Discovering How Distributed Cognitive Systems Work

DAVID WOODS

Cognitive Systems Engineering Laboratory
Institute for Ergonomics
The Ohio State University

> *Modern biology has demonstrated the enormous power of functional analysis coupled with naturalistic observation.*
>
> *Ulric Neisser, 1991, p. 35*

Abstract

In cognitive work there is a mutual adaptation of agents' strategies, affordances of artefacts, and demands of the work setting. Regularities in cognitive work concern dynamic interactions across of these sets of factors. Since cognitive work systems are not decomposable into independent basic elements, different empirical tactics are necessary, though core values of observation, discovery, and establishing warrant remain fundamental guides.

Functional analysis is a mode of research that can cope with the unique difficulties of studying and designing cognitive work in context. Functional analysis is a process that coordinates multiple techniques in order to unpack complex wholes to find the structure and function of the parts within the whole. Central to the empirical techniques orchestrated is observation. Three families of techniques emerge which vary in how they shape the conditions of observation: natural history or in situ observation, staged world or observation of performance in simulated situations as models of what is important in situ, and spartan labs settings where observation occurs in experimenter created artificial tasks. The chapter discusses how to coordinate these techniques in a discovery process that reveals the mutual adaptation of strategies, affordances, and demands and predicts how these dynamic processes will play out in response to change.

Studying Cognitive Work in Context

> *The study of cognitive systems in context is a process of discovering how the behavior and strategies of practitioners are adapted to the various purposes and constraints of the field of activity.*
>
> *David Woods, 1997*

This single sentence both describes the enterprise of understanding and designing cognitive work and points to the difficulties and challenges embedded in that enterprise. The target phenomena of interest for distributed cognition at work are reflected in the basic founding slogans of Cognitive Systems Engineering -- adaptations directed at coping with complexity, how things make us smart or dumb, and how to make automated and intelligent systems team players. Patterns on these themes in cognitive work exist only at the intersections of people, technology and work.

In studying cognitive work in context we are out to learn how the more or less visible activities of practitioners are parts of larger processes of collaboration and coordination, how they are shaped by the artefacts and in turn shape how those artefacts function in the workplace, and how their activities are adapted to the multiple goals and constraints of the organizational context and the work domain (Rasmussen et al., 1994). These factors of complex artefacts, dynamic worlds, cognitive work, coordinated activity, and organizational dynamics do not come to us pristine, isolated, one level at a time. Rather they come in association with each other, embodied in the particular, cloaked by some observer's vantage point. In other words, cognitive work in context comes in a "wrapped package" as a complex conglomerate of inter-dependent variables (Woods et al., 2002).

Agent-Environment Mutuality

Understanding the deeper patterns in this type of complex conglomerate is difficult because the processes of cognitive work are dynamic and mutual adapted–a case of agent-environment mutuality (Flach et al., 1995; Woods, 1988). Agent-environment mutuality, after von Uexkull's Umwelt (1934) and Gibson (1979), means that agents' activities are understandable only in relation to the properties of the environment within which they function and a work environment is understood in terms of what it demands and affords to potential actors in that world. Agents and environment are mutually adapted establishing a dynamic equilibrium when the environment is stable (i.e., when sources of variability are not changing). When the environment is changing (i.e., when sources of variability are changing) adaptation ensues until a new equilibrium is reached.

In the case of cognitive work in context, this mutual adaptation is a three way interplay of the strategies of agents (human and machine), the affordances of artefacts (including how they represent the processes at work), and the demands of the field of practice (both those inherent in the processes and those imposed by organizational context) as captured in the cognitive systems triad (Woods, 1988; Woods et al., 2002).

Partial decomposability in this triad means that methods can be used to build up a model of how any one is adapted to the conjunction of the other two. In other words, the methodological challenges build from the fact that two of the three factors will be confounded in any technique to tease open the interplay. For example, since the processes in the triad are mutual adapted, developing skill at cognitive work hides or obscures the underlying adaptive web. As the fluency law of cognitive work states (Woods, 2002): “well adapted cognitive work occurs with a facility that belies the difficulty of the demands resolved and the dilemmas balanced.”

However, ‘Spartan’ lab experiments would simplify such conglomerates into more manageable units for experimental manipulation. Simplifying such dynamic processes into a series of static snapshots or treating this highly interconnected set of factors as separable represents a retreat from complexity that values the means of experimental tractability over the end of adding to our understanding of the phenomena of interest (Hoffman and Deffenbacher, 1993). Such efforts fall prey to oversimplification biases which eliminate the very phenomena of interest in the process of simplification (Feltovich et al., 1997). Furthermore, deciding what to leave out or put in, whether in an experimenter created micro-world or in a scaled world simulation, is in itself a potent model of what matters in situ.

How then can we advance our understanding of only partially decomposable, mutually adapted complexes at work? Understanding what an environment affords to agents (given their goals) and how agents’ behaviour is adapted to the characteristics of the environment and how this linkage changes is *functional analysis*. Hence, the study of cognitive work in context develops a functional model that captures how the behaviour and strategies of practitioners are adapted to the various purposes and constraints of the field of activity. This chapter grounds the functional analysis of cognitive work as a process of *discovery* built on an empirical foundation of abstracting patterns from multiple sources of observation and directed toward sparking innovation of new ways to use technological possibilities to enhance cognitive work (Potter et al., 2000; Woods and Christoffersen, 2002; Woods et al., 2002).

The Functionalist Perspective

The factors of complex artefacts, dynamic worlds, cognitive work, coordinated activity, and organizational dynamics come to us as an adaptive web consisting a complex conglomerate of interdependent and dynamic processes. Patterns in cognitive work do not exist separate from particular intersections of specific people, technologies and work. Hence, field settings are more than unique domains or exotic fields of practice, but function in parallel as *natural laboratories* where these generic patterns emerge. What approaches would help us discover the essential patterns underlying diversity of people, technology and work? What techniques would help us begin to unpack or partially decompose complex wholes like these into meaningful parts and their interactions? How do we use these patterns to guide change and innovate new possibilities in concert with other stakeholders?

