow experts to become both faster and more accurate in processing information. In essence, because of these *tricks of the trade*, experts are able to avoid the information limits that bound the performance of most humans. This is the positive, good aspect of heuristics. These heuristics allow experts to focus on the patterns (or chunks) that *specify* significant aspects of situations and to coordinate their action to respond *automatically* - quickly and accurately.

The term *heuristic* is also used to describe *biases* in decision making. As the extensive research of Kahneman and Tversky demonstrates, heuristics can lead people to make choices that violate the conventions of traditional logic. Heuristics are a form of bounded rationality that apply within certain domains of activity. Thus, the *automatic* responses that work in one set of situations can seem illogical or *mindless* when they are applied in situations outside that domain. This is the negative, bad side of heuristics - they can lead to performance that Jim Reason referred to as '*strong, but wrong.*'

The ugly side emerges when we put people in domains where the circumstances lead them to develop heuristics or *implicit biases* that not only violate logic, but that violate our cultural values!

My father entered the Marines when he was 17 at the end of WW II. He only served for two years and never left the country. But in his forties, when someone tried to mug him on the

street one night, he *automatically* responded as he had been trained - *slash, kick, gouge!* It thwarted the attack and may have saved his life.

There are many parallels in the training of police and soldiers - to prepare people to respond *automatically* to defend themselves and to survive potentially deadly situations. It is very tempting to attribute these *mindless* responses to evil intent of individuals, but we might consider that these implicit biases are a product of training and socialization. These biases are the result of years of *deliberate practice!*

In some cases, the violent acts that we see on the cell phone videos are the result of many years of deliberate practice, that result in mindless responses to situations. The implication is that the problem of police violence is not simply a function of a few bad apples, but the product of a system that deliberately trains mindless responses to threatening situations.

In hindsight, it is easy to attribute the clearly *mindless* violence to evil intent. But, at least in some cases, consider that these mindless responses are the product of a system of training designed for soldiers, not for peace officers. We want to blame the officer, but this will not ultimately solve the problem of police violence. Ultimately, we must change the system and reconsider what type of skill training is most important for creating expert peace keepers, rather than expert soldiers.

Chapter 23: Change: Going with the Flow

Two questions that interest general systems thinkers are:


Why things change? 2) Why things stay the same?

Let's consider the first question in this essay, and the second question will be explored in later essays. Change is pervasive – things are flowing, moving from place to place, growing, eroding, aging. During my career, my students and I have explored problems associated with the control of locomotion. How do pilots judge speed or changes in altitude? How were they able to guide their aircraft to a safe touchdown on a runway? How are pilots and drivers able to avoid collisions? Or how are baseball batters able to create collisions with a small ball? Inspired by James Gibson, this research explored the hypothesis that people utilize patterns or structures in optical flow fields to make judgments about and to actively control their own motion.

When Gibson introduced the concept of optical flow fields and the hypothesis that animals could utilize optical invariants to guide locomotion, it was a very difficult construct for many perceptual researchers to grasp. Where is optical flow? If you study eyes or if you study light or optics you won't discover optical flow. Optical flow is not in the eye and it is not in the light. Optical flow is an example of an emergent property; it arises from relations between the components. Optical flow results when an eye moves relative to surfaces that reflect light. The flow that results from movement relative to reflective surfaces has structure that specifies important aspects of relations between surfaces and moving points of observation. For example, as you are driving down the highway – surfaces that are near to you will flow by relatively quickly, while surfaces that are far from you will flow by relatively slowly. The further a surface is away, the slower it will flow by – and something that is very far away like the moon may seem to not be moving relative to you at all. And when the surface in front of you begins to expand rapidly – you will immediately recognize that a collision is impending. Thus, the structure in flow fields specifies aspects of the 3-dimensional environment that support skilled movement through that environment.

If you look out the window of a car, a train, or an aircraft you will see the world flowing by - optical flow fields are there right before our eyes. Yet, few perceptual researchers could see it! Or if they did, they didn't recognize it as an important phenomenon to explore.

Few are aware that a key source that helped Gibson to discover the role of optical structure for controlling locomotion was Wolfgang Langewiesche who described how the relation between the point of optical expansion and the horizon helped to specify the path to a safe landing for a pilot.



Interest in change is partly motivated by our interest in making changes – in steering an aircraft to a soft landing or in managing an organization to achieve some function. In the case of controlling vehicles or other mechanical or physical technologies there is typically a linear, proportional relation between action and reaction. Small inputs to the steering wheel typically result in small changes to the direction of travel. However, this is not the case for more complex systems such as weather systems or sociotechnical systems. In complex systems, small actions can be amplified in ways that have large impacts on performance of the whole. The proverbial “Butterfly” effect illustrates this. That is, the idea that the flapping of a butterfly’s wing can impact the nature of a storm some distant time and place in the future. Or that the loss of a nail can result in the loss of a kingdom. Or that a smile can launch a thousand ships.

This is another reminder to be humble.

A sociotechnical system or complex natural system like weather cannot be controlled as simply as steering a car. And in fact, it is questionable whether such systems can be controlled at all. Yes, like the butterfly we can flap our wings, but we can’t anticipate or completely determine the consequences that will result. The performance of a complex organization depends on interactions among many people and no single person determines the outcome. While each component can have an influence, the ultimate outcome will depend on contributions from many people and it will also depend on outside forces and influences.

It is an illusion to think that we can control complex systems - to believe that we are in control or even that we could be in control. In reality, the best anyone can do is to muddle.

If we are observant and careful, we can dip our oars into the water and/or adjust our sails in ways that will influence the direction of our vessel. But the waves and currents have an important vote! And, our colleagues and friends have an impact. It is a mistake to think that we are the lone captains of our fate - but with a little luck and a lot of help from our friends - we can keep the boat upright and muddle through.

Chapter 24: Purpose or Function

We often identify natural systems with specific functions. For example, a function shared by most life forms is to survive and propagate. And this global function typically depends on many sub-functions such as obtaining nourishment, regulating temperature, and avoiding potential predators. Similarly, the people who come together to create sociotechnical organizations are generally collaborating to achieve some purpose or intention that is beyond the capacity of any of the individuals working in isolation. In fact, an organization may have multiple purposes – to provide a service, produce a product, and/or to make a profit.


The concept of *purpose* and its role in shaping performance became a primary consideration for both physical and cognitive systems with the development of Norbert Wiener's Cybernetic Hypothesis and the associated servo-mechanism metaphor. One of the critical attributes of servomechanisms is closely related to the second question mentioned in my previous essay:

Why things stay the same?

In the context of servo-mechanisms this question is typically framed in terms of stability. A system that is stable tends to resist change. For example, healthy animals typically maintain relatively consistent body temperatures, despite changes in the surrounding environments.

Similarly, engineers have designed mechanisms to regulate the temperatures in buildings. These mechanisms allow people to specify a desired temperature and the control mechanism then combines a sensor – that measures the temperature in a room, with an actuator that activates a heating or cooling mechanism whenever the temperature in the room deviates from the desired temperature. This is an elemental form of muddling – in which behavior is adjusted as a function of feedback provided by the sensors. In a similar way, pilots or drivers will adjust their behavior to keep their vehicles on an intended trajectory. For example, skilled pilots will consistently follow a similar safe path to landing at different airports and under different weather conditions. The similar results will be achieved repeatedly – yet the behaviors will be different on every occasion due to the different disturbances (e.g., winds) experienced. A servomechanism has been described as a device that achieves the same results, but in a different way each time.

Note that it is impossible to pre-specify the behaviors that will lead to a safe landing for every situation, because each situation will be different. However, the pilots need to learn or tune their perceptual and motor skills to consistently land safely. The pilots need to learn to recognize what it looks like to be on a safe path and they need to know what actions will counter specific disturbances from that path. Thus, landing an aircraft is much more complex than simply activating a heating or cooling mechanism. And of course, keeping large organizations on the appropriate paths to achieve their missions or goals is even still more complex.



The critical thing to understand about muddling is that it depends on the coupling between perception and action.

Muddling requires both situation awareness (i.e., to be able to sense states relative to goals or intentions, and relative to the capacity for action) and it requires the capacity to act (i.e., to be able to move toward more desirable states or to counter disturbances to prevent undesirable changes in state). Skill is an emergent property of the coupling that reflects relations between situation awareness and the potential for action. It can't be discovered by studying either perceptual or motor components of an organism in isolation. Further, to the extent that the perceptual and motor systems evolved to serve certain functions (e.g., to control locomotion), it is unlikely that these components can be fully understood without considering the functions that they support. This also applies to large organizations – to fully understand the components and processes of the organization it is essential to consider the ultimate functions, purposes, and value systems that the organization serves.

Another important consideration with respect to organizations that have a closed-loop coupling of perception and action is associated with explaining 'why' the organization behaves as it does. To understand an open loop organization, we typically look for causes, which are typically prior events (e.g., the preceding dominoes in the chain of actions and

reactions). However, to explain why a closed-loop organization behaves the way it does, we must ask “what is the function or purpose?” of the organization. In other words, we must ask what the organization’s goals are or what it is attempting to achieve. Whereas, behavior of open-loop systems tends to be the effects of prior causes or reactions to past events, closed-loop systems typically are attracted by goals or in pursuit of potentially satisfying future events.

Thus, the behavior of closed-loop systems does not conform to the cause-effect (or stimulus-response) narratives that have been used to describe many isolated physical systems.

Although change is pervasive, with skilled muddling many organisms and organizations actively counter many undesirable changes and often they can achieve some degree of stability around states that they find to be satisfying. These satisfying states are typically referred to as the *function* or *purpose* of the organizations. For example – achieving and maintaining a specified temperature is the purpose of a temperature regulator; keeping their vehicles in the field of safe travel is the function of pilots or drivers; achieving and maintaining a successful company is the function of corporate executives.


In simple terms, pursuing and maintaining satisfactory states are what organizations do!

However, as suggested in the previous essay, it would be naive to associate the purpose or goals of a complex organization with the ideas in an executive's head or even with the organization's mission statement. These may or may not align with the actual purpose or function of an organization. Unlike the heating control system that is designed by an engineer, complex organizations are self-designing or self-organizing. Thus, stability is an emergent property of organizations reflecting the muddling and mutual adjustments of multiple components.

