

**UNITED STATES AIR FORCE
ARMSTRONG LABORATORY**

**INFORMATION WARFARE: EVALUATION
OF OPERATOR INFORMATION
PROCESSING MODELS**

Judi See

LOGICON TECHNICAL SERVICES, INC.
P.O. BOX 317258
DAYTON OH 45431-7258

Gilbert G. Kuperman

CREW SYSTEMS DIRECTORATE
HUMAN ENGINEERING DIVISION
WRIGHT-PATTERSON AFB OH 45433-7022

OCTOBER 1997

INTERIM REPORT FOR THE PERIOD APRIL 1997 TO OCTOBER 1997

19980316 058

DTIC QUALITY INSPECTED 3

Approved for public release; distribution is unlimited.

Crew Systems Directorate
Human Engineering Division
2255 H Street
Wright-Patterson AFB OH 45433-7022

NOTICES

When US Government drawings, specifications, or other data are used for any purpose other than a definitely related Government procurement operation, the Government thereby incurs no responsibility nor any obligation whatsoever, and the fact that the Government may have formulated, furnished, or in any way supplied the said drawings, specifications, or other data, is not to be regarded by implication or otherwise, as in any manner licensing the holder or any other person or corporation, or conveying any rights or permission to manufacture, use, or sell any patented invention that may in any way be related thereto.

Please do not request copies of this report from the Armstrong Laboratory. Additional copies may be purchased from:

National Technical Information Service
5285 Port Royal Road
Springfield, Virginia 22161

Federal Government agencies and their contractors registered with the Defense Technical Information Center should direct requests for copies of this report to:

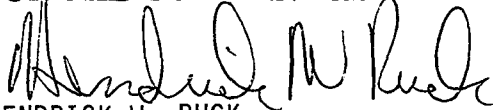
Defense Technical Information Center
8725 John J. Kingman Road, Suite 0944
Ft. Belvoir, Virginia 22060-6218

TECHNICAL REVIEW AND APPROVAL AL/CF-TR-1997-0166

This report has been reviewed by the Office of Public Affairs (PA) and is releasable to the National Technical Information Service (NTIS). At NTIS, it will be available to the general public, including foreign nations.

This technical report has been reviewed and is approved for publication.

FOR THE COMMANDER


HENDRICK W. RUCK
Chief, Human Engineering Division
Armstrong Laboratory

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE October 1997	3. REPORT TYPE AND DATES COVERED Interim Report Apr 97 - Oct 97		
4. TITLE AND SUBTITLE Information Warfare: Evaluation of Operator Information Processing Models			5. FUNDING NUMBERS C: F41624-94-D-6000 PE: 62202F PR: 7184 TA: 10 WU: 46	
6. AUTHOR(S) *Judi E. See Gilbert G. Kuperman				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) *Logicon Technical Services, Inc. P.O. Box 317258 Dayton, OH 45431-7258			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Armstrong Laboratory, Crew Systems Directorate Human Engineering Division Human Systems Center Air Force Materiel Command Wright-Patterson AFB OH 45433-7022			10. SPONSORING/MONITORING AGENCY REPORT NUMBER AL/CF-TR-1997-0166	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) The present document provides a review of 17 models of human information processing. The models were grouped into five classes: Memory and attention, artificial intelligence, visual attention, language comprehension, and situation awareness. The utility of the models was evaluated with respect to their contributions toward understanding the Observe, Orient, Decide, and Act phases of the OODA Loop, the predominant decision-making framework for the Third Wave Battlespace, or information warfare. The assessment indicated that current information processing models contribute primarily to the Observe and Orient phases; however, their contributions are disappointingly minor. Further, among the models considered in this review, effective portrayals of the Decide and Act phases were lacking altogether.				
14. SUBJECT TERMS information warfare, information processing models, OODA Loop			15. NUMBER OF PAGES . 132	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UNLIMITED	

This page intentionally left blank.

PREFACE

This effort was conducted by the Collaborative Systems Technology Branch (CFHI) of the Armstrong Laboratory at Wright-Patterson Air Force Base, Dayton, Ohio. The project was completed under Work Unit 71841046, "Crew Systems for Information Warfare." Logicon Technical Services, Inc. (LTSI), Dayton, Ohio, provided support under contract F41624-94-D-6000, Delivery Order 0007. Mr. Donald Monk was the Contract Monitor.