For studying cognitive work in natural laboratories, perhaps astronomy or evolutionary biology, not laboratory physics, provide models for progress. Independent of controls to isolate individual cause-effect relations, in astronomy and evolutionary biology investigators are able to select and shape conditions of observation and use models to form expectations about what one might see in order discover and understand the interplay of processes at work in the phenomena of interest. These are the areas that led Neisser (1991, p.35) to note, "… the enormous power of functional analysis coupled with naturalistic observation." The model for doing functional analysis of cognitive work emerged at the beginning of Cognitive Systems Engineering in several parallel lines of work (see Rasmussen, 1986; Hutchins, 1995a; Vicente, 1998 for later descriptions).

Functional Analysis as a Mode of Research

Functional analysis is a process that coordinates multiple techniques in order to unpack complex wholes to find the structure and function of the parts within the whole. Functional analysis builds a model that explains the structure in relation to the purposes, where a function specifies how structures are adapted to purposes (more formally, a functional analysis describes an object and its function as how behaviours **B** $_{1, ..., p}$ of item *i* $_{1, ..., m}$ function in system **S** in environment **E** relative to purposes **P** $_{1, ... n}$ to do **C**, according to Theory **T** and its supporting data.). Functional Design is the process of shaping structure to function.

Functional analysis is a mode for research critical to understanding systems that are only partially decomposable, dynamic, and adaptive. It is both empirical -- based on a variety of observations collected in different ways; and theoretical -- a model of how something functions to achieve a goal.

1) A functional model is context bound, but generic.
2) A functional model is tentative; it can be overturned or the patterns re-interpreted.
3) A functional model is goal oriented; purposes serve as a frame of reference or point of view; the functions of pieces change with purpose (i.e., purposes provide varying perspectives).
4) Functional analyses are inherently multi-level; phenomena at one level of analysis are situated in a higher level context and are present because of what they contribute in that context.
5) Functional analyses involve the mutual interaction of object or agent and environment (ecological).
6) Functional models do not imply correct, or ideal or fixed (i.e., that it has to be that way). Rather they are candidate explanations which attempt to meet one kind of sufficiency criteria – ability to carry out the functions in question.
7) The Coupling Principle: the more intertwined the relationships between structure and function, the more complex the system operationally (and the less the system is decomposable into almost independent parts).
8) The essence of good representation design is making function apparent (linkage of apparent structure and underlying function); this means functional modelling is linked to design–functional analysis results in models of how the system in question will adapt to change and in hypotheses about what changes would prove useful.

Functional analysis generates models based on patterns abstracted from observation and which can be subjected to a critical examination through argument based on re-interpretation of what would account for patterns and on abstracting additional patterns. Functional models motivate and guide empirical confrontations with the field of practice.

The functional model serves to abstract the particulars of a setting in a way which supports bounded generalities. Functions serve as a tentative condensation of what has been learned (subject to critical examination and revision) based on a set of empirical observations and investigations.

How does one take up functional analysis as a research program? The starting point is orchestrating a variety of different ways to observe cognition and collaboration at work.

Shaping the Conditions of Observation

Observation

The hallmark of any science is *empirical confrontation* – one observes to note patterns; one generates ideas to account for the variability on the surface; and one subjects these ideas to empirical jeopardy. Different methodological traditions are adapted to constraints on these processes related to the phenomena of interest and to be explained. All approaches have the potential to generate insight and all are subject to fundamental sources of uncertainty that define basic tradeoffs in research. Different traditions in method balance these tradeoffs differently, but the ultimate test is one of productivity – do the methods help generate patterns and ideas about what are the essential variables and processes that account for the breadth and diversity on the surface.

When we see that field settings can function in parallel as natural laboratories, one relevant question becomes: how can one shape the conditions of observation in naturally occurring laboratories? When one considers research as orchestrating different ways to *shape* the conditions of observation, a variety of issues arise about coordinating more field- and more lab-oriented techniques.

Different field studies can become part of a larger series of converging observations and operations that overcomes or mitigates the limits of individual studies. This is based on using different techniques to examine a single field of practice and from examining the same theme across different natural laboratories.

Three Families of Methods that Vary in Shaping Conditions of Observation

How can one shape the conditions of observation in naturally occurring laboratories? This question serves as a means to distinguish different classes of research methods each based on core values, but pursuing those basic values in different ways adapted to the role of each in a full program of discovery and verification.

When one sees research methods as variations in shaping the conditions of observation, three classes of research methods appear along this base dimension (Figure 1):

- Natural History methods -- *in situ*,
- Experiments-in-the field or Field Experiments -- *Staged or scaled Worlds*,
- Spartan Lab experiments -- *experimenter created artificial tasks*.

These three classes are distinguishable by determining the 'laboratory' used. Natural History techniques are based on a diverse collection of observations in situ (Hutchins, 1995a). Experiments in the field begin with Scaled World simulations that capture or "stage" what is believed to be the critical, deeper aspects of the situations of interest from the field (e.g., De Keyser and Samercay, 1998). Spartan Lab techniques focus in on a few variables of interest and their interactions in experimenter created (and therefore artificial) situations.

These three families differ in many ways, but all three begin with observation. The most basic the role of scientist is as trained observer -- one who sees with a fresh view, who wonders child-like -- why is it this way, how does it work -- making fresh connections and generating new insights. Similarly, all three classes pursue basic experimental values (Salomon, 1991) though in different ways adapted to the uncertainties and purposes of each family of methods (Hoffman and Woods, 2000):

- means to establish warrant -- how do you know? why should one accept your findings, observations, conclusions, and interpretations?
- standards of quality -- how do you know what is a competent from an incompetent investigation given the uncertainties of research and modelling?
- means to facilitate generalizability -- what is this a case of? what other cases/situations does this apply to?

Also basic to all three families is the *target-test mapping*–the mapping between the target situation to be understood and the test situation where observations are made. In Natural History techniques the target is examined relatively directly, though subject to a variety of choices about who observes what, when, and from where.

What is particularly interesting is that all of the constraints and issues that apply to Natural History methods apply to Staged World methods, plus more. And all of the constraints and issues that apply to Staged World methods flow to Spartan Lab studies, plus more. Ironically, moving from in situ to scaled worlds to artificial worlds brings sources of additional leverage, but at the cost of new dependencies related to authenticity –the target-test mapping.