Ultimately, the proof of the function or purpose is in the doing.

And the doing involves dynamic interactions across all the components, as well as the impacts from the environment. The CEO can't just flap her wings and determine the future

trajectory of the organization. Ultimately, the purpose emerges from interactions involving all the components within the organization.



Chapter 25: Better Maps or Better Muddling?



Many problems of life are ill-structured (e.g., humanitarian or military operations). That is, they are complex and rife with uncertainty, is it possible to make 'good' or 'right' decisions? Perhaps, in these complex situations success may depend on making the decision right (or making the decision work), rather than on making the normatively right decision.

In complex situations, the problem may be more analogous to adaptive control (i.e., making continual adjustments to incrementally satisfy the functional goals –muddling through), than the problem of discretely choosing a right option from a fixed set of alternatives. The loose coupling between a ‘right’ choice and a successful adaptation was illustrated by Karl Weick (1995) using the following story illustrated in the cartoon:

The young lieutenant of a small Hungarian detachment in the Alps sent a reconnaissance unit into the icy wilderness. It began to snow immediately, snowed for 2 days, and the unit did not return. The lieutenant suffered, fearing that he had dispatched his own people to death. But on the third day the unit came back. Where had they been? How had they made their way? Yes, they said, we considered ourselves lost and waited for the end. And then one of us found a map in his pocket. That calmed us down. We pitched camp, lasted out the snowstorm and then with the map we discovered our bearings. And here we are. The lieutenant borrowed this remarkable map and had a good look at it. He discovered to his astonishment that it was not a map of the Alps, but a map of the Pyrenees.

Weick concludes that this story suggests the possibility that “when you are lost, any map will do.” He continues:

The soldiers were able to produce a good outcome from a bad map because they were active, they had a purpose (get back to camp), and they had an image of where they were and where they were going. They kept moving, they kept noticing cues, and they kept updating their sense of where they were. As a result, an imperfect map proved to be good enough. The cues they extracted and kept acting on were acts of faith amid indeterminacy that set sensemaking in motion. Once set in motion, sensemaking tends to confirm the faith through its effects on actions that make material what previously had been merely envisioned [sic]. (p. 55)

Charles Lindblom (1959), in his classic paper “The Science of ‘Muddling Through,’” comes to a conclusion similar to Weick's. Lindblom noted that for complex policy decisions, the comprehensive evaluations that are suggested by normative models of decision-making are typically impossible:

Although such an approach can be described, it cannot be practiced except for relatively simple problems and even then only in a somewhat modified form. It assumes intellectual capacities and sources of information that men simply do not possess, and it is even more absurd as an approach to policy when the time

and money that can be allocated to a policy problem is limited, as is always the case. (p. 79)

Lindblom suggests the idea of an analytically correct map is a fantasy. He offered a more heuristic program of incremental adjustment as a more realistic alternative to the classical, normative approaches to policy making. He called this heuristic, trial and error process: ‘muddling through.’

Twenty-years after his classic paper Lindblom (1979) comments:

Perhaps at this stage in the study and practice of policy making the most common view ... is that indeed no more than small or incremental steps – no more than muddling – is ordinarily possible. But most people, including many policy makers, want to separate the ‘ought’ from the ‘is.’ They think we should try to do better. So do I. What remains as an issue, then? It can be clearly put. Many critics of incrementalism believe that doing better usually means turning away from incrementalism. Incrementalists believe that for complex problem solving it usually means practicing incrementalism more skillfully and turning away from it only rarely.

The question in the title of this post is actually a trick question. It is not a matter of better maps OR better muddling. We will always need both. But it is also important to realize that for many important decisions there is no correct a priori map to guide us. There is no closed-form analytical solution. Yes, the quality of decision making can be improved by utilizing the computational power of AI/ML algorithms to produce better decision aids (e.g., maps). But despite the increasing capacity of these information technologies, significant uncertainty will remain due to the nature of complexity and the difficulty in quantifying all the variables that might be important. Thus, success will ultimately depend to some extent on the skilled muddling of humans. As we begin to integrate AI/ML decision aids into our organizations, it is important to keep in mind that ultimately success will depend on the quality of the *Joint Cognitive System* - that includes smart technologies and smart people.

It seems to me that much of the hype about Artificial Intelligence (AI) reflects the potential power of AI to control things that had traditionally been controlled by humans. If not replacing humans, the expectation is that humans will be displaced from lower levels of control to higher supervisory levels. Thus, the AI systems are framed to be autonomous control systems (e.g., autopilots), with the hidden implication that they have access to the right maps.

I think this vision might be reasonable for simple and maybe even complicated domains of operation. However, I think this is the wrong vision when thinking about applying AI to

complex or chaotic domains of operation where there is no complete map. An alternative that I suggest for applications in domains that are more than complicated is to think about AI as being analogous to an incomplete map that helps guide attention. That is, ***the function of AI is to enhance observability***. The function is to use the power of advanced computations and statistics to pull partial patterns out of data that humans could not otherwise perceive to support orientation and sensemaking.

In this context, the control problem is framed as a joint cognitive system (rather than as an autonomous system). The role of AI in this joint cognitive system is to enhance observability to shape how humans orient to the problem space and frame decisions. In terms of Boyd's OODA loop framework, the value of AI is to enrich Observations in ways that constructively shape how humans Orient to a problem or frame the Decision processes. Thus, humans are engaged and empowered (not displaced), and the ultimate quality of control is enhanced.


References:

- Weick, K.E. (1995). *Sensemaking in organizations*. Thousand Oaks, CA: Sage Publications.
- Lindblom, C.E. (1959). The Science of 'Muddling Through.' *Public Administration Review*, 19, 2, 79-88.
- Lindblom, C.E. (1979). Still Muddling, Not Yet Through. *Public Administration Review*, 39, 6, 517-526.

Chapter 26: Multi-loop Control Systems


As I have discussed before, I was introduced to the quantitative methods for analyzing control systems early in my graduate career and it set an important framework for how I have approached Joint Cognitive Systems. Learning mathematical control theory was a struggle for me and I was so eager to share what I learned with other social scientists that I co-authored a [book on control theory](#) with Richard Jagacinski who was my major graduate advisor. However, as I began to look at joint cognitive systems that were more complex than the laboratory target acquisition and tracking tasks, I soon realized that everyday life is a lot more complex than the laboratory tasks and that the quantitative models that worked for simple servomechanisms and for experimental conditions that required people to act like simple servomechanisms were of limited value.

Everyday life and especially organizational dynamics typically involves many interconnected loops with non-trivial interactions. For example, Boyd's OODA Loop, that is often used to describe skilled behavior, is not a single loop but a multi-loop, adaptive control system.




Joint Cognitive Systems typically have the capacity to learn and to adapt the dynamics of the perception-action coupling to take advantage of that learning. Thus, Joint Cognitive Systems can modify or tune the dynamics of the perception-action coupling to fit the unique demands of different situations. The diagram below is one that I developed to show multiple adaptive loops that reflect some of the different strategies that control engineers have used in designing adaptive automatic control systems (e.g., gain scheduling, model-reference adaptive control, self-tuning regulators). Note that the two different styles of

arrows reflect different functions - the thin arrows reflect the flow of information that is input to and 'processed' through the dynamics of the different components of the system (i.e., the boxes). However, the block arrows actual operate on the boxes and change or tune the processes within the boxes. For example, in an engineered adaptive control system the result might be to turn down or amp up the sensitivity or 'gain' within a control element. Thus, in adaptive control systems the outer loops typically alter the dynamics of the inner loops.



While the previous figure was designed to imagine skilled motor control as an adaptive system the next diagram was designed based on observations of decision making in the Emergency Department of a major hospital. The point was to illustrate some of the ways in which organizations learn and adapt as a result of experience.

Decision Making in the ED



I hope that the preceding diagrams can help readers to get a taste of the complexity of control in the natural world, however I am not fully satisfied with them. I still have a feeling

that these diagrams trivialize the real complexity. More recently I have been inspired by the work of Elenor Ostrum on how communities adapt to manage shared resources and to avoid the "tragedy of the commons," and work by the SNAFU catchers (Allspaw, Cook, & Woods) on Dev Ops and managing large internet platforms. It has been suggested that we have to think about layers of control - or polycentric control.


The groups who have actually organized themselves are invisible to those who cannot imagine organization without rules and regulations imposed by a central authority. (Ostrum, 1999, p. 496)

Each technology shift—manual to automated control to multi-layered networks—extends the range of potential control, and in doing so, the joint cognitive system that performs work in context changes as well. For the new joint cognitive system, one then asks the questions of Hollnagel's test:

What does it mean to be 'in control'?

How to amplify control within the new range of possibilities? (Woods & Branlat, 2010, p. 101)

The following figure is my attempt to illustrate a *polycentric control system*. This diagram consists of three layers that seem to have a rough correspondence with Rasmussen's (1986) three levels of cognitive processing (Knowledge-, Rule- and Skill-based) and Thompson's (1967) three means of coordination within organizations (Standardization, Planning, Mutual Adjustment). These levels interact in two distinct ways - one is passing information through direct communication as is typically represented by lines and arrows in standard processing diagrams. However, the second important way is through the propagation of constraints. I am unaware of any convention for diagramming this. In general, higher levels set constraints on the framing of problems at lower levels - in more technical terms - the higher levels impact the degrees of freedom for action at the lower levels. For example, the standards and principles formulated at the highest level set expectations (e.g., through the way people are selected, trained, and rewarded) for the 'proper' way to do planning and the proper way to act. Or the plans, set expectations about responsibilities and actions at the mutual adjustment level. The constraints typically don't specify the actions in detail - but they do shape the framing of situations and often bound the space of possible actions that are considered.




Although it is not possible to model Polycentric Control Systems using the same mathematics that was used to model simple servomechanisms, there are important principles associated with stability that will generalize from the simple systems to the more complex systems. Perhaps, the most significant of these is the impact of time delays on the stability of these systems and the implication for the ability to pick-up patterns and to control or respond to events. The effective time delays associated with communication and feedback will set constraints on the bandwidth of the system. This is seen in models of human tracking as the Crossover Model - in which the effective reaction time sets limits on what frequencies the human can follow without becoming unstable. However, this constraint is also seen almost universally in natural systems in terms of the $1/f$ characteristic that appears so predominantly when examining performance of many natural systems in the frequency domain. In essence, there are always delays associated with the circulation of information (e.g., feedback) in natural systems and this will always set bandwidth limits on the ability to adapt to situations.