The authors wish to acknowledge Dr. Randall Whitaker, who provided valuable advice throughout the duration of the project, and Ms. Elisabeth Fitzhugh, who spent many hours perfecting all of the figures included in the document.

TABLE OF CONTENTS

	<i>PAGE#</i>
LIST OF FIGURES	vi
LIST OF TABLES	viii
THE THIRD WAVE BATTLESPACE	1
Three Types of Warfare	1
Information Warfare	3
The OODA Loop	4
Information Warfare: Summary and Implications	7
MODELS OF HUMAN INFORMATION PROCESSING	7
Information Theory	8
Models of Memory and Attention	12
The Modal Model of Memory	12
Baddeley's Model of Working Memory	16
Articulatory loop	18
Visuo-spatial sketchpad	20
Central executive	22
Summary	23
Controlled and Automatic Information Processing	23
Norman and Shallice's Attention to Action Model	29
Global Workspace Theory	34
Multiple Resource Theory	40
Artificial Intelligence Models	44
General Problem Solver	44
ACT-R	47
Soar	51
GOMS	57
EPIC	65
Models of Visual Attention	71
Uniform Connectedness	71
CODE Theory of Visual Attention	74
CODE theory of perceptual grouping by proximity	75
Bundesen's theory of visual attention	77

TABLE OF CONTENTS (cont'd.)

	<i>PAGE#</i>
CODE theory of visual attention (CTVA)	78
Evaluation of the CTVA	81
Models of Language Comprehension	82
Construction-Integration Model	82
Knowledge representation	83
Construction process	84
Integration	85
Applications	85
Summary	87
Sanford and Garrod's Model of Written Discourse Comprehension	87
Models of Situation Awareness	91
Adams, Tenney, and Pew's Model of Situation Awareness	91
Endsley's Model of Situation Awareness	96
Individual factors affecting SA	97
Task and system factors affecting SA	100
Errors in SA	100
CONTRIBUTIONS TO INFORMATION WARFARE: UNDERSTANDING HUMAN INFORMATION PROCESSING VIA THE MODELS	101
Rasmussen's Skills-Rules-Knowledge Framework	102
Model Evaluation	104
Models of Memory and Attention	105
AI Models	109
Models of Visual Attention	111
Language Comprehension Models	112
Models of Situation Awareness	114
Conclusions and Recommendations	115
REFERENCES	118
GLOSSARY	124

LIST OF FIGURES

<i>FIGURE#</i>	<i>TITLE</i>	<i>PAGE#</i>
1	The OODA loop (Boyd, 1987).	5
2	Correspondences between the OODA and SHOR models (SHOR descriptions adapted from Wohl, Entin, & Eterno, 1983).	6
3	The flow of information through a communication system.	10
4	The flow of information through the memory system as conceived by Atkinson and Shiffrin (1968).	14
5	A simplified representation of the working memory model (Baddeley, 1990, p. 71).	18
6	The overall system comprising the attention to action model proposed by Norman and Shallice (1986, p. 7).	32
7	A proposed dimensional structure of human processing resources (Wickens, 1992, p. 375).	42
8	The Soar problem space architecture (Newell, 1990, p. 161).	55
9	The Model Human Processor of GOMS (Card, Moran, & Newell, 1983, p. 26).	59
10	Overall structure of the EPIC architecture showing information flow paths as solid lines and mechanical control or connections as dotted lines (Kieras & Meyer, 1994, p. 3).	67
11	The representation of groups according to the CODE theory of perceptual grouping by proximity. Panel a shows a two-dimensional arrangement of five dots. Panel b shows the 2-D CODE surface for the five dots. Panel c shows the application of a threshold to the CODE surface (Logan, 1996, p. 607).	76
12	Architecture of the CODE theory of visual attention indicating the parameters and representations associated with the early identity and location system, the late identity system, and the late location system (Logan, 1996, p. 616).	79
13	Neisser's perceptual cycle. The inner circle depicts the perceptual cycle and the outer circle depicts the general exploratory cycle.	93
14	Adams, Tenney, and Pew's (1995) modification of the perceptual cycle (p. 91).	96

LIST OF FIGURES (cont'd.)