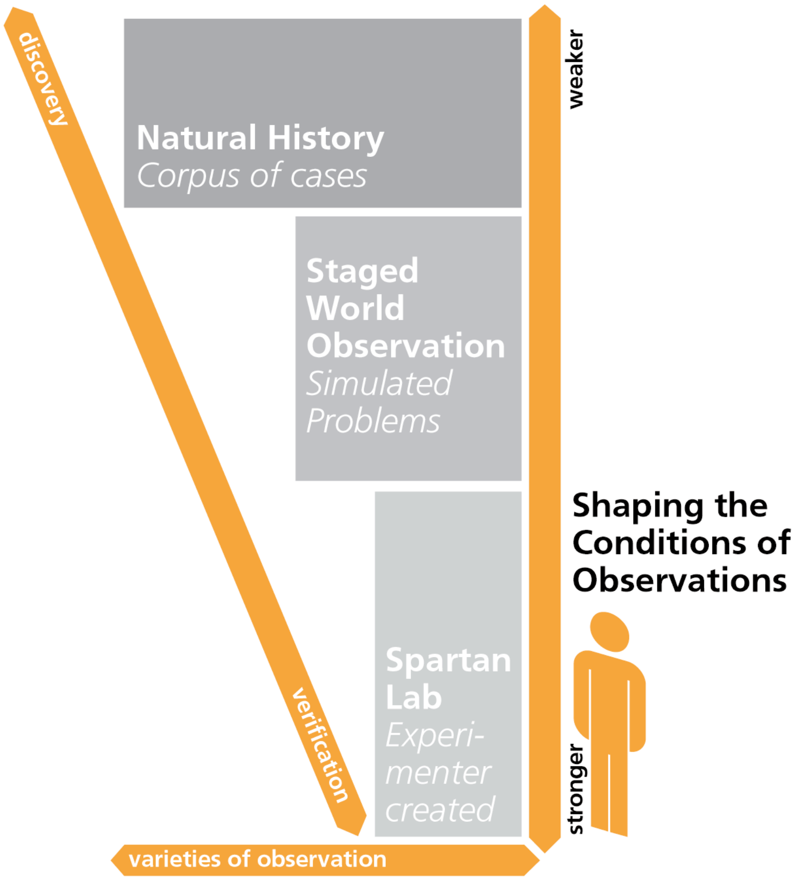

Figure 1. Three Families of Variants on Shaping the Conditions of Observation

Natural History Techniques

> *"no search has been made to collect a store of particular observations sufficient either in number, or in kind, or in certainty, to inform the understanding ..."*
>
> *from Bacon 1620*

Natural History techniques are based on a diverse collection of observations, "which agree in the same nature, though in substances the most unlike" as Bacon put it at the dawn of science.

Natural History begins with the analysis of the structure or process of each case observed and then the results are compared and contrasted across other analysed cases (requiring observations across various conditions). This is well suited to situations where variety across multiple factors is endemic in the phenomenon to be understood.

There are no subjects when one works with people who have a stake in the behaviour they engage in. Similarly, investigators become participant observers; as a result, the question of how the process of observation changes the processes observed is intensified (De Keyser, 1992).

In natural history methods a critical constraint is that observers have to wait for what happens, however, what happens is authentic by definition, though only a sample. Shaping the conditions of observation is directed at increasing the chances or frequency of encounter of the themes of interest. For example, one might decide to observe a crossroads setting (e.g., the intensive care unity in medicine) where many diverse activities, tempos, and agents interact at different times, then using patterns noted to choose a tighter focus for future rounds of observation.

In shaping the conditions of observation in this family, one considers:

- How to be prepared to observe (doing one's homework about how the domain works)?
- The role of concepts and patterns in being prepared to be surprised.
- Who observes?
- What is the role of guides, informants and others who mediate access to the field of activity?
- From where to observe (vantage point)? For example, points of naturally occurring disruption and change may be useful places to observe.

How does past knowledge (patterns, concepts, or models) help one form expectations about what one might see in order discover and understand more about the processes at work? Or *how does one prepare to be surprised?* This is the critical driving question for designing discovery oriented research programs. Observation in situ begins a process of discovering what patterns and themes in distributed cognitive work play out in this setting. What kind of natural laboratory is this or could this be? This presupposes knowledge of patterns and themes in distributed cognitive work which could be recognized as tentatively relevant for structuring these settings (Woods and Christoffersen, 2002).

How to describe a culture from the insider's perspective without getting lost in the natives' views? How is one authentically connected to practice yet able to step out of the flow to reflect on practice (a reflective, but practitioner-centred account) especially when the nature of practice is under vectors of change? For example, one might observe at those points and times practitioners gather in reflective forums or one might re-shape a naturally event to become a reflective forum. Coming to think like a native to some extent, to understand the distinctions they make, is required. But, the observer confronts the problem of referential transparency, which refers to the difficulty in seeing what one sees with, both for the practitioners in their context and for observers themselves given their own pressures and goals. After all and fundamentally, people may not have access to the critical factors that influence behaviour, but people never fail to have some model of how they work (and of the people, devices and processes they interact with).

It is particularly important to keep in mind that in cognitive work the apparent variables are in fact complex processes in their own right, rather than mere primitive elements. Discovering/identifying the processes that drive performance and adaptation is always a critical part of research that becomes explicit in functional analysis. This relates to the problem of discovering what may be the effective stimuli that control behaviour out of the infinite possible ways one could represent the stimulus world. Natural History methods are needed to identify initial places to focus on in this quest (Bartlett, 1932).

Staged World Studies

What has been generally labelled field experiments or experiments in the field lies in-between Natural History techniques and Spartan Lab techniques (Woods, 1993). This family is often viewed as a weak or limited version of real experiments, degraded and limited by the context and (by implication) the skill of the investigator. Rather, this area on the dimension of shaping conditions of observation is a unique, substantive family of methods different from lab experiments and natural history though overlapping with each in several ways (see as examples, Roth et al., 1987; Layton et al., 1994; Nyssen and De Keyser, 1998, Johnson et al., 2001; Dominguez et al., in press).