A key attribute of the different layers shown in the diagram of the polycentric control system is that the higher layers will have effective time delays that are progressively longer than the lower layers. On the one hand, this means that high frequency events require the capacity for elements at the mutual adjustment level to be in control (a requirement for subsidiarity). For example, it means that in highly dynamic situations the people on the ground (at the mutual adjustment level) may need to be free to adapt to unexpected situations without waiting for instructions from the higher levels (even if this may require them to deviate from the plan or violate standard operating procedures). On the other hand, this means that higher levels may be better tuned to pick-up patterns that require integration of information over space and time that is outside the bandwidth of people who are immersed in responding to local events (slowly evolving patterns or general principles). Thus, for example an incident command center may be able to provide top-down guidance that allows people at the mutual adjustment level of a distributed organization to

coordinate and share resources with people who are outside of their field of regard. Or standard operating procedures are developed and trained to prepare people at the mutual adjustment level to deal efficiently with recurring situations.

So, I hope this short piece has heightened your appreciation for the complexity of natural control systems and wet your appetite to learn more about the dynamics of complex joint cognitive systems. There is a lot more to be said about the nature of polycentric control systems and the implications for the design and management of effective organizations.

Chapter 27: Polycentric Control, Degrees of Freedom & Resilience



Most organisms and essentially all organizations have multiple layers of interconnected perception-action loops. And it will be generally true that the loops all operated at different time constants, have differential access to information, and have differential levels of authority. One way to think about the coupling across authority levels is that higher levels set the bounds (or the degrees of freedom) for activity at lower levels. In this context, a key attribute of an organization is the *tightness* of the couplings between different levels.

For example, the classic Scientific Management approach involves a very tight coupling across levels. With this approach, higher levels (e.g., management) are responsible for determining the 'right way' to work and the management levels have the responsibility to implement training and reward systems to ensure that the lower levels (operators - workers) strictly adhere to the prescribed methods. This is an extremely tight coupling leaving workers very little degrees of freedom, as they are typically punished for deviating from the "one best way" prescribed by management. In essence, the organization functions as a clockwork mechanism. This approach can be successful in static, predictable environments, where the assumptions and models of managers are valid. However, in a changing, complex world (e.g., VUCA environments) this approach will be extremely brittle - because the cycle times for the upper levels of the system will always be too slow to keep pace with the changing environment. This approach will surely fail in any highly competitive environment, with intelligent, adaptive competitors.

One way to loosen the coupling between levels is to replace the "one best way" with a playbook, or a collection of smart mechanisms that are designed to address different situations that the organization might face. Typically, higher levels develop the playbook

and are responsible for training the plays and for rewarding/punishing workers as a function of how well they are at selecting the right play for the situation and implementing that play.

The coupling can be relatively tight if the playbook is treated as a prescription such that any deviations or variations from the plays are punished, especially when negative results ensue. Or the coupling can be looser if the plays are treated as suggestions, deviations (e.g., situated adaptations) are expected or even encouraged, and the lower levels are not punished for unsuccessful adaptations.

A still looser coupling is an approach of *Mission Command*. With Mission Command higher levels in the organization set objectives and values, but they leave it to lower levels of the organization to determine how best to achieve those objectives given changing situations. This approach requires high levels of trust across levels in an organization. Higher levels must trust in the competence and motivations of the lower levels, and lower levels must trust that higher levels will have their backs so they will not be blamed or punished when well-motivated adaptations are not successful.


Also, Elinor Ostrum's construct of polycentric governance describes how loosely coupled social systems self-organize to manage share resources in a way that avoids the tragedy of the commons. Note that Ostrum's work contradicts the common notion that top-down control is essential to prevent competition between local interests from resulting in complete exhaustion of a resource. Her research discovered many situations where coordination and cooperation emerged bottom-up without being imposed by a centralized, external authority.

These are simply some examples of different levels of tightness in the couplings between levels intended as landmarks within a continuum of possibilities. On one end of the continuum are tightly coupled organizations where the couplings across levels are like the meshing of gears in a clock. As couplings become looser the organization becomes less machine-like and more organismic. Metaphors that are often used for looser couplings include coaching and gardening. In organizations with looser couplings higher levels introduce constraints (e.g., guidance, suggestions, resources), but don't determine activities. The higher levels tend the garden, letting the plants (workers) free to grow on their own.

In an increasingly complex world - organizations with tight couplings across levels will tend to be brittle and vulnerable to catastrophic failures due to an inability to keep pace with changing situations. When couplings across levels are looser, the potential for bottom-up innovation and self-organization increases. However, if the coupling is too loose - there is an increasing danger for inefficiencies, conflict, and chaos. Thus, organizations must

continuously balance the tightness of the couplings trading off top-down authority and control to increase the capacity for exploration and innovation (adaptive self-organization) at the operational level. In a complex world, loosening the couplings can allow the organization to muddle more skillfully.

Chapter 28: Degrees of Freedom: A blessing and a curse



At the heart of any consideration of questions of controllability and resilience is the question of degrees of freedom. This was a central question of concern for Nikolai Bernstein in his explorations of perceptual-motor skill. In the context of motor skills, degrees of freedom refer to the constraints/possibilities for moving our bodies. But more generally - one might think of the degrees of freedom in a system as setting bounds on the potential for action. When degrees of freedom are high - then there will be many different ways to accomplish the same task or goal.

Generally, the more degrees of freedom a system has - the more resilient it can be. This is because when one route to a goal is blocked, there will be other paths for reaching that goal. However, degrees of freedom pose a challenge for control - because the more degrees of freedom a system has - the more variables have to be considered in the control logic. In essence, increasing the degrees of freedom increases the potential for disturbances that have to be managed by the controller.

So, the degrees of freedom problem refers to a trade-off between flexibility (resilience potential) and control complexity. More degrees of freedom mean there are more ways to be successful in accomplishing a goal, but also more ways to fail. In the context of perceptual motor control, the constructs of *coordinative structure* and *smart mechanism* were introduced as hypotheses of how to effectively manage this trade-off.

A coordinative structure or smart mechanism is created by constraining or locking out a subset of the degrees of freedom to create a simpler 'mechanism' (less complex control problem) that is specialized for a particular task or situation. This is somewhat analogous to the smart heuristics described by Gerd Gigerenzer in relation to ecological rationality.


The skill of golf provides a good example of how skilled athletes manage the degrees of freedom problem. Trying to drive a golf ball is a difficult control problem because of the many degrees of freedom in the body. For example, there are many potential disturbances to the path of the club head during a swing - a bend at the elbow of the leading arm, a twist

of the head or shoulder, the positions of the legs, etc. It is simply impossible for a person to control all of the different degrees of freedom in real time. So, learning to drive a golf ball consistently involves learning to lock out many of the degrees of freedom (e.g., locking the elbow of the leading arm, holding the head in a fixed position with eyes on the ball, etc.). Thus, pro golfers learn to become a simple 'driving machine' that involves realtime control of only a few degrees of freedom that are specifically chosen to generate power to hit the ball great distances in a specific direction.

So, what good is having a lot of degrees of freedom if we are going to lock them out. The high degrees of freedom allows golfers to become many kinds of simple mechanisms. They can lock out a different set of degrees of freedom to become a chipping mechanism, designed to hit a ball with greater control over distance and direction. Or they can lock out another set of degrees of freedom to become a putting mechanism. Thus, a pro golfer will learn to become many different simple control mechanisms that are specialized for different situations they will face on a golf course. Successful golfers not only have different types of golf clubs - but also have specialized control mechanisms that have been 'designed' through extensive practice for specific situations.

Thus, to be resilient it is desirable for organizations to maximize degrees of freedom to increase the space of potential possibilities or to increase the possible routes to satisfying ends. However, the organizations also need to develop (train) smart mechanisms tuned to the demands of different situations - so that they will not be overwhelmed by the need to control too many variables in realtime. And they have to learn to choose the right mechanism for the right situation. Just as *smart heuristics* reduce the computational burden, while maintaining high levels of effectiveness with regards to decision making, *smart mechanisms* reduce the computational demands on control, while maintaining high levels of effectiveness with respect to coordinated action.

Chapter 29: Polycentric Control: Another View



I am still struggling to find the right way to represent the multiple layers of constraint that shape the dynamics of organizations. The image above is framed in terms of a nuclear power plant. As with previous depictions, I include three layers of constraint - recognizing that there may be other layers within layers. However, I think these three layers represent important distinctions that are common across many different organizations.

In this representation - the intention is to emphasize that the layers are nested - such that outer layers set the context or the degrees of freedom for operations at inner layers. Each layer is monitoring operator actions and plant performance - but they are filtering the information to reflect different spheres of action, different degrees of granularity, and different time spans.

At the **Industry Management** level the focus is on standards associated with regulation, investment, and training. At the **Plant Management** level the focus is on anticipating events, planning, and supervision. At the **Operator** level, the focus is on operating, monitoring, and maintaining the equipment.

A key aspect of the coupling across the layers is the degrees of constraint. That is, how loosely or tightly the inner loops are constrained by outer loops. The stability or success of an organization will be a function of finding the right balance of constraint relative to the demands of the work environment (e.g., the pace of change, the degree of complexity/uncertainty). The balance is reflected in attributes such as subsidiarity and

trust - which describe whether organizations give more or less discretion (degrees of freedom, decision authority) to inner layers.

A general principle is that when the pace of change or the degree of uncertainty is high, organizations that give more discretion/flexibility to people in the inner layers to directly adapt to situations will be more stable. This is largely due to the longer time constants required for sensemaking at the outer layers. In other words, when events on the ground are rapidly changing in surprising ways, there is no time to wait for direction from upper levels of management. The shifting of military organizations such as the US Marines toward more Mission Command styles of leadership is motivated by a desire to be more effective in a highly dynamic, uncertain domain.


On the other hand, when the work domains are relatively stable, the broader perspective and longer integration periods (i.e., broader experience) of outer layers are less likely to lead the organization to over-react to transient changes (e.g., chase fads). Thus, more top-down control or regulation leads to stable, more efficient performance. For example, commercial aviation and nuclear power domains typically benefit from more top-down regulation and constraint that reflect extensive experience in relatively stable domains of operation.

