

1 Introduction

The value of experience and of understanding how to apply knowledge is described in a wonderful story, a story shared so often it has become one of those enduring urban legends. The story has many different versions,¹ involving various professions and famous people. In one version of the story the scientist and inventor Nikola Tesla visited Henry Ford at his automobile factory. The factory was having some kind of difficulty with its systems, and Ford asked Tesla if he could help identify the problem area. Tesla walked up to a wall of boilerplates, scanned them briefly, and then made an "X" in chalk on one of the plates. Examination of the boilerplate showed that it was indeed faulty. Ford was impressed, and told Tesla to send an invoice. The bill arrived, for \$10,000. Ford, never known for his generosity, was astonished at the cost of writing an "X" on the boilerplate, and asked for a breakdown. Tesla sent another invoice, which read:

Marking wall:	\$1
Knowing where to mark:	\$9,999

This story speaks directly to the purposes and goals of this book in two respects.

First, the story illustrates the "why" of Cognitive Task Analysis (CTA). What is it that Tesla knows, and how does he know it? What tells him what to do, with Henry Ford (not the most patient of men, by many accounts) looking over his shoulder? Capturing that knowledge and reasoning is one of the things CTA can do.

Second, the story illustrates the "how" of CTA. Cognitive Task Analysis can be thought of as a set of tools in a toolkit. Like any tool, CTA can be employed well and wisely, or it can be employed poorly or inappropriately. What tool would you use if you wanted to understand how Tesla was able to grasp the nature of the problem so quickly?

This book is about having the tools and the toolkit to understand how people think: how their minds work, what they struggle with, and how they manage to perform

complex work adeptly and pluck inventive solutions out of difficult, sometimes dangerous, situations. Our purpose in writing the book is to help people learn how to do CTA—how to collect data about cognitive processes and events, how to analyze it, and how to communicate it effectively.

What CTA Offers

All CTA procedures have the general goal of helping researchers understand how cognition makes it possible for humans to get things done and then turning that understanding into aids—low or high tech—for helping people get things done better. The “work” may be that of a consumer who is using a product for the first time, or that of a weather forecaster who is trying to cope with data overload during a thunderstorm, or that of a firefighter who must figure out in seconds or minutes what to do about a dangerous situation. In all these cases, performance depends on what people know, what they perceive, what they believe, and how they think.

In many applications of CTA, the work is conducted in what are called “complex cognitive systems” (Hoffman and Woods 2000). These are work settings in which the knowledge and reasoning of individuals play a role (of course), but so do the cognition and reasoning of larger groups of people, including teams and even entire organizations. In addition, these complex cognitive systems often involve people interacting with computers and also interacting with each other via computers in intricate networks of humans and technology. Cognitive Task Analysis can show what makes the workplace work and what keeps it from working as well as it might.

Over the past several years an unusually broad population of individuals has become interested in CTA. People want to know how to do it, how to use it, and how to make it work for their organizations. Systems analysts need CTA methods to develop user specifications for new computer technologies. Trainers and instructional systems designers imagine applying CTA in order to describe the cognitive processes that need to be trained and how best to train them. Market researchers clearly understand the benefit of a lens into the minds of consumers and are discovering that CTA can offer ways to expose the thought processes involved in purchase decisions and product use. Program managers tasked with building new or improved technologies for military clients are embracing the notion that front-end analyses of the operators can help ensure that their systems work effectively. They look to CTA as a tool for understanding the cognitive requirements of those operators and the most effective combinations of humans and technology. Employers faced with a range of personnel issues wonder whether CTA could provide insights into selecting and retaining personnel. Healthcare pro-

viders and medical technology developers have begun to look to CTA to assist with enhancing patient safety and to identify and apply lessons learned from errors and accidents. Military commanders, faced with increasingly complex and dangerous missions, seek ways to best support planning and decision making in the field.

Across these many different types of work, there is recognition that CTA yields information people need. It provides leverage on deeply challenging problems, and when done well it provides solutions that can make a difference.

In writing this book, we hope to increase greatly the number of people with the skills and knowledge to conduct high-quality CTA. There are individuals and organizations with problems they cannot solve and opportunities they want to take advantage of. They need CTA tools and methods, and people who know how to apply them skillfully, across a range of problems and issues.

Unpacking Cognitive Task Analysis

A good place to begin is with some definitions.

Cognitive When the tasks that people are doing are complex, it is not enough to simply observe people's actions and behaviors—what they do. It is also important to find out how they think and what they know, how they organize and structure information, and what they seek to understand better. This is a principal reason why the word “cognitive” begins the phrase Cognitive Task Analysis. Cognitive Task Analysis is a family of methods used for studying and describing reasoning and knowledge. These studies include the activities of perceiving and attending that underlie performance of tasks, the cognitive skills and strategies needed to respond adeptly to complex situations, and the purposes, goals, and motivations for cognitive work.

Task What about the second word in CTA, the notion of a “task?” It may seem straightforward to think about “task” as people engaged in discrete activities or sequences of activities aimed at achieving some particular goal. This is a traditional notion of “task.” But in complex cognitive systems, it is not always the literal action sequences—the steps—that matter as much as the fact that practitioners are trying to get things done; they are not simply performing sets of procedures. Therefore, we define task in this broader sense as the outcomes people are trying to achieve.

Analysis We use the term “analysis” deliberately. Literally, to analyze something is to break it into parts in order to understand both the component parts and their relationship in making up the whole. Cognitive Task Analysis methods provide procedures for systematic, scientific examination to support description and understanding. For scientists interested in pursuing research questions, we believe CTA presents a number

of opportunities and challenges (Klein, Phillips et al. 2003). For practitioners who are interested in CTA primarily to develop tools and technologies, the "analysis" component of CTA is particularly important. Cognitive Task Analysis provides a process for systematically identifying key cognitive drivers in many types of applications.