The key in this family of methods is that the investigators stage situations of interest through simulations of some type. The concrete situation staged (the test situation where observation occurs) is an instantiation of the generic situations of interest (the target situation to be understood). The mapping between the target situation to be understood and the test situation where observations are made must be explicit as this represents a hypothesis about what is important to the phenomena of interest. Classically this is called the problem of the effective stimulus. Thus, the fundamental attribute to Staged World studies is the investigators' ability to design the scenarios the participants face. Problem sampling is the focus of this family of techniques.

Staging situations introduces a source of control relative to Natural History methods as it allows for repeated observations. Investigators can seize on this power to design scenarios and contrast performance across inter-related sets of scenarios--if something is known about *what makes situations difficult* and about the problem space of that domain. Though it must be remembered that while one may try to stage the situation the same each time, each problem solving episode is partially unique, as some variability in how each episode is handled always occurs.

Second, in staging situations investigators can introduce disturbances – or *probes* -- and observe the reaction of the system tracing the process as the system moves toward a resolution or new equilibrium. Probes can be fashioned as complicating factors in scenarios, points of change, and the introduction and absorption of new artefacts. Probes can be built into the problems posed for practitioners, or they can be changes in the artefacts or resources you provide/allow access to, or changes in the distributed system (e.g., adding a new team member; linking a collaborator at a distance through a technology portal).

When investigators probe the nature of practice by changing the kinds of artefacts practitioners use to carry out their activities, the intent is not to evaluate the artefact as partially refined final product, but rather to use them as tools for discovery, i.e., as wedges to break into the cognitive system triad. Artefact-based methods have great potential because introducing new technology into fields of ongoing activity always serves as a kind of natural experimental intervention (Flores et al., 1988; Carroll and Campbell, 1988). Prototypes of new systems are not simply partially refined final products; they also can function as a kind of experimental probe to better understand the nature of practice and to help discover what would be useful (Woods, 1998). Artefact-based methods depend on a functional model of the artefacts as affordance for cognitive work or how the change in artefacts changes the cognitive work necessary to meet demands(see Watts-Perotti and Woods, 1999, for one example). In this way functional models are a means for representing the research base in a way that is useful to frame or focus new studies and to guide provisional action

Critically, in the final analysis, staged world studies are still fundamentally observational and discovery oriented -– letting the world tell us how it works rather than playing 20 questions with nature in an artificial lab. Building on top of the fundamental observational nature of this family, investigators can engineer contrasts by designing sets of scenarios, changes in artefacts, and changes in collaborative links, again, if they can draw on knowledge about the relevant problem space, affordances, and functions of distributed cognitive work. But it is important not to be confused into now thinking these contrasts convert a Staged World observational study into a Spartan experiment. Each of the things contrasted (problem, affordance, collaborative function) is still a complex conglomerate in its own right with many degrees of freedom. Avoid the error of mistaking complex factors as primitive elements, which require no further analysis.

Process Tracing in Staged World Observation

Because the investigator designs the scenarios and probes the participants confront, the investigator now has the ability to trace the process by which practitioners handle these situations. Process tracing methods follow the steps specified by Hollnagel et al. (1981) as sketched in Figure 1 and summarized in Table 1 (see Roth et al., 1987 for a sample analysis in full):

- The starting point is analysis of the process for handling the situation in its specific context.
- Second, shifting to analysis of process abstracted using concepts about the aspects of cognitive work instantiated in the concrete case (e.g., Roth et al., 1987; Johnson et al., 2001).

- Third, various levels of cross-protocol contrasts (across probes, scenarios, domains) to abstract patterns and build possible functional accounts (e.g., Woods and Patterson, 2000).

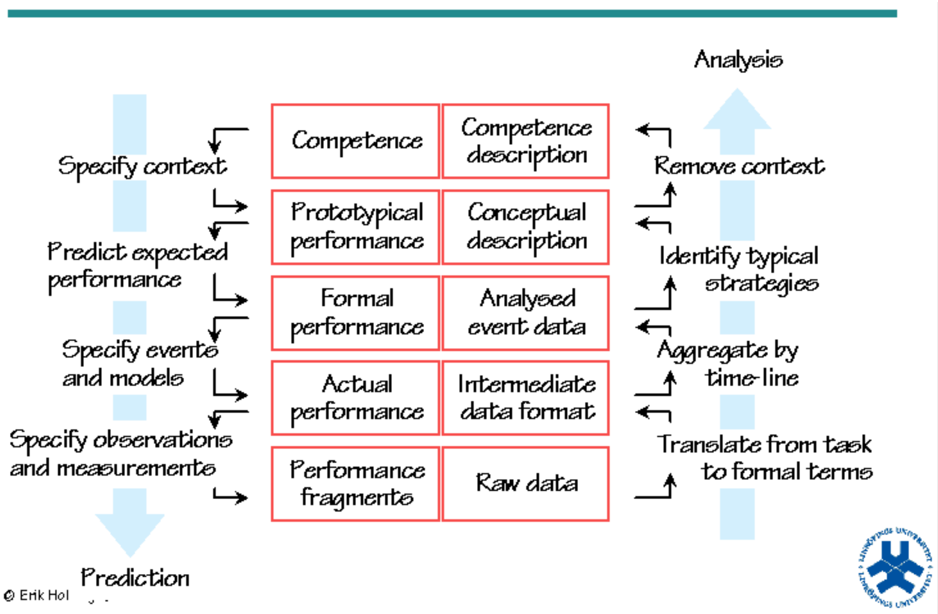

The ability to probe practice through Staged World methods is a powerful tool to observe how change in demands, pressures, complexities, and capabilities triggers vectors of adaptation and resettling into new equilibria. Staged World techniques build on Natural History methods which can be used to begin to identify what are the important aspects to preserve in the scaled world (De Keyser and Samercay, 1998; or see the linkage across Sarter and Woods, 1995; 1997; 2000).

Staging situations of interest also occurs in Spartan Lab methods though these attributes may be overshadowed by new sources of control. Simulations vary from replica worlds, to scaled worlds, to the micro worlds experimenters create for artificial labs (sometimes these artificial labs are given cover stories as a pretence to authenticity). In any case, what is staged, whether scaled or artificial, is a potent model of the target phenomena based on insight and evidence, tentative, and subject to contention. While the target is examined relatively directly in natural history techniques, the target-test mapping is still subject to a variety of choices about who observes what when and from where (an extreme form of a 'replica').