It is important to recognize that the effectiveness of a polycentric control system (or any control system) can only be judged relative to the functional demands of specific work contexts. So, be skeptical of any models that propose to set the 'ideal' or 'optimal' general recipe for success. Solutions that may have been well tuned to the past, may become unstable when the demands of the work domains change.

Chapter 30: Re-envisioning Boyd's OODA Loop as a Layered (Polycentric) Adaptive Control System

The figure below is an attempt to illustrate three important dimensions of any cognitive system as layers in a polycentric control system. The three layers are designed to reflect John Boyd's dimensions of the strategic game: Physical, Mental, and Moral. Note that all three layers are closed loops (effectively OODA loops) that get feedback as a result of the consequences of acting on the ecology. Further, the outer layers set up constraints or determine the degrees of freedom available to inner layers. Finally, it is important to recognize that each layer is working continuously and simultaneously with the other layers, but at a different time scale (i.e., these are not sequential operations).

The inner most layer (green) in this diagram represents the direct coupling with the ecology that reflects perceptual-motor skills. This “represents the world of matter-energy-information all of us are a part of, live in and feed upon” (Osinga, 2007, p. 210). This type of coupling was the focus of J.J. Gibson's research and constructs such as affordance and optical invariant play important roles. Much of the action at this level is ‘automatic’ requiring little conscious awareness.



The next outer layer (orange) in this diagram reflects conscious thinking related to problem solving and decision making. At this level an intention is formed that will frame subsequent actions in the layer below. This “represents the emotional/intellectual activity we generate to adjust to, or cope with, that physical world” (Osinga, 2007, p. 210). The functioning of

this layer is the focus of researchers such as Gary Klein (Recognition-Primed Decision Making, Naturalistic Decision Making) and Gerd Gigerenzer (Ecological Rationality). Important constructs to consider at this level include abduction and heuristic. The intention provides the framing for attention, recognition, and the possible actions at the skill-based level.

The most outer layer (blue) in this diagram reflects the influence of long-term experience and learning. It “represents the cultural codes of conduct or standards of behavior that constrain, as well as sustain and focus, our emotional/intellectual responses” (Osinga, 2007, p. 210). This layer reflects the values, principles, assumptions, and beliefs that shape performance at the other layers. This layer reflects the largely ‘implicit knowledge’ that shapes analysis and synthesis at the conscious thought level below.

This system is a multilayered, adaptive control system that is continually tuning to achieve a stable relation with the ecology. This tuning is happening simultaneously at all three levels, but at different time scales. Ultimately, stability depends on coordinated interaction across the three levels. The general notion is consistent with C.S. Peirce's logic of *abduction*, Piaget's constructs of *assimilation* and *accommodation*, EJ Gibson's construct of *attunement*, and Friston's *Free Energy Principle*. Thus, the dynamic is essentially one of minimizing "surprise" (i.e., reducing uncertainty) relative to the continuously changing Ecology.


The unique feature of this model relative to other models is that most models envision the interactions between levels as sequential. That is, there is a precedence relation between the three layers so that the outer layers need to *infer* the 'ecology' based on cues provided by the inner layer. In contrast, the model presented here suggests that all three levels have direct feedback of the consequences of action - though this feedback is integrated over different time constants. In essence - each layer is tuned to different bandwidths (in terms of the patterns that matter most).

Chapter 31: Boyd: An Organizational Perspective

In the previous post - I offered a model for a cognitive system in which I attempted to integrate the three dimensions of the Strategic Game described by Boyd (Moral, Mental, and Physical) as nested layers in a polycentric control system. A unique and critical aspect of this model is that each layer functions effectively as an OODA loop closed through the ecology. That is - each layer is involved in directly interacting with the ecology. The difference between layers is the time span over which they are integrating information. The outer layer in the model is integrating across events - to derive general principles that have broad implications for stability. The middle layer is integrating over a specific event - typically taking into account patterns that allow it to recognize possibilities, to apply heuristics, to plan, and to steer activities toward a satisfying outcome. Finally, the inner layer is controlling actions and reacting to stimuli and situations in the context of the beliefs and intentions specified by the outer layers.

However, I think this is difficult for many to grok, because it is difficult to get past conventional assumptions about information processing. Information processing models have typically been formulated as a series of information processing stages with a precedence relation such that later stages in the process only "see" the products of earlier stages. Thus - the deeper stages are required to utilize inference to construct an 'internal model' of the actual ecology from the 'cues' provided by the earlier stages. In other words - there is typically assumed to be a gulf (blanket or filter) that separates our mental experiences from our direct physical interactions with the ecology. Many people can't imagine that abstract principles or beliefs such as "do unto others as they would do unto you" could be actually tested directly as a result of the consequences of actions in the ecology.

However, I hypothesize that this cognitive interplay between moral, mental, and physical dimensions is scale independent and that it applies equally to individuals and to organizations as cognitive systems. Thus, the figure below frames exactly the same model - but in terms of organizational dynamics. The hope is that the organizational perspective will be a little easier for people to grok, because the interplay of the different levels is more observable in an organization than in an individual organism.



In this diagram the inner loop (green) represents the front-line workers who are directly acting in the ecology. For example, these might be the first responders (fire, police, and medics) who are directly involved in managing the emergency response.

The next outer layer (orange) represents the incident command center. This includes the incident commanders and managers who are monitoring the situation on the ground, who are managing the logistics, and who are developing plans to allocate responsibility and coordinate the actions of the first responders on the ground. Note that the incident command is getting information from the responders on the ground, but they also have access to information about the ecology (e.g., the availability of external resources) that is not available to the people who are immersed in the details of local situations.

The outer most layer (blue) represents the organizational homes of the various first responders (e.g., the fire department, the police department and the medical system). This outer system is less active in dealing with the specific immediate emergency event - but this loop will be responsible for after action analysis of the event, where it will have information about the ecology that might not have been available to either of the inner layers during the event. This layer is responsible for evaluating the response in relation to other events (e.g., after action reviews and historical analyses of prior events) and to extract the lessons and general principles that will potentially improve responses to future events. These lessons will then be

integrated into training processes that will shape how planning and emergency responses will be conducted in future events.

A key aspect of the dynamic is that the (beliefs, intentions, and actions) at the different levels are shaping and being shaped by each of the other layers - though the changes are happening at different time scales. Each layer is operating at a different tempo or bandwidth than the other layers, and most importantly this means that each layer will be able to 'see' or 'pick-up' patterns in the ecology that cannot be seen in the other layers. And it is important to emphasize that these patterns are not 'in the head' but are emergent properties of the direct physical interaction with the ecology. Each layer is an abductive system in that it is forming hypotheses that are then directly tested by future actions in the ecology (not by inferences on a mental model). Every loop is 'tested' against the practical consequences it has toward achieving stable (i.e., satisfying) consequences in the actual ecology.

In sum, I am proposing that the classical partitioning of cognitive systems into a sequence of processing stages is wrong. Yes - information is being processed at multiple different levels - but each level functions much like an OODA loop and the stability of the complete system is dependent on whether the various loops shape interactions with the ecology in productive or satisfying directions.

Chapter 32: A Polycentric Control Story

Three paramedics are driving along a rural road in central Ohio when a truck in front of them suddenly seems to go out of control causing several other cars to crash, then hitting some pedestrians who were skating along the road and eventually hitting a tree. The paramedics get out to survey the scene and two of them each go to different victims and begin treating them, the third paramedic surveys the situation and heads back to the car, where he remains throughout the event.

What is he doing in the car?

In fact, after the event was over, the other two paramedics were puzzled, and ask him why he did not immediately begin treating the other victim? His response was "which other victims?" It turns out that neither of the other two paramedics realized that there were more than three victims.

What was the third paramedic doing in the car? He was taking incident command. He realized that for some of the victims to survive it would be essential to get them to trauma hospitals, where they could get more extensive treatment, within the "golden hour." He realized, that to ensure that all the victims were saved they would need to get ambulances and a care flight helicopter into the scene as quickly as possible. He was on the radio calling for support and providing them directions so that they could get to the rural scene as quickly as possible. This included identifying a potential landing site for the care flight helicopter.

While he was in the car, a number of other people stopped to offer help, including a nurse. He was able to direct these volunteers. Asking the nurse to attend to one of the injured and directing another person to attend to the truck driver who was wandering from his truck in a daze.

After the incident, this paramedic explained to the other two, "Yes, I could have started to treat one of the victims, but I wanted to make sure that all the victims survived. And I realized that to do that, we needed more resources and somebody would have to take command and coordinate these resources."


This story illustrates the functions related to three different layers in a polycentric control system.

At the *mutual adjustment level*, two of the paramedics, the nurse, and other volunteers were directly acting to address the immediate demands of the situation.

At the *planning level*, the third paramedic was functioning as an incident commander. He wasn't directly treating patients, but he was attending to larger patterns and trends to anticipate needs and to coordinate the resources that would ultimately be critical for achieving a satisfying outcome.

And the reason this paramedic took this role was because he had been trained in the National Incident Management System which had been developed based on generations of firefighting experience with large forest fires. This has become a national standard and many fire and police departments are required to have NIMS training. This establishment and training of a standard reflects contributions at the *Standardization or Knowledge-based layer* of the emergency organizations. Additionally, the idea of a 'Golden Hour' is a principle based on extensive experiences in emergency medicine. The Standardization Layer reflects a capacity to integrate across an extensive history of past events to pull out principles that provide guideposts (structure, constraints) for organizing operations at the other levels.

The figure below is an alternative way to represent a polycentric control system to the [layered diagram](#) used in prior essays. This representation was developed based on observations and interviews with emergency response personnel. Here the layers are illustrated as nested control loops. These loops are linked through communications (the thin lines) but also through the propagation of constraints (the block arrows), in which outer loops shape or structure the possibilities within the inner loops (e.g., standard practices, command intent, a plan, distributing resources). The coordination between layers is critical to achieving satisfactory results and this coordination depends both on communications within and between layers and the propagation of constraints (sometimes articulated as common ground, shared expectations, shared mental models, or culture).