Topics and Focus

This book is about both the "why" and the "how" of methods for studying thinking and reasoning in the course of performing real-world tasks in complex and dynamic work settings.

We primarily study adults in the workplace, and the methods described here have been developed within the world of adult work. The tasks that make up the working life of firefighters, nurses, military commanders, weather forecasters, or pilots may seem far from commonplace to you and me, but they are what fill the work lives of people in each of these occupations. Everyday tasks can also mean decisions and choices about products that face consumers on a daily basis.

Much of what we have written about here focuses on people's reports about their own, lived experience—their stories and examples and their understanding of the work they do. As we will show, CTA study can reveal the risks, time elements, opportunities, and mistakes that confront people as they work. It can help us understand the workplace: the technologies, tools, work conditions, stressors, and team interaction modes that all contribute to cognitive performance. Cognitive Task Analysis can help us consider hypothetical conditions, such as the influence of system X or technology Y, or a work practice that increases tempo by a factor of two. These are all questions that have been posed by people using CTA.

We also share some of our own stories and experiences as CTA practitioners, including what we have learned about how to apply CTA methods and how to get them to work well. What factors make the difference between a great interview and a folder full of notes that don't really say much? What techniques can determine whether data analysis yields an elegant Concept Map that conveys crucial knowledge about a domain or a bewildering mass of lines and arrows? How do seasoned CTA practitioners use a given set of methodologies to investigate problems and issues?

We have several goals in writing this book. First, the book is intended to help people learn how to do a CTA: how to collect data about cognitive work, how to analyze it, and how to use and communicate it effectively. We offer examples, guidance, and our

own experiences using CTA to investigate a wide range of problems, questions, issues, and domains. Our express purpose is to give people knowledge and tools that will allow them to conduct CTA studies.

Second, the book is intended to convey the reasons to do CTA: What is CTA good for? What sorts of questions can it answer? What problems can it address?

We offer a look at how experienced practitioners apply CTA in such arenas as aviation, the military, national security, health care, firefighting, emergency response, manufacturing, nuclear power, consumer research, and many others. We hope to give readers an understanding of why CTA matters and how it can make a difference. Details of methods do not mean much if the reasons for using the methods are not clear.

We chose to write about CTA methods that we have found useful for understanding cognitive aspects of work in order to share the expertise we have gained in performing CTA and provide an insider's perspective on the process of CTA data collection. These methods are the tools we use in our own work to pursue certain types of questions and to explore the cognitive landscape. For us, core questions concern how people think and reason in complex, dynamic settings that characterize real-world tasks. That said, every method carries particular assumptions with it. Every method opens some doors and leaves others tightly shut. Knowing the strengths and limitations of a particular method is critical to using it well.

We also believe that skilled and effective use of CTA methods means understanding something about the cognitive issues they have been designed to illuminate. Our experience helping people learn to use CTA methods has left us convinced that skilled CTA practice has to combine knowledge of specific methods and techniques with some conceptual grounding in cognitive theory and research. It simply isn't sufficient to pick up a tool and place it in your toolbox. To use it well, you have to also understand why the tool was fashioned in a particular way and how the tool came to be. Doing CTA well requires knowing what a cognitive perspective can offer for understanding problems and issues of work. With that in mind, we provide an overview of current work on cognition—of how people think, reason, and make decisions—in the real world.

If we are successful in meeting these two goals, we will satisfy a third one: to expand the circle of CTA practitioners. We hope to foster a community of research practitioners who have the necessary skills and knowledge to conduct CTA, who can provide useful information and effective application to individuals and organizations, and who can advocate for its use. Although learning to do CTA can be a demanding experience, the insights and perspective to be gained make the challenges of learning to do it

worth the investment. One primary reason for writing this book is to present the details of the CTA process, to provide a road map for how to conduct a CTA. We want to make the methodology more accessible and the skills involved in CTA practice more attainable.

In terms of coverage of the types of CTA methods and applications, the book is selective rather than inclusive. We chose not to write a survey volume with brief summaries of many different methods. There are some excellent survey volumes and review articles available, and we recommend that you spend some time with them in order to gain an overview of the breadth of methodologies available (Bonaceto and Burns 2003; Cooke 1994; Hoffman 1987; Jonassen, Tessmer, and Hannum 1999; Patrick 1992; Schraagen, Chipman, and Shalin 2000). Instead, we have presented an overview of the field and then homed in on a smaller number of methods to provide detailed descriptions of the CTA process, offer specific guidance, describe examples from our own work, and supply practical tips. We believe the narrower focus and specificity of detail will be particularly helpful for people who are new to CTA. We also expect that people who are experienced with other forms of behavioral task analysis, or who are seasoned interviewers, will find this book interesting and useful for expanding their skills to encompass cognitive components of performance.

Talking to the Reader

Across the pages of this book we present many suggestions. Some are pretty firm guidance about the "how to" of CTA. Some convey lessons learned or cautionary tales. Some are specific descriptions of steps in procedures and may seem rather prescriptive. Others are best regarded as advice.

In our efforts to present advice and guidance, we occasionally speak directly to you, the reader, as an individual who is interested in learning about and possibly conducting CTA. Thus, for example, we say in chapter 2:

The three primary aspects of CTA are knowledge elicitation, data analysis, and knowledge representation. Each of these aspects is critical to a successful CTA study. Many people equate CTA with the first aspect, eliciting the knowledge, because traditionally that has received the most attention. But if you don't do a good job of analyzing your data, why bother collecting them? And if you don't represent your findings so that others can understand them and why they matter, what have you accomplished?

It is our hope that this style of directly addressing the reader is not perceived as overly familiar or informal. We are simply trying to communicate clearly and in a way that is meaningful to you, the reader, as we present the concepts covered in this book.

Organization of the Book

The book is organized into three major sections:

Part I, "Tools for Exploring Cognition in Context," provides detailed guidance for planning and carrying out CTA. It includes chapters on capturing knowledge and on capturing the way people reason. We rely on this distinction throughout the book: CTA investigates what people know and how they think.