Table 1. A Checklist for Designing Staged World Field Studies (see Hollnagel et al., 1981; Woods 1993).

STEPS	PROBLEMATIC ITEMS
1. What are the cognitive system issues in question?	Pitfalls: superficial labels
2. Explicitly map target and test situation relationship	Problem space, scenario design, stimulus sampling
	Pitfalls: psychologists fallacy, cover stories
3. Shape the test situation to address the general cognitive system issues	Scenario design
	Artefacts as tool for discovery
	Procedure to externalise cognitive activities
4. Map canonical behaviour; how does one prepare to be surprised?	
5. Collect the data (update 4 as needed)	
6. Collate multiple raw data sources to construct base protocol	Analyse each participant/team's process
	Pitfalls: get lost in the raw data
7. Use concepts from 1 to build formal protocol	Pitfalls: get lost in the language of the domain; excessive micro-encodings; hindsight bias
	Distinguishing interpretations from observations
8. Cross-protocol contrasts	Aggregate across teams and problems
	Role of cognitive simulation, neutral observers
9. Integrate with other converging studies to build a functional model	

To summarize, staged world studies are fundamentally observational. Their power comes from the design of the situation to focus the conditions of observation on places where the patterns of interest are more likely to play out. This means, not just any 'interesting' situation, but working through the target–test mapping explicitly and in detail. For example, typically this is done by considering what makes the problem difficult (demands) for any system of cognitive agents. In designing the situation, probes, and the conditions of observation, the goal is to uncover what has been encoded in practice. The general problem relates to any adapted system, i.e., the fluency law.

1. Focusing observation by designing and staging situations based on the abstract patterns of interest.
2. Introducing disruptions and observing how the distributed cognitive system responds (probe techniques).
3. Contrasting conditions relative to the artefacts, relative to the levels of experience of practitioners, relative to the difficulty of the demands of the situation, or on other dimensions.

Converging Operations in Discovery

> *Doing a cognitive task analysis is like trying to learn the secret of a magic trick: once you know the trick and know what to look for, you see things that you did not notice when you did not know exactly what to look for.*
>
> *Woods' (1997) paraphrase of Orville Wright's comment about discovering the secret of flight*

> *Fluency law: Well adapted cognitive work occurs with a facility that belies the difficulty of the demands resolved and the dilemmas balanced.*
>
> *Woods, 2002, Laws That Govern Cognitive Work*

Different studies within and across these families become part of a larger series of converging operations that overcome or mitigate the limits of individual studies. This is based both on using different techniques to examine a single field of practice and from examining the same theme across different natural laboratories. As a result, we speak of studying cognitive work in context as a bootstrap within a field of practice and set of converging operations across natural laboratories as one hones in on a cognitive system theme (Potter et al., 2000).

Note the use of the words weak and strong to describe the vertical dimension in Figure 1. Weak shaping of the conditions of observation is both a strength and a weakness; strong shaping of the conditions of observation is both a weakness and a strength, though each in complementary aspects. Hence the bootstrap or orchestration of varieties of observation that are diverse in how they shape (and therefore distort) the processes we wish to observe and understand.

As one orchestrates a series of investigations that vary in how they shape the conditions of observation there are a variety of pitfalls that must be evaded.

The Psychologist's Fallacy in Research on Cognitive Work

> *“Because technical work ... is so poorly understood, policy makers routinely fall back on stereotypes or images of work … in order to make sense of technical work. The potential cost of misunderstanding technical work” is the risk of setting policies whose actual effects are “not only unintended but sometimes so skewed that they exacerbate the problems they seek to resolve”. “Efforts to reduce ‘error’ misfire when they are predicated on a fundamental misunderstanding of the primary sources of failures in the field of practice [systemic vulnerabilities] and on* ***misconceptions of what practitioners actually do***”
>
> (Barley and Orr, 1997, p. 18, emphasis added).

It is easy to commit what William James (1890) called over one hundred years ago the Psychologist’s Fallacy. Updated to today, this fallacy occurs when well–intentioned observers think that their distant view of the workplace captures the actual experience of those who perform technical work in context. Distant views can miss important aspects of the actual work situation and thus can miss critical factors that determine human performance in that field of practice. Integrating Natural History techniques is a strong guard against the Psychologist’s Fallacy, as a functional analysis, like an ethnography, pursues access to ‘authenticity’ in understanding practice.

There is a corollary to the psychologist's fallacy that is another pitfall in pursuit of authenticity.

> “Avoid assuming that those we study are less rational or have a weaker grasp on reality than we ourselves.” “This rule of method asks us to take seriously the beliefs, projects, and

> resources of those whom we wish to understand." (Law and Callon, 1995)

Managers, designers, researchers, and regulators often feel they are immune from the processes they study and manage, especially when the trigger for investing in studies is human "error", "limits" and "biases." This fallacy blocks effective functional analysis as it blinds the observer from seeing how behaviour is adapted in a universe of multiple pressures, uncertainty, and finite resources. Learning about cognitive work begins with adopting the point of view of practitioners in the situation before outcome is known or what anthropologists call an emic perspective (Fetterman, 1989).

The Danger of Whig Histories

> *A 'Whig' History – "produces a story which is the ratification if not the glorification of the present."*
> *Butterworth, 1965/1931, p. v)*

Historians refer to the same dangers that worried James as Whig histories which divide the world into the friends and enemies of progress or at least the teller's model of progress (which today is a uncritical anticipation of benefits to come from investment in the next new technology). They look at the past to tell a story that justifies and celebrates an aspect of the present state of affairs as the best. A Whig history does not examine how participants in those events viewed the world, but picks out those parts of the past that can be seen now as anticipating or retarding development of the story teller's model of the present. Interestingly, technology advocates often develop Whig accounts of cognitive work as a means to recruit resources and to pursue status in our culture.