Note that neither the previous representation showing layers or the present representation showing nested control loops is complete. Each representation makes some aspects of the dynamic salient - while hiding other aspects of the dynamic and perhaps carrying entailments that are misleading with respect to the actual nature of the dynamic. An assumption of general systems thinking is that there is no perfect representation or model. Thus, it is essential to take multiple perspectives on nature to discover the invariants that matter - to distinguish the signal from the noise.

Three important points:


- **First** - the power of a polycentric control system relative to addressing complex situations is that each layer has access to potentially essential information that would be difficult or impossible to access at other layers. However, without effective coordination between the layers some of that information will be impotent.
- **Second** - it is easy for people operating within a layer to have tunnel vision, to take the functions of the other layers for granted, and to underestimate the value that the other layers contribute.
- **Third** - be careful not to get trapped in any single perspective, model, or metaphor. Be careful that your models don't become an iron box that you force the natural world to fit into. Be careful not to fall prey to the illusion that any single model will provide the requisite variety that you need to regulate nature and reach a satisfying outcome.

Now, the rest of the story: All the victims from the accident above survived.

Over the Edge

Chapter 33: Truth? A Pragmatic Perspective on Meaning

Is there such a thing as an 'objective truth' that is independent of an individual's experience? For many, the point of science is to discover such a truth that is invariant across observers. On the other hand, one might ask whether the unique subjective experiences of individuals are valid?



Consider the question of whether a frozen pond will support an individual. We can 'objectively' measure the thickness of the ice. We can do experiments to objectively determine what weight the ice will support. But we can't say 'objectively' whether the ice will afford walking, because the experience of 'support' depends both on the thickness of the ice and the weight of an individual. The same frozen pond can mean different things to different individuals. So, the affordance of support or the functional meaning of the ice cannot be specified independently from an observer's weight.

Thus, individuals of different weights might experience the frozen pond differently. The lighter individual might validly see the pond as a safe field of travel, and the heavier individual might validly see the pond as an unsafe field of travel. The subjective experiences are different, but both might be valid.

On the other hand, the experience for both individuals might be ambiguous. The lighter individual might experience the pond as potentially dangerous, and decide to walk around the pond, rather than take a more direct path across the pond. Or the heavier individual might perceive the pond as a safe path but then discover his error when the ice cracks and he falls through.

If the pair walks together, the lighter individual might convince the heavier individual that the pond is safe (which it is for her) and they might then both walk together across the ice, and both fall into the lake when the ice fails to support their combined weight.

Thus, different experiences might be valid - though different. And common experiences/interpretations might be valid for one individual, but not for another or not for the group - though the interpretations are identical.

Note in an attempt to achieve objectivity, science establishes measures/rulers like pounds and inches that are observer independent. But these measures are not invariantly related to peoples' experiences. Do people experience the world relative to these observer independent metrics/standards? Or do people experience the world pragmatically? Does a surface support locomotion? Is a gap pass-through-able? Is an object reachable or graspable? We don't directly experience the world in pounds or inches. Rather we experience the world directly in functional terms - we directly experience what the world affords. The measures of the physicist are abstractions - not direct experiences!

The observer independent metrics serve the physical sciences well, but are they appropriate for the social sciences? Does it make sense to gauge whether people's experiences or perceptions are 'right' or 'wrong' using the metrics of the physical sciences? This is the question that James Gibson asked? And he concluded that perception is not based in "the world of physics, but the world at the level of ecology." Similarly, Jacob Johann von Uexkull introduced the construct of 'umwelt' to illustrate how functional ecologies can be describe in relation to the unique capabilities of different animals.

The implication of this is that the truth, validity or meaning of our experiences is not observer independent. Yet, they also are not arbitrary. The fact of whether the frozen pond will support us is a pragmatic truth, and the ultimate proof is in the doing. The ultimate tests of our perceptions (i.e., the meanings that are most significant) are the potential consequences of our actions!

Chapter 34: Getting Past the Boxes

Since the introduction of the information processing model of human performance - box diagrams have been a dominant form for representing cognitive processes. Unfortunately, these diagrams have many misleading entailments:

- The boxes (or circles) suggest that the 'system' can be reduced to a number of relatively independent processes that can be studied in isolation.
- The diagrams suggest that the processes can be ordered, such that some processes have precedence over other processes. This is typically implied by the left to right sequence in the diagram.
- Related to the previous point, there is an implication that performance depends on a sequence of discrete steps in which the output from one process is passed discretely to the next (analogous to a string of dominoes).
- The boxes tend to be visualized as processes that can be localized in the head of an agent, and a significant component of the circular dynamic is generally missing. That component is the ecology (or situation) that perception and action are closed around.

The problems associated with these entailments are discussed throughout this book. But here, I simply want to suggest an alternative form of representation. To illustrate an alternative, I will use Rasmussen's Decision Ladder Figure 34.1 which I think captures aspects of cognition that are not well represented in other box diagrams – even though Rasmussen utilized a box diagram.

There are several aspects of the Decision Ladder that I find valuable:

- The Decision Ladder shows that there are many different paths linking perception (the left side of the ladder) with action (the right side of the ladder)
- The Decision Ladder also suggests that there are different levels or depths of analysis corresponding to knowledge-based, rule-based, and skill-based processing.
- The levels help to account for important qualitative differences in human performance – such as 'automatic' and 'controlled' processing.




Figure 34.1. Rasmussen's (1986) Decision Ladder

However, there are also misconceptions associated with the Decision Ladder. The most significant problem is that the levels are often treated as if they are discrete categories for classifying performance. That is, that performance is either knowledge-based, rule-based, or skill-based. When the reality is that performance is always simultaneously shaped by all three processes. For example, the performance of a musician simultaneously involves knowledge-, rule-, and skill-based processes. The knowledge and rule processes don't stop when the fingers of the guitarist are performing at a skill-based level. Information is always simultaneously flowing through all levels of the ladder. However, the levels are operating at different time constants.

To combat some of the potential misconceptions associated with the box diagram representation of the decision ladder, I suggest that the levels should be represented as a nesting of constraints such that knowledge-based processing sets the constraints on rule-based processing, and rule-based processing sets the constraints on skill-based processing (Figure 34.2). For example, a musician's knowledge suggest what rule or heuristics are relevant to playing a particular style or musical piece. In turn, the rules or

heuristics guide attention in a way that allows actions to be 'automatically' triggered by events. And of course, the resulting quality of performance provides feedback that is simultaneously being processed at all three levels but at different time scales (or different chunks).




Figure 34.2

It is also important to realize that the mental processes are being constrained by a domain (e.g., musical performance) that includes the physical properties of the instrument, the physics of sound, and the cultural expectations of an audience.


Again - there is a lot to unpack here. But hopefully there is enough here to whet your appetite for further discussion and exploration. Chapter 29 presents a similar nested control model, but for an organization. In an organization the nesting of levels is more observable than within a single organism. In an organizational context, James Thompson (1967) discusses three levels for coordinating action: Standardization, Planning, and Mutual Adjustment. The functions and means of constraint for these levels are similar, if not analogous, to the levels in Rasmussen's Decision Ladder.

References:

Rasmussen, J. (1986). *Information Processing and Human-Machine Interaction: An approach to Cognitive Engineering*. New York: North Holland.

Thompson, J. D. (1967). *Organizations in action: Social science bases of administrative theory*. McGraw-Hill.

Chapter 35: The Myth of 'Control'



In 1978, Rich Jagacinski hired me as a graduate research assistant to help with a project to compare peoples' ability to track a target based on kinesthetic feedback with the ability to track the same target based on visual feedback (1). I was totally unprepared at the time, and even after 6 years of graduate school, I was only beginning to understand the technical aspects of 'control theory.' However, eventually, I was able to co-author a text with Rich to help introduce the technical aspects of control theory to other social scientists (2). But this was only the beginning of a long journey to explore the dynamics of complex couplings of humans, technologies, and ecologies.

Eventually, I have come to the conclusion that *control theory* has absolutely nothing to do with 'control.' Further, the metaphor of the 'steersman' associated with cybernetics is completely misleading as a framework for understanding human performance. This metaphor suggests that humans determine the behavior of an organization, when in fact, many other factors contribute to shaping the ultimate performance of any organization. The myth that humans are 'in control' leads to humans getting too much credit when things work well (the mythical hero leader) and too much blame when things don't work well (the myth that human error is the 'cause' of most accidents).

On the other hand, behaviorism is the myth that human behavior is 'controlled' by the situations (i.e., the stimulus). Despite the fact that even the transfer function for the human tracker must change as a function of the dynamics of the plant being controlled (as reflected in McRuer's Crossover Model); and despite the fact that 'context matters,' it is erroneous to think that the stimulus or the context determines or controls behavior.

So if nothing is 'in control,' how can we possibly understand or explain the performance of even simple organizations. If nothing is in control, does anything 'determine' performance? What matters? (3) The answer lies in the dynamics of coupling or of networks. Ultimately, performance reflects the demands of stability. The behaviors that persist are the behaviors that lead to network stability. Ultimately, control theory and dynamical systems theory are NOT about 'control,' but about 'stability.' Further, stability is an emergent or relational property of a network. It cannot be localized in any of the elements of the network. This is why any theory that tries to attribute 'control' or 'cause' to any component will simply be wrong!

A better metaphor for understanding the coupling within a closed loop is the prey-predator relation. Who determines the population in an ecology, the predator or the prey? This is obviously a silly question. The population depends on a delicate balance between predators and prey. If this relation is out of balance the ecology will eventually collapse.

A necessary step toward understanding performance is to abandon the illusion of 'control' (or causality) and to begin framing questions to explore the properties that contribute to the emergence of stability. It is important to shift our focus from the *elements* in an organism or organization to include the *relations* across the elements that lead to a healthy balance.

References

- (1). Jagacinski, R.J., Flach, J.M., & Gilson, R.D. (1983). A comparison of visual and kinesthetic tactual displays for compensatory tracking. *IEEE Transactions on Systems, Man, and Cybernetics*, 13(6), 1103-1112.
- (2) Jagacinski, R.J. & Flach, J.M. (2003). *Control Theory for Humans: Quantitative approaches to modeling performance*. Mahwah, NJ: Erlbaum. ISBN-13: 978-0805822939
- (3). Flach, J.M. & Voorhorst, F.A. (2020). *A Meaning Processing Approach to Cognition: What Matters?* New York: Routledge.