Part II, "Finding Cognition," provides a perspective on studying cognition in real-world settings and what an expanded view of cognition—a macrocognitive framework—offers. We describe some of the issues that surround CTA and what it means to study cognition in context. We end the section by exploring the challenges of rapidly changing technology.

Part III, "Putting CTA Findings to Use," describes key issues in applying CTA findings to several applications areas: technology development, training and instructional design, and market research. We also present a chapter on the role of CTA in the development of measures for evaluating cognitive work.

Our intent in writing this book is to share what we have learned about CTA, from our experience in the field to the concepts and models we draw on. We have offered examples and suggested ways to apply CTA findings to real-world problems and issues. We hope this book provides you with some tools you can use in your own practice and that the CTA methods can help you discover how people like Tesla know where to put their chalk marks.

2 Overview of Cognitive Task Analysis Methods

In this chapter, we briefly survey the leading methods for conducting Cognitive Task Analysis (CTA). The purpose of CTA is to capture the way the mind works, to capture cognition. The researcher or practitioner carrying out a CTA study is usually trying to understand and describe how the participants view the work they are doing and how they make sense of events. If they are taking effective action and managing complex circumstances well, the CTA should describe the basis for their skilled performance. If they are making mistakes, the CTA study should explain what accounts for the mistakes. Cognitive Task Analysis studies try to capture what people are thinking about, what they are paying attention to, the strategies they are using to make decisions or detect problems, what they are trying to accomplish, and what they know about the way a process works.

The three primary aspects of CTA are *knowledge elicitation*, *data analysis*, and *knowledge representation*. Each of these aspects is critical to a successful CTA study. Many people equate CTA with the first aspect, eliciting the knowledge, because traditionally that has received the most attention. But if you don't do a good job of analyzing your data, why bother collecting them? And if you don't represent your findings so that others can understand them and why they matter, what have you accomplished?

One way to get an overview of CTA is to understand how many methods there are, the sorts of labels applied to them, and what types or categories they belong to. But describing the larger picture of CTA can be quite a challenge. Cognitive Task Analysis has developed from many diverse traditions (see chapter 9) with differing root metaphors, terminologies, prevailing methodologies and testbeds, areas of application, and standards for what qualifies as worthwhile—or even what qualifies as “cognitive.” In the first section of this chapter, we review methods of knowledge elicitation and present some ways to distinguish among them. In the second half of the chapter, we discuss approaches to CTA data analysis and representation.

Knowledge Elicitation Methods

Knowledge elicitation is the set of methods used to obtain information about what people know and how they know it: the judgments, strategies, knowledge, and skills that underlie performance. There are many different knowledge elicitation methods, so many that simply tracking them all down is a challenge. Tables 2.1¹ and 2.2² illustrate the diversity of tools and techniques available to CTA practitioners. They contain the methods that can be found at two different websites created with the express purpose of providing information about CTA. The CTA Resource site (<http://www.ctaresource.com>) is maintained by Aptima, Inc. All methods identified as "knowledge elicitation" within the CTA methods summary information provided at CTA Resource are presented in table 2.1. The Survey of Cognitive Engineering Methods and Uses was developed by the MITRE Corporation and can be accessed through their Mental Models website (<http://mentalmodels.mitre.org/index.htm>). Information provided at that website is presented in table 2.2.

Both websites provide descriptions and references for individual methods, along with a number of other resources. In addition, they each organize methods into classes or types of knowledge elicitation, and those categories and method assignments are included. However, tables 2.1 and 2.2 are by no means exhaustive. Other sources present additional methods and various ways of organizing and categorizing them (e.g., Cooke 1994; Hoffman et al. 1995; Jonassen, Tessmer, and Hannum 1999; Schraagen, Chipman, and Shalin 2000).

The first thing to notice about these two tables is the sheer number and variety of knowledge elicitation methods and tools. Even though the CTA Resource and Mental Models websites have similar goals, the methods they list and the categories they use to organize them are considerably different. Some of that difference may be due to the lack of generally accepted definitions and qualifiers for what counts as CTA in the first place.

Another reason for the diversity we see in the tables is that methods have been assimilated into the family of knowledge elicitation techniques by a number of different pathways. Some methods have been developed specifically for CTA (e.g., Goal Directed Task Analysis; PARI method); others have been purposefully adapted from methods initially created for other uses (e.g., Concept Mapping; Cloze Technique; Table Top Analysis). Still others have migrated into the field as researchers and practitioners began applying tools developed for purposes such as task analysis and instructional systems design to cognitive issues (e.g., Repertory Grid; Activity Sampling; Hierarchical Task Analysis).

Table 2.1

Knowledge elicitation categories and methods

<i>Interview</i>	<i>Focused observation</i>
Applied Cognitive Task Analysis	Interruption analysis
Cloze Experimental/Minimal Scenario technique	Job analysis
Cognitive Function Model	Operator Function Model
Comparing two or more representations	Process Tracing/Protocol analysis
Critical Decision Method	Role play
Critical Incident Technique	Shadowing another
Critical Retrospective	Simulator/mockup
Crystal ball/stumbling block	Structured observation
Diagram drawing	Time line analysis
Distinguishing goals	Unstructured interview
Dividing the domain	Walk-throughs and talk-throughs
Focus groups/joint application development	<i>Textual</i>
Functional Flow Analysis	Content analysis
Group discussion	Management Oversight Risk Tree technique
Group interview	<i>Psychometric</i>
Hazard and Operability Analysis	Cloze Experimental/Minimal Scenario technique
Identifying aspects of the representation	Concept listing
Information Flow Analysis	Controlled association
Interaction analysis	Drawing closed curves
Interruption analysis	Eliciting estimations of probability and utility
Job analysis	Free association
Operator Function Model	Function Allocation issues and tradeoffs
Precursor, Action, Result, Interpretation method	Graph construction
Questionnaires	Hierarchical sort
Reclassification/goal decomposition	Laddering
Retrospective/aided recall	Likert scale
Self critiquing/eidetic reduction	Magnitude estimation
Step listing	Multidimensional card sorting
Tabletop analysis	Nonverbal reports
Teachback	P Sort
Think-aloud	Paired comparison
Twenty Questions	Q Sort
Workflow model	Repeated sort
<i>Observation</i>	Repertory Grid
Active participation	Statistical modeling/Policy capturing
Activity sampling	Step listing
Cognitive Function Model	Structural analysis techniques
Controlled simulated observations	Triad comparison
Field observations/ethnographic methods	