> "Instead of seeing the modern world emerge as the victory of the children of light over the children of darkness in any generation, it is at least better to see it emerge as the result of a clash of wills, a result which often neither party wanted or even dreamed of, a result which indeed in some cases both parties would equally have hated, but a result for the achievement of which the existence of both and the clash of both were necessary." (Butterfield, 1965/1931, p. 28).

Following Butterfield's advice places a premium on 'authenticity' in the orchestration of methods (Woods and Christoffersen, 2002). The search for authenticity reveals a tradeoff: with stronger shaping of the conditions of observation, the greater disruption of the thing observed and greater risk of psychologist fallacy and its analogs. Natural History methods then are not ‘weak’ as in poor substitutes for the real methods of science or necessary evils due to practical limitations. Rather, because they are weak in shaping the conditions of observation, Natural History methods are a fundament al contributor to the goal of transcending limits to authenticity to capture how the strategies and behaviour of people are adapted to the constraints and demands of fields of practice. Because Natural History methods are weak in shaping the conditions of observation, they need to be coordinated with investigations that more strongly shape the conditions of observation based on tentative, partial knowledge of the essential factors in cognitive work (e.g., attempts such as Woods, 2002). But Staged World and Spartan lab techniques are weak in the sense that they depend on strong commitments to models of how cognition adapts to artefacts and demands of work.

Discovery and Functional Analysis

Ultimately the role and value of Natural History and Staged World techniques in a full program of research has been downplayed because the community of science has downplayed explicit consideration of the processes of discovery in science. Re-examining these families of techniques re-awakens the challenge of describing and facilitating processes of insight (remember the corollary to the Psychologist’s Fallacy applies to scientists who are not immune from the processes they study–limits on scientists’ ability to reflect on the practice of discovery they engage in).

To focus on the role of insight, the best method is to examine sets of functional analyses of cognitive work in context in detail. For example, consider a trio of natural history studies drawn from three different domains (Cook, 1998 in medicine; Hutchins, 1995b in aviation; Patterson et al., 1999 in space mission control) and how each discovered (or present possible accounts of) how the behaviour and strategies of practitioners are adapted to the various purposes and constraints of the field of activity (see Woods et al., 2002).

Each uses natural history methods to explore the interplay of agents, artefacts, and demands of the field of activity, but each enters the complex interplay of practice from different points. One is Richard Cook’s “Being ‘Bumpable’” (1998) study of the intensive care unit which explores the significance of a phrase used by practitioners. In this study modelling the significance of insiders‘ use of a highly coded argot of practice does one begin to appreciate the adaptive power of the intensive care unit. Another is Ed Hutchins’ “How a Cockpit Remembers Its Speed” (1995) study of cockpits, which begins with the speed management tasks during a descent. The cockpit is understood as a distributed system by understanding how certain apparently insignificant artefacts turn out to be highly adapted to the pressures of critical tasks and how these artefacts can be used in ways that make cognitive work more robust. A third begins with an artefact in space mission control centre – voice loops – and explores the role of this artefact in supporting successful collaborative activities when disruptions to plans occur (Patterson et al., 1999). As you examine the trio reflect on how the 8 points of functional modelling and analysis are illustrated in these specific functional analyses.

Warrant in Functional Analysis

I chose these three because of how they illustrate discovery processes at work in functional analysis and modelling. I chose three natural history studies as examples to eliminate any confusion with methods of verification.

First, note the striking absence of discussion of methods. To the degree method is discussed at all it is in terms of intimacy and extent of observation. In other words, claims to authenticity for natural history studies begin with intimate connection to the field of practice, i.e., to be able to observe from inside. The first kind of warrant for these functional accounts is one kind of claim to authenticity derived from the close contact for observation.

Second, each study provides observations and derives patterns to be accounted for. They note deeper aspects of the field of practice which are in need of explanation in a functional account. The warrant for these is related to the sets of observations and patterns noted in previous studies of this setting and the cognitive system themes invoked (is it expanding or modifying the patterns to be explained?). The latter presupposes a means to share observations and patterns from various field studies across a community of researchers interested in those cognitive system themes in settings such as the one in question.

Third, each study also makes a claim for warrant because the study provides a coherent account of the field of practice that goes behind the veil of surface characteristics to propose deeper dynamics. Providing an account or candidate functional explanation makes a sufficiency claim relative to the sets of observations and relative to other functional accounts previously offered--how well do they cover patterns noted in observation (note how this links to the role of computer simulations in functional modelling, e.g., Johnson et al., 2001). The latter presupposes a means to share functional models across a community of researchers interested in the relevant cognitive system themes and in characterizing the setting in question.

Ultimately, what is most striking in this trio of functional accounts is insight. From a discovery point of view, insight has no need of justification by pure method. It is a rare and precious commodity in need of facilitation; as such, insight is its own justification. Insight does not demand repeatability or indifference to the investigator or investigative team (what some criticize erroneously as 'subjectivity' in field research). Insight generates patterns from observations significant in their demand for explanation and it generates new or modified candidate explanations significant in the claim they stake to cover the relevant patterns derived from observation.

Insight is valuable; indeed, it is primary, in itself, for science though almost universally ignored as ineffable in accounts of science and in teaching of research methods. All programs of research in science embed processes of insight and methods to facilitate such insights; the methods and difficulties are generally just ignored in the descriptions of research process that are socially approved for their community.

The absence of, or even taboos on, discussion of insight and discovery do not facilitate these essential processes. As goes for any and all aspects of cognitive work, while there are individual differences in acquiring and demonstrating skill, discovery is trainable and aidable. One can view using and orchestrating field research techniques (the top two families in Figure 1) as a pragmatic program for facilitating insight.

Insight generates ideas which stand on their own and are open to argument and revision through new observations and alternative covering accounts. Insight as a discovery process is measured by models of argument (such as Toulmin's structure of argument or accounts of abductive explanation as a model of justification in argument), not by the syntax of inductive or deductive method. In other words, insight provides the linkage between field research techniques as discovery processes and laboratory research techniques as verification processes -- hence, Bartlett's dictum (1932) to begin experimental studies with observations of ongoing activities in situ.