Chapter 36: Explaining Mechanisms vs. Organisms

Early breakthroughs with respect to explaining mechanisms or inanimate things have shaped the narratives that have dominated early science. The scientific explanations of mechanisms were based on causal narratives with a strict ordering in time such that causes always precede effects. Thus, explaining an action (an effect or a response) involved tracing back along a metaphorical string of dominos until you found a satisfying explanation (i.e., the cause of the action). This resulted in a metaphorical string of dominos that is conveniently decomposable in time. That is, it is possible to isolate pairs of dominos, and it is possible to identify local and remote causes of an action. This led to experimental approaches that isolated pairs or small sequences of dominos to test specific hypotheses, with the promise that the understanding of these isolated parts (e.g., gears) could be combined to explain the entire sequence (e.g., behavior of the clock). In fact, there was an implication that it would ultimately be possible to explain the entire universe by isolating the first, primary domino and then tracing forward along the causal chain to predict the future.

It is not too surprising that early biologists and social researchers adopted this same narrative in their attempts to develop a scientific approach to explaining the actions of organisms. This is clearly illustrated by the stimulus response (SR) explanations of behavioral psychologists. Or the decomposition of cognition into an ordered series of discrete information processing stages. It is also why analogies to mechanisms – clockworks, servomechanisms, and computers were satisfying explanations for early social scientists.


However, it is becoming increasingly evident that the causal narrative that was successful for explaining the behavior of mechanisms will not lead to satisfying explanations for the behavior of organisms. In fact, there is growing suspicion that this causal narrative also falls short with respect to understanding complex physical systems (e.g., weather systems or dripping faucets). Thus, if the goal is to explain the behavior of organisms it becomes necessary to develop an alternative to the causal narrative.

Let's consider some of the reasons why the causal narrative does not lead to a satisfying explanation for the behavior of organisms. First, the behavior of organisms is shaped by the potential consequences of actions. This is reflected in the fact that organisms pursue goals. This does not fit with the assumption that causes precede action. Thus, rather than stimuli causing actions, it could be argued that in many situations actions cause

consequences which in turn lead to sensations and perceptions. The implication is that the coupling between stimuli and actions can be bi-directional or circular.

This circular coupling can be represented by adding feedback to the typical information processing model of cognition. Although early representations of information processing with feedback tended to preserve a fixed ordering in time consistent with a causal narrative in which sensing precedes perception, which precedes decision making, which precedes action. This is inconsistent with the circular dynamic that results from adding feedback.

In the circle there is no sense in which any stage necessarily precedes any other information processing stage. For example, Gibson observed that feet were a critical component of visual perception because movement is necessary to create the visual flow fields that specify our relations to the surfaces in our ecology. The figure below shows the conventional representation of information processing and an equivalent model that implies an alternative precedence between stages. The key point is that these representations are equivalent – illustrating that within the circle there is no fixed precedence in time between the stages of processing.




A second problem with the mechanistic narrative is that there is an implication that the transfer functions of the elements (e.g., the dominos or information processing stages) are fixed. This implies that pairs of dominos or stages of processing can be isolated and

studied independently from the other dominos or stages, and that eventually the behavior of an organism could be explained by combining knowledge obtained by studying the independent elements. For example, it implies that decision making could be understood independently from sensing, perception, or action. However, research on Naturalistic Decision Making illustrates that decision making is not independent from recognition stages (e.g., selective attention and framing of the situation). And research on perception and action illustrates that a primary function of perception is to tune into the information that specifies the opportunities or constraints on action (i.e., the affordances).

This was recognized by Piaget with respect to cognitive development and by Boyd with respect to the development of expertise. Piaget described two dynamics with respect to cognitive processes – assimilation and accommodation. The assimilation is reflected in the forward loop. This reflects how internal schemas (intentions and expectations) guide our interactions with the ecology. That is, the schema directs our attention and helps us to guess what actions will lead to satisfying interactions with our ecology. For Boyd this is analogous to creative action. If our guesses are correct, then the resulting feedback will confirm or reinforce the schema as a satisfying framing for productive action.

However, sometimes we will be surprised. That is the consequences will not be consistent with our expectations. The upper loop reflects what Piaget called accommodation and what Boyd called destruction. The feedback in this loop changes our schema (e.g., it adapts or destroys our old model to accommodate the unexpected consequences). That is, the transfer function for the box labeled schema is not fixed or stationary. The transfer function changes as a result of surprising consequences reflecting the capacity of organisms to adapt or learn from experience.



An Adaptive Cognitive Dynamic

Thus, organisms are more complex than simple servomechanisms or computers with fixed transfer functions or programs specified by the engineer or programmer. They are engaged in a constant process of attunement to create a stable coupling with the ecology. In other

words, the survival of the organism depends on them attuning to their ecology. A critical insight here is that stability is a constraint on the whole system (i.e., the coupling of organisms with their environment). This reverses the narrative with respect to parts and wholes. Rather than the whole being the sum of the parts, in a circular dynamic, the parts must satisfy the stability constraints on the whole coupled system. For example, if the ecology changes, the organism must also change in a way that maintains stability. Otherwise, the organism will become extinct. In other words, survival of the organism depends on maintaining a stable coupling with its ecology – the transfer function of the organism must adapt to the stability constraints of the coupling.

McRuer's crossover model of how humans adapt to changing dynamics in a tracking task is a great illustration of how the transfer function of a human tracker must change to meet the demands of stability. If the transfer function of the plant is a simple gain (position control), then the human transfer function must have the properties of an integrator. But if the plant transfer function is an integrator (1st-order, velocity control), then the human transfer function must have the properties of a gain. The combined dynamic of human plus plant must meet the stability requirements of the entire loop (i.e., must be a low-pass filter).

You could also think about this in the context of competitive sports like tennis. To be consistently successful against different opponents it is necessary to adapt your style of play (your transfer function) in a way that counters the different strengths and weaknesses of your opponents.


Another aspect of the dynamic of circles that is inconsistent with the causal narrative is that the dynamics of circles cannot be dissected into a series of discrete instances in time (e.g., one domino or billiard ball hitting another in a series of discrete collisions). There is not a single signal in the information processing circle that then moves from one stage to the next. Rather, all stages are getting signals continuously. For example, you don't have a sensation that becomes a perception that then drives a decision which then becomes an action. Rather, the organism is continuously sensing, which is continuously being integrated into ongoing perceptions, that inform ongoing decisions, that are integrated into ongoing actions. Further, the constraint on stability is not in time, it is over time. In other words, organisms are analog.

Finally, organisms are not processing single signals in one or two loops. Rather organisms are simultaneously dealing with a multitude of signals that are percolating through a multilayered network of multiple loops. This includes sensing through multiple modalities, perceiving multi-dimensional situations, and directing and coordinating the actions of multiple effectors. Thus, organisms are not reacting to discrete particles like colliding billiard balls. Rather, they are managing fields of information and coordinating fields of

effectors. In motor control, this is reflected in the degrees of freedom problem – how to coordinate multiple muscles and joints to consistently achieve the same ends (e.g., an accurate throw to first base) under a wide range of different situations (e.g., reflecting different postures resulting from the need to first retrieve a ground ball).

The bottom line is that much of the biological and social sciences are framed within a causal narrative that was designed to explain inanimate mechanisms. However, it is becoming increasingly obvious to many researchers that this narrative does not lead to satisfying explanations for how organisms behave. Today a new narrative is being created based on the capacity of the organism to resiliently adapt to situated constraints. Rather than looking for causes, researchers are beginning to explore how behavior is shaped by constraints on perception and action. They are asking how patterns in fields of information (e.g., perceptual invariants) specify the opportunities and risks for action (e.g., affordances), relative to the successful pursuit of satisfying ends (e.g., healthy consequences, stable balance with the ecology, survival).

Chapter 37: Playing Twenty Questions with Nature



Heinz von Forster observed that

Objectivity is the delusion that observations could be made without an observer.

I think this has important implications for those of us interested in studying human experience. The idea that we can be objective observers of human experience may be a delusion. In fact, in many contexts the constraints we impose as observers may account for most of what we see. For example, early in my career I did research on human tracking and was at first amazed that classical or optimal control models could account for up to 90% of the variance. Now I look back and think - why was I surprised that when I constrained the situation so that the person had to behave like a simple servomechanism to succeed the resulting behaviors could be accounted for using models of simple control systems.

The physicist John Wheeler used the Surprise Version of the 20 Questions game to illustrate the impact of observers and the epistemological implications at the Quantum level.

In contrast to the normal Twenty Questions game where the people in the room decide on a single word to be the target, in the Surprise Version the people in the room each independently choose their own words with the constraint that once the game begins, the word that they have in mind must not contradict any previous answers to questions. This may require that the people in the room must change their word over the course of the game to be consistent with the responses of other people. Thus, the "correct" word is a moving target that is shaped, in part, by the questions asked. In fact, the final target may not have been in anyone's mind when the questioner re-enters the room. Wheeler concludes his story that


In the real world of quantum physics, no elementary phenomenon is a phenomenon until it is an observed phenomenon. In the surprise version of the game no word is a word until that word is promoted to reality by the choice of questions asked and answers given. 'Cloud' sitting there waiting to be found as we entered the room? Pure delusion!

The lesson of the Surprise Version of 20 Questions is particularly relevant to exploring human experience. Just as at the quantum level, the actions of the observer (i.e., scientist) can be a significant source of variance. This impact is typically referred to as the demand characteristics. These demand characteristics (e.g., the awareness of participants that

they are being observed) can be important factors in shaping the experience (i.e., performance) that results.

Also, it is important to realize that human actions (e.g., innovative technologies) are shaping the opportunities for experience more globally. For example, new forms of travel and communication greatly extend our abilities to explore the world and to collaborate with people in distant locations. Cell phones have dramatically changed how people coordinate social activities – “Just call me when you get out of class and I will let you know where I am, so that we can meet.” Thus, we should not think about human experience as an “object” that exists “out there” independent from the experience of the scientist or designer. Experience is not a stationary object that can be isolated and studied in a vacuum as if it is independent of an observer. Experience is constantly evolving in response to changing physical and social contexts.