Source: www.ctaresource.com

Table 2.2

Survey of cognitive engineering methods and uses

<i>CTA Methods</i>	Discourse/conversation/interaction analysis
Applied Cognitive Task Analysis	Exploratory Sequential Data Analysis
Critical Decision Method	Interruption analysis
Cognitive Function Model	Minimal scenario technique
Cognitive-Oriented Task Analysis	Retrospective/aided recall
Decompose, Network and Assess method	Shadowing another
Goal-Directed Task Analysis	Shadowing self
Hierarchical Task Analysis	Simulators/mockups and microworld simulation
Interacting cognitive subsystems	Tabletop analysis
Knowledge Analysis and Documentation Systems	Think-aloud problem-solving/protocol analysis
Precursor, Action, Result, Interpretation method	Wizard of Oz technique
Skill-based CTA framework	<i>Knowledge Elicitation—Conceptual Methods</i>
Task knowledge structures	Cluster analysis
<i>Knowledge Elicitation—Interview/Observation</i>	Conceptual graph construction
Field observations/ethnographic methods	Decision analysis
Group interview	Diagramming
Questionnaires	Hierarchical sort
Step listing	Influence diagram construction
Structured interviews	Laddering
Teachback	Likert scale elicitation
Twenty Questions	Magnitude estimation
Unstructured interviews	Multidimensional scaling
<i>Knowledge Elicitation—Process Tracing Methods</i>	P Sort
Activity sampling	Q Sort
Cloze Experimental technique	Rating and sorting tasks
Critical Incident Technique	Repertory Grid
Critiquing	Structural analysis techniques
Crystal ball/stumbling block	

Source: www.mentalmodels.mitre.org

In addition, the methods identified here vary from very specific tools and techniques (e.g., Applied Cognitive Task Analysis [ACTA]; Cognitively Oriented Task Analysis) to entire classes of methodologies used across a wide range of problems in psychology and human factors (e.g., interviews, error analysis, questionnaires). None of this is wrong, but the jumble of terms and descriptive levels is certainly confusing.

Given the mixture of terms, sources, and levels it is probably not so surprising that there is no single, well-accepted taxonomy of methods available. In fact, both of the classification schemes in tables 2.1 and 2.2 make sense, but neither seems to capture fully the multiple dimensions that exist within the overall class of knowledge elicitation methods.

We have found it useful to divide knowledge elicitation tools along two separate, intersecting dimensions: how the data are collected, and where a particular method is focused.

Types of Data Collection Methods

One way of classifying CTA knowledge elicitation is by the way the data are collected—what sort of activity is involved in eliciting information? We can distinguish four ways to gather data: interviews (i.e., asking people questions), self-reports (i.e., people talk about or record their behavior and strategies), observations of performance or task behavior, and automated collection of behavioral data. Each of these activities is discussed in the sections that follow.

Interviews

The most common CTA method is a structured interview. Interview methods are widely used in CTA practice, and for good reason. Interviews are efficient—they avoid the investment of time and effort and the logistical complications that often occur with observations. Interviews can also elicit information about issues that are easily missed by the other methods. For example, if you do not conduct an observation at precisely the right moments, you might miss key dynamics or critical elements of task performance. Anticipating when those moments are likely to occur is much more easily done in laboratory settings than in real-world data collection. Moreover, in a study of naval officers (Kaempf et al. 1992), we found that virtually every incident we studied via interview hinged on some subtle issue of personality clash or lack of confidence in the skill of a cohort. These types of dynamics are rarely incorporated into simulated task scenarios and can be difficult to discern in behavioral observations.

Many CTA practitioners view interview data as extremely rich, but best treated as exploratory data and as a source of hypotheses. Findings from one interview can be treated with greater confidence when they are replicated across interviews with other participants or are corroborated by other methods.

Interviews have disadvantages as well. Many CTA methods require interviews with highly skilled professionals, and scheduling even an hour of time with busy professionals can be difficult. Moreover, getting good data depends on participants' being able and willing to reflect deeply on their performance and their work. People may be reluctant to divulge details of some events, they may be mistaken, or they may have limited information about what happened or why. Another drawback to interview methods is

that many of them require well-trained interviewers. That training requires knowledge and skill that goes well beyond understanding of standard data collection and analysis procedures.

Self-Reports

A second variety of methods are based on participants generating data on their own. These methods vary from highly structured formats, such as surveys and questionnaires, to open-ended formats such as diaries and logs. Clearly, self-report formats have an efficiency advantage, because the data collection doesn't require an interviewer or skilled data collector to be present. The quality of data generated by questionnaires and rating scales obviously depends in part on the instrument itself. There is an entire scientific field and set of methodologies that surround development of scales and questionnaires that are psychometrically sound—that are valid and reliable and can be counted on to measure what they claim to measure. Simply compiling some questions and providing the list and a pencil to participants is not necessarily going to produce insights.

Questionnaires and rating scales can be valuable tools for gathering information on the concepts and items they contain. There can be advantages in knowing what sort of information you are likely to get. The disadvantage is that structured questionnaires and rating scales do not allow for the elements of discovery and exploration that are available in more open-ended reporting formats. Diaries and logs can offer those opportunities because they provide greater flexibility of format and content. However, data quality depends a lot on participants' motivation and willingness to complete entries consistently.