Prediction and Innovation in Cognitive Work

Studies of cognitive systems as dynamic, adaptive systems have revealed technology change as a process of transformation creating new roles, changing what is canonical and what is exceptional, changing the kinds of errors and paths to failure, where practitioners and stakeholders adapt to achieve their goals and avoid failure (Woods and Tinapple, 1999; Woods and Dekker, 2000; Woods et al., 2002). This process of transformation and adaptation creates a challenge of prediction more daunting than the challenge of understanding the current adaptive balance in cognitive systems:

- How does one envision or predict the relation of technology, cognition and collaboration in a domain that doesn't yet exist or is in a process of becoming?
- How will envisioned technological change shape cognition and collaboration?
- What are the new roles for people in the system?
- How will practitioners adapt artefacts to meet their own goals, given new capabilities and mismatches to the actual demands and pressures they experience?
- How can we predict the changing nature of expertise and new forms of failure as the workplace changes?

Asking questions like these reveals that studies of cognitive work are going on around us constantly as developers and technology advocates make claims about the impact of objects-to-be –realized and that the adaptive response of practitioners and organizations to these changes provides data about the accuracy of those predictions (e.g., Woods and Sarter, 2000). The artefacts under development are not simply objects in the process of realization , but also hypotheses about how technology shapes cognition, collaboration and performance (Woods, 1998) subject to empirical jeopardy. In this light, functional analysis is a programme to harness and direct the ongoing natural experiments and to provide a base to accelerate progress in understanding the dynamics of distributed cognitive systems.

However, studies of cognitive work are not justified solely on the basis of abstracting patterns and building explanations of the adaptive interplay of people, technology and work. Ultimately, their purpose is to inform the design of systems for cognitive work (Woods et al., 2002). These studies occur as a part of process of organizational and technological change spurred by the promise of new capabilities, the dread of paths to failure , and continuing pressure for higher levels of performance and greater productivity (systems under 'faster, better, cheaper' pressure). This means investigators, through observing and modelling practice, are participants in processes of change in meaningful fields of practice. To become integral participants who represent the goal of supporting effectiveness in cognitive work given limited resources, the potential for surprise, and irreducible uncertainty, the families of methods need to be extended to meet the challenges of design (Woods et al., 2002):

- the *leverage* problem -- how do studies of cognitive work help decide where to spend limited resources to have the greatest impact (since all development processes are resource limited)?
- the innovation problem – how do studies of cognitive work support the innovation process (it is probably necessary but not sufficient as a spark for innovation)?
- the *envisioned world* problem – how do the results that characterize cognitive and cooperative activities in the current field of practice inform or apply to the design process *since the introduction of new technology will transform the nature of practice*? (a kind of moving target difficulty).
- the *adaptation through use* problem – how does one predict and shape the process of transformation and adaptation that follows technological change?
- the problem of *'error' in design* - designers' hypotheses, as expressed in artifacts, often fall prey to William James' psychologist's fallacy – the fallacy of substituting the designer's vision of what the impact of the new technology on cognition and collaboration might be, for empirically based but generalizable findings about the actual effects from the point of view of people working in fields of practice.

Summary

Functional analysis is a discovery programme of research that can support four basic values that define practice-centered design of what will be useful in cognitive work (Woods and Christoffersen, 2002).

- Transcending limits to authenticity to capture how the strategies and behaviour of people are adapted to the constraints and demands of fields of practice.
- Meeting the challenge of abstraction to find and explain patterns behind the surface variability.
- Sparking inventiveness to discover new ways to use technological possibilities to enhance human performance, to identify leverage points, and to minimize unanticipated side effects.
- Creating future possibilities as participants with other stakeholders and problem holders in that field of practice.

Functional analysis is the means to achieve these delicate balances--generic but relevant; finding in the particular the existence and expression of universal patterns; linking understanding to usefulness as fields of practice change in the face of new pressures and new capabilities.

Acknowledgments

This work was carried out in part through participation in the Advanced decision Architectures Collaborative Technology Alliance sponsored by the U.S. Army research Laboratory under Cooperative Agreement DAA19-01-2-0009.

References

Barley, S. and J. Orr. J., editors, (1997). *Between craft and science: Technical work in US settings*. IRL Press, Ithaca, NY.

Bartlett F. C. (1932). *Remembering: A study in experimental and social psychology*. Cambridge University Press, Cambridge, England.

Butterworth, H. (1965). *The Whig Interpretation of History*. W. W. Norton, New York (original 1931).

Carroll J. M. and Campbell, R.L. (1988). Artifacts as psychological theories: The case of human-computer interaction. IBM Research Report RC 13454, Watson Research Center, Yorktown Heights, NY.

De Keyser, V. (1992). Why field studies? In M. G. Helander and N. Nagamachi (Eds) *Design for manufacturability: A systems approach to concurrent engineering and ergonomics*. Taylor and Francis, London.

De Keyser, V. and Samercay, R. (1998). Activity theory, situated action and simulators. *Le Travail Humain*, 61(4), 305—312.

Dominguez, C., Flach, J., McDermott, P., McKellar, D., & Dunn, M. (in press). The conversion decision in laparoscopic surgery: Knowing your limits and limiting your risks. In J. Shanteau, K. Smith, & P. Johnson (Eds) *Psychological explorations of competent decision making*. Cambridge University Press, New York.

Feltovich, P.J., Spiro, R. J. and Coulson, R. L. (1997). Issues of expert flexibility in contexts characterized by complexity and change. In P. J. Feltovich, K. M. Ford & R. R. Hoffman (Eds) *Expertise in context*. MIT Press, Cambridge, Mass.

Fetterman, D. M. (1989). *Ethnography: step by step*. Sage Publications, Beverly Hills, CA.

Flach, J., Hancock, P., Caird, J. and Vicente, K., editors, (1995). *An ecological approach to human machine systems I: A global perspective*, Erlbaum, Hillsdale, NJ.

Flores, F., Graves, M., Hartfield, B. and Winograd, T. (1988). Computer systems and the design of organizational interaction. *ACM Transactions on Office Information Systems,* 6, 153-172.

Gibson, J. J. (1979). *An ecological approach to perception*. Houghton Mifflin, Boston.

Hoffman, R. R., & Deffenbacher, K. A. (1993). An analysis of the relations of basic and applied science. *Ecological Psychology*, *5*, 315-352.