Chapter 38: What Time is it?



Have you ever become so engrossed in an event that time just flew by? When an hour seemed like minutes? Have you ever been surprised to find that a companion did not share your feelings? That the event that you found to be intensely interesting, she found to be mind-numbingly boring? What to you seemed like moments, seemed to her to be an eternity?

Whose impression was right? What was the *real* time?

When people experience the same event differently – the differences are said to be *subjective biases*. The people are not being *objective*. Obviously, the *real time* is that measured by the clock. And any experiences that don't correspond with what the clock specifies must be *biased* or *illusionary*.

When our experiences of time don't correspond with the ticks on a clock; or when our experiences of distance don't correspond with the ticks on a ruler, conventional wisdom suggests that we have made a mistake, that we are in error – that our experiences are illusions.

Obviously, all the different impressions of the duration of an event by a diverse group of people can't be *real*. Only one impression can be right – and that is the impression that is consistent with the objective measurement made with a clock. An hour is an hour. And if the experience does not agree with the ticks of the clock, the experience must be wrong.

I wonder – did time exist before there were clocks?

Clocks and rulers are wonderful inventions. Together they help us to map out our world and to coordinate and synchronize our actions with others. They allow us to plan our days and map out our world – to schedule when and where to join others for work, to share a meal, or to play.


Yes, clocks and rulers are wonderful constructions, but they are constructions non-the-less. They are constructions that were designed specifically to objectify the world, to eliminate the subjective dimensions of experience from the equation.

Conventional wisdom accepts the measures of clocks and rulers as the objective reality. These are some of the tools we use to describe the objective world of *matter*. But what about peoples' experiences? If they are not consistent with the objective world of matter – are they *real*?


Conventional wisdom suggests that there are two choices: EITHER our experiences are an illusionary by-product of matter; OR they may be part of a separate, different reality that we call *mind*.

But what if conventional wisdom is wrong? What if we have mistaken the map that we constructed with our clocks and rulers as reality? What if our experiences are actually more real than the ticks on the clock?

Are there alternatives to ticks on the clock for thinking about events, durations, and synchronicity? One alternative is the frequency domain - that is characterizing patterns over time as composed of sums of sine waves. For example, insights into the performance of many systems can be gained from examining the frequency response. Many systems exhibit 1/f characteristics and stability of control systems can be inferred from Bode plots that show the frequency response. In the frequency domain, ticks of the clock are no longer meaningful. Rather, the focus is on the capacity to detect patterns over time. These patterns are indexed by sine waves of various frequencies. The fundamental question is: which patterns is the system tuned to? What is the system bandwidth? It also turns out that the stability of a dynamical system depends critically on the relation between gain (sensitivity) and processing delays (lags), that can be visualized in terms of the gain and phase margins represented in a Bode plot.




Another framework is to take events as the objective reality - and then frame time in event relative terms reflecting timing and windows of opportunity - too soon, too late, too fast, too slow, too short, too long ...



Clocks are useful for coordinating social activities - but maybe we should consider alternative ways to think about how events are experienced by biological systems.

Chapter 39: All Circles are Not the Same



With respect to the logic of block diagrams, the four diagrams above are all identical. Thus, the similarities illustrate that the four paradigms all accept that cognition involves a circular coupling among the various component processes. Despite this consensus, there are striking differences about how the four different paradigms view the dynamics of cognition. This is reflected in which components are considered to be 'driving' the forward loop and which components are considered to be consequences in the feedback loop.

Classical Information Processing - tends to start with perception as the stimulus that leads ultimately to a choice (decision) that sets the goal for action. This reflects the residuals of a stimulus-response model. In this model, the function of cognition is to generate a 'plan' that directs action.

Cognitive Science - tends to focus on the internal cognitive processes (e.g., mental models or schema) and the associated intentions as the 'central processor' where all the interesting action happens and tends to treat perception and action as peripherals. In this approach, the mind is a disembodied computational system and focus is on the brain as the central controller of human performance.


Hypothesis/Abduction - reflects the long neglected pragmatic perspective associated with Peirce, James, and Dewey. This perspective starts with action as a 'guess' or 'test of a hypothesis' - that is measured against the practical consequences that follow. If the

hypothesis leads to satisfying consequences the 'idea' or 'habit' is strengthened, otherwise an alternative hypothesis is considered. Note that while this has clear similarities with behaviorism - the key difference is that 'ideas' are strengthened - not mindless actions. These ideas reflect an ability to generalize about the world in ways that lead to improved fitness. A key distinction from traditional cognitive science is that cognition makes sense of the relation between action and consequences after the fact. Cognition is part of the feedback loop. Mercier and Speber (2007) have more recently made the case that reasoning is a post hoc phenomenon for making sense of our behavior.

Ecological Psychology - reflects the intuitions of JJ and EJ Gibson. They focused on the coupling of perception and action as the central issue for psychology - and delegate cognition and intention to the periphery. This is the central thesis of JJ's theory of *direct perception* which posits that there are structures in the perceptual array that directly specify the possibilities and consequences for action (i.e., the affordances) without the need for interpretation or mediation via mental models. This is also reflected in EJ's theory of perceptual learning in which expertise results from perceptually tuning to the relevant perceptual information, rather than in storing memories and constructing mental models.

For me, the key is to recognize the circularity of the coupling among the various components. Once this is recognized, the contradictions and points of contention among the various approaches take on a quite different flavor. Once it is recognized that all of these different approaches are talking about the same circle - then it becomes easier to appreciate the value of each perspective, to discover the common ground, to pull the signals from each of the four perspectives, and to filter out some of the debates and arguments that unproductively circle the same tree over and over.

Chapter 40: Experience: Reality or Illusion



It's the same temperature, as measured by the thermometer, but one person experiences miserable cold, while another experiences a refreshing chill. Which experience is 'true' or 'real.' Or are both experiences illusions (i.e., purely subjective)?

Most introductory psychology texts describe an experiment in which the participant leaves one hand in a bucket of ice water and another hand in a bucket of hot water for a few minutes, then simultaneously puts both hands into a bucket of water at room temperature. What happens? One hand experiences the water as being warm, while the other hand experiences the bucket as being cold.

Does this demonstrate that perception is illusory? Or does it suggest that the experience of temperature is dynamic. That it reflects change, or relations over time (as opposed to isolated events in time).

Imagine two intersecting lines on a graph one sloping down and the other sloping up. At the intersection, both point values are identical - but the slopes are different. Classically, we have tended to treat human experience as if it were a collection of isolated points in time. Thus, we have failed to consider that the slopes may be different.

If experience is dynamic, then it is essential that we consider the points as integrated components of the line. Rather than isolating behavior in time, we need to examine behavior over time (we need to consider the lines).

From the perspective of dynamics, the experiences of both people can be real. Both experiences are partial functions of the current physical temperature (as measured by a thermometer), but they are also functions of different past histories and different potential futures (e.g., intentions). Though the temperature 'points' may be the same for both, the slopes may be very different. The experiences may be situated on different trajectories.

When the different experiences are dismissed as 'subjective,' there is an implication that all experience is illusory. That experiences are groundless with respect to the objective physical situations (e.g., the objective temperature). That experiences are 'in the head.' However, if you view experience as a dynamic property over time (e.g., with both a position and a velocity or slope), then you can see that the differences may be due to the fact that both are grounded, but in different ways. In this context, both experiences can be 'real' (in the sense that they are grounded, but with respect to objectively different situations).


Note that (as with constructs such as affordance) the 'situation' is not purely physical (objective) or purely mental (subjective). Both the situation and the experience are dynamical phenomenon reflecting both where agents are coming from and where agents are going. And that trajectory is shaped by both objective physical properties (e.g., the temperature) and by mental properties (e.g., an intention).

Cognitive Science will resolve many mysteries and contradictions, if it will only connect the dots and begin to consider trajectories as the fundamental units of analysis.

Chapter 41: Do we Make or Discover Meaning?

In discussions about the nature of cognition, a central question focuses on how meaning emerges from interactions between agents and their environments. It seems clear that the 'meaning' of any object depends in part on properties of the object, in part on the observer, and in part on the situation. For example, consider the following observations from Rasmussen (1986)

The way in which the functional properties of a system are perceived by a decision maker very much depends upon the goals and intentions of the person. In general, objects in the environment in fact only exist isolated from the background in the mind of a human, and the properties they are allocated depend on the actual intentions. A stone may disappear unrecognized into the general scenery; it may be recognized as a stone, maybe even a geologic specimen; it may be considered an item suitable to scare away a threatening dog; or it may be a useful weight that prevents manuscript sheets from being carried away by the wind - all depending on the needs or interests of a human subject. Each person has his own world, depending on his immediate needs. (p. 13)



There are two subtly different ways to think about the dynamics of experience that underlies the emergence of meaning. Conventionally, constructivist approaches to cognition talk about *making meaning*. This makes a lot of sense in the context of language,

where arbitrary signs such as a sequence of marks on a page (e.g., C - A - T) are interpreted relative to prior learning about alphabets and word definitions. The suggestion is that the meaning is the result of adding prior knowledge to the arbitrary sign to make (or construct) meaning. The implication is that the symbols are meaningless until they are interpreted.

An alternative way to think about the dynamic of experience, that reflects ecological or situated perspectives on experience, is that *meaning is discovered*. This perspective makes a lot of sense in terms of perceptual-motor skills. For example, we discover affordances like graspable and reachable by interacting with the objects in the environment. The underlying relations that determine whether an object will fit comfortably in the hand are not arbitrary (though the affordances of a specific object like a basketball may vary from individual to individual as a function of hand sizes). Affordances reflect *meaning-full* properties of the ecology - that exist independent from perception or interpretation. The intention will not be realized if the affordance is not detected, but the affordance exists and can be specified objectively, whether it is ever realized in action. Further the meaning can be mis-perceived but will be corrected through the feedback that results from acting on the misperception.

The framework of *meaning making* makes sense if you think about the stimuli of experience as punctate instances in time (e.g., isolated frames in a movie reel). In this case, experiencing a melody requires that the significance of a particular note be *constructed by* retrieving the prior notes from memory and mentally adding them together to re-construct the melody.