<i>FIGURE#</i>	<i>TITLE</i>	<i>PAGE#</i>
15	Model of situation awareness in dynamic decision making (Endsley, 1995, p. 35).	97
16	Mechanisms of situation awareness (Endsley, 1995, p. 41).	98

LIST OF TABLES

<i>TABLE#</i>	<i>TITLE</i>	<i>PAGE#</i>
1	Characteristics of Six Models of Memory and Attention	106
2	Characteristics of Five AI Models	109
3	Characteristics of Two Models of Visual Attention	112
4	Characteristics of Two Models of Language Comprehension	113
5	Characteristics of Two Models of Situation Awareness	114
6	Contributions of Five Classes of Cognitive Models to the Phases of the OODA Loop	116

THE THIRD WAVE BATTLESPACE

In the aftermath of the Gulf War of 1991, a great deal of attention began to be devoted to what has come to be known as the Third Wave battlespace, or information warfare (IW) (DiNardo & Hughes, 1995; Jensen, 1994; Toffler & Toffler, 1991, 1993). The Gulf War, the world's first "Third Wave" war, served to emphasize the growing importance of the role of technology in warfare. As Col. Owen Jensen (1994) points out, the clearest and most accurate account of how this new type of warfare evolved is provided by Alvin and Heidi Toffler (1993). According to the Tofflers, warfare follows wealth; i.e., the culture, technology, communication, technical skill, and organizational pattern that develop in a society and define its economy also dictate the manner in which that particular society will wage war.

Three Types of Warfare

Three basic types of warfare have evolved in human history: agrarian, industrial, and informational. Agrarian warfare predominated during the agrarian age when farming replaced hunting and gathering. People settled more or less permanently in one place, and populated towns developed. Unlike hunting and gathering, agriculture enabled communities to produce and store an economic surplus. It also expedited the development of the state. With the proliferation of these conditions, conflict first took on the true character of war as a battle between organized states. Wars were motivated by the goals of capturing additional wealth and land and were fought according to the agrarian schedule; i.e., during the intervals between planting and harvesting. Because people were needed primarily for tending the land, First Wave armies were usually formed only when needed and were not maintained throughout the year. Weapons were unstandardized and, like farm implements, were designed for hand-use. Like the manual labor in the fields, combat in the battlefield was hand-to-hand.

When the agrarian age gave way to the industrial age, agrarian warfare was replaced by industrial warfare. The economic and military climate began to change in the seventeenth century with the introduction of steam power and the manufacture of interchangeable, machined parts, launching the Second Wave of historical change. Mass production in industry was paralleled by the use of mass armies in wartime, professional full-time fighting forces paid by the state to do nothing else. The most dramatic change in industrial warfare was the introduction of

standardized weaponry, made possible by the new methods of mass production and mass distribution. The goals of warfare were to annihilate and subordinate, achieving unconditional surrender. The hallmark of Second Wave warfare was mass destruction, and World War II remains the prime example.

While some areas of the world remain in the agrarian and industrial realms, others such as the U. S. have moved unequivocally into the age of information. The pattern of life in information-age societies is controlled by information technology. People make their living by exchanging information via computers, cellular phones, and fax machines. Products are designed with computer assistance. Customized production has replaced mass production. Third Wave production relies on customization, precision, and waste or damage reduction. As the Tofflers point out, these economic changes are reflected in the nature of warfare in information-age societies. Their military forces use “smart” weapons that support precision aiming and the minimization of collateral damage. “Information warfare relies on sophisticated communication, imbedded intelligence, access to space, and real-time decision loops. It is permeated by information feeding precision weaponry, multispectral sensors providing real-time data about the battlefield, and tightly woven command and control of combined arms elements” (Jensen, 1994, p. 36). In Third Wave war, time is even more critical than in the past. Events in the Third Wave battlespace are accelerative, demanding rapid decision-making and instantaneous communication and response. The need for speed translates into an emphasis on rapid deployment, mobility, and surprise. Data from a variety of sources must be rapidly integrated and transformed into useable information. In short, Third Wave warfare is knowledge-driven, knowledge-intensive warfare waged by a knowledgeable, professional fighting force.