Finally, self-report methods assume that respondents are capable of "self CTA" and of reporting tacitly held knowledge, subtle cues and perceptions, and other cognitive elements on their own. That assumption is not backed up by research—in fact the evidence suggests quite the opposite: people have considerable difficulty reporting on their own cognitive processes (Nisbett and Wilson 1977; Wilson 2002). And as people gain experience and higher levels of skill, it becomes increasingly difficult for them to articulate the basis for their expertise and the judgments, decisions, and assessments they make so capably (Chi and Bjork 1991; Chi, Glaser, and Farr 1988; Feltovich, Ford, and Hoffman 1997; Klein and Hoffman 1993).

Observation

Observing people perform their work offers advantages and unique opportunities. If on-site observations are feasible (they often are not), we strongly advocate that the

requires knowledge
ction and analysis

data on their own.
eys and question-
self-report formats
require an inter-
rated by question-
t itself. There is an
lopment of scales
d and reliable and
y compiling some
necessarily going

ing information on
owing what sort of
red questionnaires
d exploration that
gs can offer those
content. However,
gness to complete

of "self CTA" and
nd other cognitive
earch—in fact the
culty reporting on
02). And as people
ifficult for them to
s, and assessments
rr 1988; Feltovich,

ie opportunities. If
advocate that the

CTA researcher take advantage of the opportunity. There are insights and types of information that it is simply not possible to get any other way. Observations provide opportunities for discovery and exploration of what the actual work demands are; what sorts of strategies skilled workers have developed for coping; how work flows across the environment, the team, and the shift; and communication and coordination issues (Roth 2002).

Observation can be particularly effective when the researchers are well trained in the phenomenon they are studying and do not require a lot of structure for their data-collection activities. Structured observation procedures, such as predetermined formats for sampling activities, may be desirable if the research demands some degree of quantification. Without an observational checklist or other predetermined format, the researchers may wind up figuring out the coding categories afterward and wrestling with category descriptions and coding instructions. They may also find uneven coverage in their data because observers were unaware of its significance. However, advance structuring can also render the observer less sensitive to what is actually going on or unable to take advantage of a rich opportunity—particularly if what is occurring is different from what was expected.

The primary disadvantage of observational methods is that they simply may not be feasible, either because the observation opportunity represents unacceptable risk to observers, or because observers get in the way and impede the ability of personnel (e.g., firefighters, medical personnel, military forces) to respond fully to a critical situation. Other issues in observational data collection are that the events observed may not be typical and that the observers have to be highly skilled in order to capture what is going on.

In our view, observation is best coupled with other forms of data collection such as interviews to find out how the participants were viewing the events. Merely recording the events and actions taken can result in a misleading or cognitively shallow account.

Automated Capture

The collection of CTA data can be handled by computers. This approach has not been widely used to date, but we expect that to change. One example is the Situation Awareness Global Assessment Technique (SAGAT) developed by Endsley (1988b; Endsley and Garland 2000). Previously, de Groot (1946/1978) had described a strategy for comparing chess players at different skill levels. The de Groot method was to have a player study a game in progress and then unexpectedly remove all the pieces. The player would be asked to reconstruct the board. De Groot found that players were more accurate when reconstructing actual board positions than they were in reconstructing

randomly placed pieces and that more skilled players were more accurate in reconstructing coherent board positions than novices. The SAGAT method is an adaptation of de Groot's technique, basically a form of "time freezing." In the midst of a computer-driven simulated mission, all of the instruments go blank and the pilots are asked to reconstruct the instrument values. SAGAT is a measure of situation awareness. According to Endsley, the better a person's situation awareness, the more accurate the reconstruction.


Advantages and drawbacks to automated capture are similar to those we noted for questionnaires and surveys. Automated capture offers ease and precision of data collection. The potential naturalness of embedding data capture in the computer-guided flow of events has benefits and appeal. Disadvantages include the effort to program the system, the difficulty of determining when to interrupt task performance, and the insensitivity of the knowledge capture to nuances, confusions, and questions that the participant might raise. Another limitation is that the automated capture is not well suited for follow-up interrogation or deeper probing to follow up participants' comments. Automated capture doesn't lend itself to the back-and-forth, interactive data gathering that is possible in interview and observational settings.

Types of Data Targets: Where Are Methods Focused?

A second set of CTA categories addresses *where* to look for data, rather than how to get them. Here, we consider four different facets of the data collection target: its location along a continuum in time, in realism, in difficulty, and in generality.

The two ways of categorizing knowledge elicitation methods intersect. Table 2.3 illustrates the intersections between how to look and where to look. For example, one can conduct an interview about a retrospective event or observe a videotape of a past event; one can observe exercises and events as they are occurring in present time and interview participants as it unfolds, and so on.

Time



How close to "here and now" is the data target?

Studying cognitive performance, we can work with participants as they are in the midst of performing a task or working with a problem and collect data concurrently

Table 2.3

Key attributes of CTA methodology

How to Look: → Where to Look: ↓	Interview	Self-Report	Observation	Automated Capture
Where in TIME: past/present/future				
Where in REALISM: real world/simulation or scenarios				
Where in DIFFICULTY: routine tasks/challenging tasks				
Where in GENERALITY: abstract knowledge/specific events				

in present time. We can also elicit data about events that have happened at some point previously (i.e., retrospectively). We might ask about events that are likely to happen in the future, or about hypothetical possibilities. Each of these possibilities has advantages and drawbacks.

One of the most powerful means of eliciting knowledge is to study prior incidents that were extremely challenging, to see what made them so difficult and to learn why decision makers succeeded or failed. Flanagan (1954) introduced the idea of using critical incidents to describe the nature of work, and Hoffman (1987) showed that the study of tough cases resulted in high degrees of efficiency in eliciting knowledge from experts. Klein, Calderwood, and MacGregor (1989) and Hoffman, Crandall, and Shadbolt (1998) described the use of a Critical Decision Method for knowledge elicitation that relies on retrospective accounts. We discuss CDM in detail in chapter 5.