Hoffman, R. and Woods, D.D. Studying cognitive systems in context. *Human Factors*, 42(1), 1-7, 2000.

Hollnagel, E., Pedersen, O. and Rasmussen, J. (1981). Notes on Human Performance Analysis. Roskilde, Denmark: Risø National Laboratory, Electronics Department.

Hutchins, E. (1995a). *Cognition in the wild*. MIT Press, Cambridge, MA.

Hutchins, E. (1995b). How a cockpit remembers its speeds. *Cognitive Science*, 19, 265--288.

James, W. (1890). *Principles of psychology*. H. Holt and Company, New York.

Johnson, P. E., Grazioloi, S., Jamal, K. and Berryman, R. G. (2001). Detecting deception: Adversarial problem solving in a low base-rate world. *Cognitive Science*, 25, 355—392.

Layton, C., Smith, P.J. and McCoy, C.E. (1994). Design of a cooperative problem-solving system for en-route flight planning: An empirical evaluation. *Human Factors*, 36:94--119.

Neisser, U. (1991). A case of misplaced nostalgia. *American Psychologist*, 46(1), 34—36.

Nyssen, A. S. and De Keyser, V. (1998). Improving training in problem solving skills: Analysis of anesthetists' performance in simulated problem situations. *Le Travail Humain*, 61(4), 387—402.

Patterson, E. S., Watts-Perotti, J.C. and Woods, D. D. (1999). Voice loops as coordination aids in Space Shuttle Mission Control. *Computer Supported Cooperative Work*, 8, 353—371.

Potter, S. S., Roth, E. M., Woods, D. D. and Elm, W. (2000). Bootstrapping multiple converging Cognitive Task Analysis techniques for system design. In J.M.C. Schraagen, S.F. Chipman, and V. L. Shalin,. (Eds) *Cognitive Task Analysis*. Erlbaum, Hillsdale, NJ

Rasmussen, J. (1986). *Information processing and human-machine interaction: An approach to Cognitive Engineering*. North-Holland: New York, 1986.

Rasmussen, J., Pejtersen, A. M., and Goodman, L. P. (1994). *Cognitive Systems Engineering*. Wiley, New York.

Roth, E.M., Bennett, K., and Woods. D.D. (1987). Human interaction with an 'intelligent' machine. *International Journal of Man-Machine Studies*, 27:479--525.

Salomon, G. (1991) Transcending the qualitative-quantitative debate: The analytic and systemic approaches to educational research. *Educational Researcher*, August-September, 10-18.

Sarter, N.B. and Woods, D.D. (1995). How in the world did we get into that mode? Mode error and awareness in supervisory control. *Human Factors*, 37(1), 5-19.

Sarter, N.B. and Woods, D.D. (1997). Team play with a powerful and independent agent: A corpus of operational experiences and automation surprises on the Airbus A-320. *Human Factors*, 39, 553—569, 1997.

Sarter, N.B. and Woods, D.D. (2000). Team play with a powerful and independent agent: a full mission simulation. *Human Factors*, 42, 390—402, 2000.

Toulmin, S., Rieke, R. and Janik, A. (1984). *An introduction to reasoning*. Macmillan, New York.

Vicente, K. (1998). *Cognitive Work Analysis: Toward safe productive and healthy computer based work.* Erlbaum, Hillsdale, NJ.

Von Uexkull, J. (1934). A Stroll through the Worlds of Animals and Men. Reprinted in C. Schiller (Ed) *Instinctive behavior*, International Universities Press, New York, 1957.

Watts-Perotti, J. and. Woods, D. D. (1999). How experienced users avoid getting lost in large display networks. *International Journal of Human-Computer Interaction*, 11(4), 269—299.

Woods, D. D. (1988). Coping with complexity: The psychology of human behavior in complex systems. In L.P. Goodstein, H.B. Andersen, and S.E. Olsen (Eds) *Mental models, tasks and errors*, Taylor & Francis, London, (p. 128-148).

Woods, D. D. (1993). Process-tracing methods for the study of cognition outside of the experimental psychology laboratory. In G. Klein, J. Orasanu, R. Calderwood, and C. E. Zsambok, C. E. (Eds) *Decision making in action: Models and methods* Ablex, Norwood, NJ, pp. 228-251.

Woods, D. D. (1998). Designs are hypotheses about how artifacts shape cognition and collaboration. *Ergonomics*, 41, 168—173.

Woods, D.D. (2002). Steering the Reverberations of Technology Change on Fields of Practice: Laws that Govern Cognitive Work. Proceedings of the 24th Annual Meeting of the Cognitive Science Society.

Woods, D. D. and Christoffersen, K. (2002). Balancing practice-centered research and design. In M. McNeese and M. A. Vidulich (Eds) *Cognitive Systems Engineering in military aviation domains.* Human Systems Information Analysis Center, Wright-Patterson AFB, OH.

Woods, D.D. and Dekker, S. W. A. (2000). Anticipating the effects of technological change: a new era of dynamics for Human Factors. *Theoretical Issues in Ergonomic Science*, 1(3), 272—282.

Woods D.D. and Patterson, E. S. (2000). How unexpected events produce an escalation of cognitive and coordinative demands. In P. A. Hancock and P. Desmond (Eds) *Stress, workload and fatigue*. Erlbaum, Hillsdale, NJ.

Woods, D.D. and Sarter, N. (2000). Learning from automation surprises and going sour accidents. In N. Sarter and R. Amalberti (Eds) *Cognitive Engineering in the aviation domain*, Erlbaum, Hillsdale, NJ.

Woods, D.D. and Tinapple, D. (1999). W^3: Watching Human Factors Watch People at Work. Presidential Address, 43rd Annual Meeting of the Human Factors and Ergonomics Society, September 28, 1999. Multimedia Production at url: http://csel.eng.ohio-state.edu/hf99/

Woods, D. D., Tinapple, D. Roesler, A. and Feil, M. (2002). Studying Cognitive Work in Context: Facilitating Insight at the Intersection of People, Technology and Work. Cognitive Systems Engineering Laboratory, Institute for Ergonomics, The Ohio State University, Columbus OH at url: http://csel.eng.ohio-state.edu/woodscta