In contrast, the framework of *meaning discovery* suggests that perceptions are not punctate, but that they are extended over time so that the pattern of notes is experienced as a whole (as a chunk). This extension may go beyond the notes heard to include prior experiences with a particular memory that allow prediction or anticipation of the entire melody. The metaphor does not have to invoke memory in terms of adding the prior notes. Rather, the metaphor is one of attuning or resonating to a pattern - and recognizing a melody.


Note that the meaning discovery framework does suggest the existence of mental structures (schema or frames) - but these structures function more like filters - that resonate to some properties or patterns, as a function of prior experience. In this framework, the function of experience or learning is not about storing past instances (that can be added to new instances to construct meaning), rather it is about tuning attention to those properties of experience that have functional significance (e.g., tuning the weights in a neural net).

Back to the processing of (C-A-T). These symbols may be *arbitrary* in that there is no obvious physical or analogical relation to the animal that they represent. But they are NOT arbitrary in a cultural sense. If we assume that the meaning of C-A-T is created by a culture - not by a mind. Then the meaning discovery framework could be sensibly applied to language as well as to perceptual motor skills. In this sense, learning language is not about creating meaning from arbitrary signs - but about discovering the cultural significance of the signs (in the same way that discovering affordances is about discovering the significant action properties of an object).

The danger of the constructivist framework where minds make meaning is the implication that everything is meaningless until it comes in contact with a mind. There is a subtle implication that we live in a meaningless world. I can't accept that implication - and thus prefer to think of the dynamic of learning and experience as one of *discovering meaning*. There is a subjective dimension to meaning, but I can't accept that meaning is purely subjective.

About the Author

Despite my best judgement and due to some serendipitous events (including a presentation given by Herbert Simon at University of Dayton) I ended up in graduate school at The Ohio State University working with Rich Jagacinski in 1978. I distinctly remember sitting in front of an analog computer and pretending to understand as Rich explained to me that linking two integrators in a closed circuit would produce a sinusoidal signal. The happy accident of a research assistantship with Rich set me on a journey that involved taking undergraduate math and engineering courses along with the normal graduate load in Psychology - so I could learn the analytical tools of control theory and apply them to modeling human motor skills in Rich's lab. Two other important influences at OSU were Dean Owen and Rik Warren who were collaborating to explore the implications of Gibson's intuitions about perception for understanding piloting using innovations in the technology of graphical displays and flight simulators. Six years and two kids later, with support and encouragement from my wife I finally completed the Ph.D. in Experimental Psychology.



The unusual blend of social science and engineering led to the next happy accident. In 1984 I began my academic career at the University of Illinois where I had joint appointments in three departments [Mechanical & Industrial Engineering (50%); Psychology (25%); and Institute of Aviation (25%)]. Although people warned me of the difficulties associated with split appointments, it seemed like an ideal opportunity to apply my unusual prior experiences to explore cognition in the context of evolving information technologies and complex work domains. While at Illinois I had an opportunity to work with some outstanding people, including Chris Wickens, Henry Taylor, Art Kramer, Neville Moray, Penny Sanderson, Peter Kugler, and Gavan Lintern.

Chris Wickens was particularly influential in building the aviation research laboratory and getting the resources that allowed me to continue to explore cognitive skills associated

with piloting and how the information in optical flow fields can be used as feedback for controlling locomotion. Peter Kugler helped me to better appreciate the challenges associated with complex systems and nonlinear dynamics. Also, while at Illinois I was introduced to some of the people who helped to frame the field of Cognitive Systems Engineering, including Jens Rasmussen, Erik Hollnagel, Dave Woods, James Reason, Morten Lind, Henk Stassen, and Sydney Dekker. Perhaps, the work with Kim Vicente and Jens Rasmussen was most influential on my future work. Jen's vision from above and Kim's pushing from below helped me to integrate the intuitions from control theory with an ecological perspective on skilled performance into a framework for thinking about how to engineer more effective joint cognitive systems.

After six years, however, I found that I was unable to construct a coherent narrative about joint cognitive systems that would satisfy the expectations of the University. I was not offered tenure at Illinois. In essence, I was banished to the ["Pointless Forest."](#) Now with three kids and thinking that my academic career might be over, my wife and I began searching for a new home.



Fortunately, the next happy accident led to our new home in Dayton, OH. In 1990 I started a new position in the Psychology Department at Wright State University with joint funding from the Air Force Research Laboratories at Wright-Patterson AFB. I was very fortunate to have the opportunity to get in on the ground floor of a new Ph.D. program in Human Factors and Industrial/Organizational Psychology and to help shape a program that better aligned with my own vision for an applied Cognitive Psychology.

I particularly valued my collaborations with Rik Warren (at AFRL) to frame an approach to "active psychophysics" that integrated concepts from Ecological Psychology (structural

invariants in optical flow fields) with principles of control theory; and also, my collaborations with Kevin Bennett in relation to developing pedagogies for teaching Cognitive Systems Engineering (CSE) and Ecological Interface Design (EID). My students and I also benefited greatly from being in proximity and learning from Gary Klein as he explored naturalistic decisions. His stories about how experts cope with complex work were an important source of inspiration.

In many ways, my situation in Dayton turned out to be the realization of what I was hoping to achieve at Illinois. Like Oblivion, I was pleasantly surprised to discover that everything in the Pointless Forest had a point!

Over the years I have benefited greatly from numerous international collaborations. Early collaborations with Jens Rasmussen and his research group at Riso helped set the foundation for much of my thinking. Collaborations with the Industrial Design Engineering group at TU Delft (Gerda Smets, Kees Overbeeke, PJ Stappers, and Fred Voorhorst) gave me my first real exposure to design thinking and research through design. Additionally, collaborations with the Human Systems Group in Aeronautical Engineering at TU Delft (Max Mulder, Rene van Paassen, Matthijs Amelink, Clark Borst) have given me an opportunity to continue to hone my understanding of control theory and to explore the application of EID concepts in the aviation domain.

In 2004, another series of accidents led to me becoming the Chair of the Department of Psychology at Wright State. Despite my inexperience and muddling style of leadership the department continued to grow and thrive. I was extremely lucky to have an excellent Dean (Michele Wheatly) - she and I were able to collaborate with the Engineering College (Forouzan Golshani and John Gallagher) to win several large NSF awards in the area of Technology-based Learning with Disabilities. The experiences of managing a faculty - all of whom were smarter than me - piqued my interest in the challenges associated with managing a complex organization. Thus, I became increasingly curious about the dynamics of polycentric control.

In 2012 I happily handed off the responsibilities of administrivia and returned to my role of professor. Then in 2017, another accident allowed me to take advantage of an early retirement Incentive offered by the university to take a deeper journey into the Pointless Forest. I left the Ivory Tower to join a small startup design company (Mile Two) that focused on improving the fit between people and technology. I soon discovered that building software to support cognitive systems is a lot harder than armchair quarterbacking from behind the walls of the Ivory Tower. It was exciting to be part of a growing startup - that grew from 12 to 120 people in the span of 6 years. However, I failed in making an adequate business case for CSE to customers or colleagues. It seems that it is not feasible to do a

reasonable work domain analysis in the context of a one-year contract. Customers were satisfied with modest, incremental improvements, and were not willing to make the investment to explore deeper into work domains. So, in 2023 I was compelled to move even deeper into the Pointless Forest as an independent consultant. I teamed with my brother, Joe, to form Perspicacity LLC. Our goal was not so much to create a business - as to have a vehicle that would allow us to continue to collaborate with smart people to explore interesting problems.

At the end of the day, despite many years of formal education and extensive reading about human performance and design, my understanding of human performance stems most directly from my failures at sports, ineptitude with music, and general uneasiness with modern technologies and organizations. I came to psychology seeking reasons why things that seemed to come so easily to others were difficult for me. This has led me to explore general issues of coordination and control in cognitive systems across a wide range of domains - including aviation, process control, military command & control, and healthcare.

When I look back on my journey so far - I feel more like a leaf being carried forward by a turbulent current, than like the captain guiding a ship. Perhaps, this is why, in my search to understand cognitive systems, I always look first to the ecological constraints, before considering internal mental explanations of behavior and outcomes. For my life - the direction has typically been shaped more by happy accidents (situation constraints), than by deliberate choices. It has taken almost all my energy to just keep the boat upright, and thus I have had only a very limited role in determining the ultimate destination of the journey.

One of my greatest pleasures along the journey has been participating on multidisciplinary teams to explore how information technologies can be used to help people manage the complexities of modern life. I particularly welcome opportunities to watch and talk with domain experts about their work and the tricks of the trade that allow them to manage complexities that exceed the capabilities of mere humans. I have been lucky to have had many interesting adventures in other peoples' domains within the Pointless Forest. However, the original problems remain unsolved - I remain a clumsy athlete, an inept musician, and am always the last to learn to trust and use new technologies.

I recently was surprised to learned that ScholarGPS ranks me as the 2nd highest contributor to the specialty Human-Machine Systems. Also, I am ranked in the top 1% of all scholars, and in the fields of social science, and psychology:

<https://scholargps.com/scholars/41852767616510/john-m-flach>

Google Scholar : [John Flach](#)

Links to some recorded interviews and podcasts:

Here is a link to a podcast with Nippin Anand in which we talk about his book "Are we learning from accidents?" and a wide range of topics related to human experience.

[Nippin Interview](#)

To learn more, here is a recent interview I did for the Naturalistic Decision Making Podcast:

<https://naturalisticdecisionmaking.org/2021/02/03/john-flach/>

Here is a more recent podcast that Alex Kirlik and I did for the Human Factors Interest Group at the University of Toronto.

<https://anchor.fm/humanfactorsinterestgroup/episodes/HFIG-Talks----HF-Past--Present--and-Future-with-Dr--John-Flach--Dr--Alex-Kirlik-e18bipv>

Here is a recent podcast on Ed Brenegar's Eddy Channel:

<https://www.youtube.com/watch?v=R2G08cP1MLU>

Here is a recent conversation on the NoWayOut podcast with Brian "Ponch" Rivera and Mark "Moose" McGrath:

<https://nowayout.buzzsprout.com/2109174/13198366>

Here is a recent interview with Laksh Raghavan:

Part 1: https://www.youtube.com/watch?v=BL_0cxnwZsY

Part 2: <https://www.youtube.com/watch?v=hJ11Xg6IID0>