Retrospective data can provide access to particular types of incidents. For example, we might elicit data about critical incidents of a particular type (e.g., emergency response to tornados) or to a specific event (e.g., emergency response to Hurricane Katrina). Retrospective accounts are usually studied via interview, but it is also possible to rely on self-report. A guided questionnaire could enable a person to review a prior incident and provide some description of how judgments and decisions were made. Retrospective data collection allows researchers to focus on particular types of events and aspects of cognitive performance.


The primary disadvantage of retrospective incident accounts is that people may forget or even distort key details. Memory is fragile. Therefore, data from retrospective accounts should be treated as a source of hypothesis or as a record of events that

requires independent verification. For that reason, we recommend the use of converging operations and other forms of cross-checking of results of retrospective inquiries.³

Collecting data concurrently in time avoids many of the memory difficulties noted above (Ericsson and Simon 1984). It also allows data collectors to observe and document aspects of the situation independently of the participants' perceptions. However, concurrent data collection does not necessarily ensure better access to cognitive processes if interviewing or other types of self-report are part of the data collection process. Depending on the type of activity, the act of reporting about ourselves and our behavior can introduce biases and distortion into the data. Moreover, reporting on an activity while one is performing it can disrupt and alter the very cognition we're attempting to study (Melcher and Schooler 1996; Schooler and Engstler-Schooler 1990). The distortions and disruptions may limit the circumstances and types of tasks in which these methods can appropriately be used. A fireground commander might be willing to "think out loud" during a field exercise, but doing so during an actual event would be an unacceptable distraction.

Asking participants to report on hypothetical or imagined future events can provide interesting data when those reports are tightly linked to actual events (for example, asking participants what it would have meant if a key aspect of an incident they experienced were altered in a particular way).

Realism



How much like the real world is the data target?

CTA data are often gathered in real-world settings. Most CTA studies are focused on performance of real-world tasks, and collecting data in real-world circumstances remains the gold standard for CTA researchers. However, there are many other types of settings that can provide CTA data that are valuable, interesting, and informative. Cognitive Task Analysis data collection is frequently carried out in simulations or in contrived or created settings. Many military field exercises contain virtually every aspect of real events save for live ammunition. Highly sophisticated flight simulators are capable of mimicking actual events—in fact, simulators have been used to "replay" incidents, including accidents and near misses, in order to better understand how they may have occurred. Similar high-fidelity simulators are now available in many fields and are used for training individuals and teams and for design, test, and evaluation of new tools and

the use of converg-
pective inquiries.³
by difficulties noted
observe and docu-
ceptions. However,
ss to cognitive pro-
a collection process.
lves and our behav-
orting on an activ-
on we're attempting
chooler 1990). The
s of tasks in which
er might be willing
actual event would


events can provide
vents (for example,
incident they expe-

udies are focused on
d circumstances re-
many other types of
id informative. Cog-
nulations or in con-
ually every aspect of
nulators are capable
"replay" incidents,
how they may have
y fields and are used
on of new tools and

technologies. Simulations range from these very high-fidelity versions to paper-and-pencil scenarios that present key aspects of an incident and ask participants to play out action, detect problems, assess situations, and make decisions. What matters in lower-fidelity simulations is the degree of cognitive authenticity the scenario is capable of creating, regardless of its technical simplicity. Computers are being increasingly used to create gaming environments and to present humans with varieties of experience in simulated settings and artificial worlds. All of these settings offer the potential for putting humans into cognitively complex and challenging circumstances in order to understand how we perform tasks, make sense of what is going on, act, and react.

A disadvantage of using simulations is that they require a great deal of effort and expense to set up compared to going out into the field and watching people in action. Another limitation is that simulations are inherently artificial. No matter how meticulously detailed they are, the researchers will only learn about conditions that have already been tagged as important and inserted into the scenarios. Simulations are inherently constrained to a simplified version of reality. Without validation, one can never be entirely sure that the behaviors and reactions in the simulation would also occur in a natural setting. Researchers like to say that a simulation was so close to real that the participants reported being "totally wrung out" when they finished. But participants know, just as we do, that the situation isn't the real thing. Simulations do not fully capture the stress of putting lives in jeopardy or the feeling of mental exhaustion from balancing a range of difficult tasks.

Difficulty




How close to everyday events is the data target?

The tasks that we seek to understand may be highly routine, reflecting aspects of people's work environment that they encounter every day. Observational methods are useful in these instances to understand and document the full range and extent of activities that may be involved in carrying out a set of tasks. In contrast, we may want to focus data collection on tough cases and seek opportunities to observe or interview or collect reports about situations that were particularly challenging, where people's skills and knowledge were pushed to the very edge. Incident-based methods are particularly well suited for these purposes. However, tasks that are cognitively challenging are not necessarily rare events. We may want to focus on atypical or unusual

occurrences in order to understand how people make sense of them and respond to them. Aviation, health care, and the nuclear power industry are three fields in which investigation of atypical events have been critically important for understanding accidents and errors and improving safety. An example is the extensive investigation that followed the nuclear accident at Three Mile Island in 1979, which helped create the field of cognitive systems engineering. Clearly, researchers may need to use different methods for studying unusual tasks. In cases where they may not have enough resources in the form of time, funding, and energy to wait for lightning to strike, researchers should use more productive and efficient methods such as simulations or retrospective techniques that do not depend on data collectors' being present for the atypical event.

Generality



Is the data target to gather abstract knowledge or specific events?

Some forms of knowledge elicitation center on mapping the declarative knowledge people have in a domain. In some cases, data collection is directed at simply surveying the participant's knowledge base of factual information (which supports other cognitive processes). In other cases, the researchers examine the conceptual relationships a person has formed. One example of general knowledge capture is Concept Mapping (Novak 1991), a technique for depicting core concepts and their relationships. We discuss Concept Mapping in detail in chapter 4.

Knowledge elicitation can search for general themes within a specific job or task. Concept maps can be focused at this level. Some other methods for surveying general knowledge center on the goals people have in performing a task, and the hierarchical relationships between these goals. Annett (1996) has described a Hierarchical Task Analysis (HTA) method that elicits goal hierarchies. Goal-Directed Task Analysis (GDTA) (Endsley, Bolte, and Jones 2003) is another example of elicitation methods that focus on goals and goals structure. An advantage of many general data gathering methods is that they are fairly well structured and can be performed by researchers who have not had a great deal of experience with the techniques. The disadvantage is that they may elicit the broad, surface features of the cognitive landscape rather than the deeper layers involved in resolving competing goals or carrying out cognitive functions under complex conditions.

Illuminating these aspects of cognition requires collecting data about actual events, about specific instances where people had to make sense out of the situation, and figure out what to do and how to do it. Understanding cognition in context means understanding both the cognition and the context that surrounds it. The depth and richness of detail means that data is more fine-grained and more tightly linked to specific cues and factors, goals, settings, and people's experience. It can be a significant challenge to identify general themes and overarching meanings in data at this level of specificity.

Combination of CTA Methods

So far in this chapter, we have discussed individual methods and the various ways to classify and categorize them as separate strands. However, in many CTA projects methods are used in combination. Using various tools and techniques in conjunction provides greater leverage and deeper insight. Understanding what the various methods offer and how they can work together in various data collection settings is part of developing expertise as a CTA practitioner.

For example, interviews can be conducted while a participant is performing an actual task or a contrived task or as part of a recall of a challenging task. Interviews can also cover general knowledge that is not related to any specific incident or task. Furthermore, interviews can vary from highly structured formats to totally unstructured, "think-aloud" techniques. One might use think-aloud problem solving with test case materials derived from archived interviews. One might take the probe questions used in a particular interview technique and use them while shadowing skilled performers and conducting observations in the workplace, and so on. New methods and combinations of methods appear in the research literature all the time. Knowledge elicitation clearly does not involve an easy listing of a handful of clearly delineated methods.

Knowledge elicitation is a critical step in performing CTA, but it is only the first step. Knowing what to do with data once it is in hand is entirely as important as knowing how to get it in the first place. We turn now to a discussion of the other two elements of CTA: analysis and representation of CTA data.

Data Analysis and Knowledge Representation

The analysis phase of CTA is the process of structuring data, identifying findings, and discovering meaning. Knowledge representation includes the critical tasks of displaying data, presenting findings, and communicating meaning. Methods for analyzing

and representing CTA data have not received the same level of attention that has been directed at knowledge elicitation. Many knowledge elicitation methods have analysis processes and representational formats contained within an overall methodology so that the output of the elicitation process is a particular analysis product (i.e., a representation). An obvious example is a Concept Map, which is the product associated with the elicitation and analysis process of Concept Mapping. Other examples include the "blackboard structures" produced by COGNET (Nii 1986a, 1986b; Zachary, Ryder, and Hicinbothom 1998) or the hierarchies produced by methods such as GOMS (Kieras 1988) or HTA (Annett 1996; Shepherd 2000). In much of the CTA literature, analysis and representation are inherently linked to knowledge elicitation and are not treated as separate processes at all. Instead, distinctions between analysis and representational tools and formats are embedded in comparisons of various approaches to knowledge elicitation (e.g., Cooke 1994; Hoffman 1987). We are aware of only a handful of articles or chapters that focus specifically on the analysis and/or representation phases of CTA, providing examples and comparison among tools and formats (e.g., Hoffman, Crandall, and Shadbolt 1998; Hutton et al. 1998; Militello 2001; Wong 2004).

However, many knowledge elicitation methods produce data that can be analyzed in many different ways and represented using a variety of formats. Treating CTA data analysis and representation separately from knowledge elicitation allows us to see different analysis processes, products, and representation formats more clearly. We can think about the range of possibilities available and how they might be brought together in a project to take full advantage of the CTA data.

What sorts of analysis and representation products are available to CTA practitioners? The CTA Resource website provides a catalogue of sixty different analysis tools and approaches (presented in table 2.4). Approximately one-third of the methods identified are linked to specific knowledge elicitation methods, and many are further linked to specific types of analysis processes and representations. The types of analytic products they yield include:

- Textual descriptions
- Tables, graphs, and illustrations
- Qualitative models, such as flowcharts, and
- Simulation, numerical, and symbolic models, including computer models.

Many of the methods identified in table 2.4 have predetermined analytic products. But how to proceed when this isn't the case? A challenge in data analysis comes in working with semistructured or unstructured knowledge elicitation methods. Here,

Table 2.4

Methods that produce an analytic product or representation

Applied Cognitive Task Analysis	Information Flow Analysis
ACT-R	Interaction analysis
Barrier and work safety analysis	Job analysis
Clustering routines	Laddering
COGNET	Link analysis
Cognitive Function Model	Magnitude estimation
Cognitive Work Analysis	Management Oversight Risk Tree technique
Comparing	Man-Machine Integration Design and Analysis System
Conceptual graph analysis	Multidimensional card sorting
Content analysis	Multidimensional scaling
Control task analysis	Network scaling
Correlation/covariance	Operator Model Architecture
Diagram drawing	Operational sequence analysis
Discourse/conversation/interaction analysis	Operational sequence diagrams
Discrete event simulation	Operator Function Model
Distinguishing goals	P Sort
Executive Process/Interactive Control	Paired comparison
Event trees	Process tracing/protocol analysis
Exploratory Sequential Data Analysis	Reclassification/goal decomposition
Failure models and effects analysis	Repeated sort
Fault trees	Repertory grid
Free association	SOAR
Functional Abstraction Hierarchy	Social organization and cooperation analysis
Functional Flow Analysis	Statistical modeling/policy capturing
Goals, Operators, Methods, and Selection	Strategies analysis
Graph construction	Structural analysis techniques
Grounded theory	Time line analysis
Hazard and Operability Analysis	Work domain analysis
Hierarchical sort	Worker competency analysis
Hierarchical Task Analysis	Workflow model
Identifying aspects of the representation	
Influence diagrams	

Source: www.ctaresource.com

attention that has
ion methods have
1 overall methodol-
ysis product (i.e., a
the product associ-
g. Other examples
5a, 1986b; Zachary,
ods such as GOMS
the CTA literature,
itation and are not
alysis and represen-
ous approaches to
are of only a hand-
d/or representation
and formats (e.g.,
itello 2001; Wong

can be analyzed in
Treating CTA data
allows us to see dif-
ore clearly. We can
ght be brought to-

ple to CTA practi-
erent analysis tools
the methods iden-
y are further linked
s of analytic prod-

models.

analytic products.
analysis comes in
n methods. Here,

the CTA practitioner faces the task of structuring the data, often in a series of analytic steps, to arrive at a set of findings and representational products.

There are some approaches to data analysis and representation that are useful for working with less-structured knowledge elicitation methods. Chapter 7 provides a more detailed description of different analysis and representational products.

Capsulizing Incidents

Incident-based knowledge elicitation methods, such as the Critical Decision Method, can produce voluminous data records. A single two-hour interview that is tape-recorded can run twenty to sixty transcribed pages. Even a small project can produce a lot of material to think about. One technique is to reduce incident accounts to a few pages, perhaps even to a graphic on a single page that captures the key decisions and the prominent cues. These encapsulated descriptions are easier to work with and compare than the full incident description. Narrative descriptions can also be effective as representations, because they can be created in ways that highlight cognitive content while retaining context and chronology of the event.

Cataloguing Cues and Patterns

Data records and interview notes can be examined for the cues that go into effective performance. These can be compiled by individual incident or combined across similar incidents. The cue sets can include obvious cues that novices would notice as well as subtle cues that only experts would readily detect. They can include cues that are easy to articulate as well as complex cues that require illustration. They can include relatively unitary cues as well as patterns of cues. The resulting critical cue inventory can be compiled from the notes, from transcripts, from situation awareness records, or from any other form of data. As a representation, it can convey the detail associated with specific cues, along with the pattern of a configuration of cues.

Identifying Themes

The simplest, most flexible, but most demanding approach to data analysis is to carefully review the data in search of major themes. This strategy is inductive—to work from particulars in order to discover general themes. For example, you may find the strategy for handling a particular cognitive challenge that occurs in one incident account or set of observational data being repeated in other parts of the data set, suggesting a more general finding. Key themes can be organized into a table that lists the dominant themes and cross-references them to interviews or observations or incident

a series of analytic

that are useful for
apter 7 provides a
products.

Decision Method,
view that is tape-
project can produce
it accounts to a few
e key decisions and
work with and com-
also be effective as
t cognitive content

at go into effective
bined across similar
ld notice as well as
e cues that are easy
y can include rela-
l cue inventory can
wareness records, or
he detail associated
s.

a analysis is to care-
inductive—to work
; you may find the
in one incident ac-
he data set, suggest-
table that lists the
rvations or incident

accounts. In this way you create an audit trail for the thematic analysis and develop the basis for additional analyses.

Coding the Data

Cue categories, thematic analysis, and other sorts of analysis products lend themselves to simple quantitative analysis because researchers can code the data and tabulate frequencies. For example, it might be interesting to create a frequency count of the typicality of themes or cue patterns: do you always see them, or rarely see them, or are they linked to certain conditions in the task or the environment?

Be advised that coding activities often lead to discovery of ambiguities that may lead back to the start and additional coding of the data, but that is a part of the learning process. The more explicit the coding rules are, the faster you will discover ambiguities. Because data coding can be so subjective, it can be important to share the task with another analyst (to seek replication of your ideas), or to turn it over to two or more coders who were not part of the effort to define the categories.

Describing Cognitive Sequences

In data that have a dynamic quality, where timing and sequence are an important part of events, data can be depicted to reflect the flow of cognitive activities of the actors. For example, sequences might be created to show the types of decisions made at various points in the incident, the cues that were present, the types of demands for identifying problems or categorizing situations, the types of strategies for gathering evidence, and so forth. Chronologies can provide temporal representations of events, specific cognitive processes, and/or cognitive requirements.

Summary

We have reviewed many different approaches to CTA in order to show the possibilities that exist and to provide a context for the methods we describe in part I. Another reason for this overview is to demonstrate that there is no single right way to do CTA. Practitioners of CTA have a wide range of choices in the strategy to use in knowledge elicitation, data analysis, and knowledge representation. Instead of worrying about following an official program, practitioners are better served by tracking the cognitive phenomenon they want to understand. Getting an insightful account of this phenomenon is far more important than preserving methodological rigor that might interfere with the investigation.

Cognitive Task Analysis research is often conducted as field studies, since it comprises the initial exploration of a cognitive process or strategy that is not well understood. We argue that it is misplaced rigor—rigor mortis, in fact—to let the choice of methods overshadow the phenomenon being studied. For field research, scientific values dictate that the methods you employ be documented and that your analyses be described in sufficient detail so that others can review your efforts and replicate your findings. You will also want to document evidence that runs counter to your hypotheses. These are all appropriate measures to increase scientific validity. In contrast, a rigid adherence to experimental control during a CTA study is an inappropriate attempt to mimic the psychology laboratory.

Therefore, we recommend that CTA researchers be prepared with a range of methods that they can use or adapt. Researchers have many choices in the strategy they use in knowledge elicitation, data analysis, and knowledge representation. The remainder of this book, in a sense, is aimed at providing enough information so that researchers can make those choices.