Many view the Gulf War as the first Third Wave war (Toffler & Toffler, 1993). In large part, the war was waged by controlling information and the enemy’s ability to gather information. Iraq’s ongoing aerial reconnaissance was suppressed, and satellite intelligence was denied. Indirect channels such as the public media were also manipulated in order to mislead Iraq into focusing on the eastern end of the Kuwaiti front. This effectively suppressed information indicating that the Allies were building up in the desert to the west of what Iraq believed to be the front lines. Further, the use of cruise missiles during Desert Storm was based on the application of precise information from terrain maps and recon photos for the delivery of precision weaponry with maximum target destruction and minimum collateral damage. As

Whitaker and Kuperman (1996) note, "the key success factors were best explained with regard to the acquisition and processing of information, the integration of this information into a base of knowledge, and the conduct of war-making activities based on this evolving knowledge" (p. 13).

Information Warfare

It should come as no surprise then that *information warfare* is the label most commonly used to refer to this emerging form of warfare. In general, the term refers to the fact that information in all its forms has become an increasingly critical component in the battlefield. More formally, IW is defined as "any action to deny, exploit, corrupt, or destroy the enemy's information and its functions; protecting ourselves against those actions; and exploiting our own military information functions" (Widnall & Fogleman, 1995, pp. 3-4). Thus, IW involves actions designed to attack, defend, or exploit information. According to Widnall & Fogleman (1995), information attack and defense may be accomplished via one of six activities: (1) psychological operations, using information to affect enemy reasoning; (2) physical destruction of enemy information systems and networks; (3) military deception, misleading the enemy about capacities and intentions; (4) information attack, direct information corruption without physical damage; (5) security measures, preventing enemy knowledge of capacities and intentions; and (6) electronic warfare, denying the enemy accurate information. Information exploitation may be accomplished through information operations: "any action involving the acquisition, transmission, storage, or transformation of information that enhances the employment of military forces" (Widnall & Fogleman, 1995, p. 11).

Within IW, two other distinctions are important. First, IW can be either offensive or defensive information warfare. Offensive information warfare tactics serve to attack or exploit the enemy's ability to gather or use information, whereas defensive information warfare measures protect our own ability to carry out information operations. Thus, at the same time that we are trying to degrade our adversary's informational capabilities, they are attempting to reciprocate. Second, IW can take the form of either information systems warfare (ISW) or information dominance warfare (IDW). ISW includes offensive and defensive actions directed at structures of command and control; i.e., the media or vehicles by which command, control, and intelligence functions are achieved. IDW, on the other hand, is information warfare aimed at

manipulating the data, information, or knowledge themselves as opposed to the channels by which they are conveyed or processed.

The term *information dominance* itself is used to refer to an operational advantage that arises as a result of superior acquisition and processing of data and information. According to the Joint Chiefs of Staff (1996), information dominance is “the capability to collect, process, and disseminate an uninterrupted flow of information while exploiting or denying an adversary’s ability to do the same” (p. 16). Unlike most others in the IW literature, Whitaker and Kuperman (1996) further specify that the informational superiority must manifest itself in instrumental superiority in order to be termed information dominance. That is, according to their definition, informational superiority in and of itself is of little value unless it is applied to our advantage. Information superiority that results in no effect (e.g., a data base of irrelevant information that has no bearing on the situation) or a negative effect (e.g., information overload that interferes with instrumental action) would not be classified as information dominance.

The OODA Loop

Finally, a discussion of IW would not be complete without mention of the most-cited theoretical construct in the IW literature: the OODA Loop (Boyd, 1987). OODA stands for the four stages of a cyclical model of the perceptual, cognitive, and enactive factors involved in the decision-making process: *Observation*, *Orientation*, *Decision*, and *Action*, as shown in Figure 1. As Whitaker and Kuperman (1996) point out, the OODA Loop is used primarily to illustrate the practical payoff of information dominance; i.e., the ability to act and react in an informed, knowledgeable manner faster than the adversary. The attainment of this temporal decision-making advantage is referred to as “operating within the enemy’s OODA Loop.” The aim is to act so as to provide the enemy with a scenario that is actually conducive to one’s own goals and deny the adversary sufficient time for assessing its validity, the options that might be available, and the potential consequences of each option.

The OODA Loop is useful in large part because it depicts the decision cycle in its entirety from perception to response. During the initial *Observation* phase, the individual engages phenomena in the environment and transforms them into data. This phase concludes when the individual begins integrating the data with his/her knowledge base. The second stage,

Orientation, occurs when the individual engages data derived from observation. Relevant information is abstracted from the stream of data and integrated with existing information to achieve at a coherent state of knowledge. Orientation concludes once this coherent state has been achieved. During the *Decision* phase, the individual engages situational knowledge from the previous stage and begins evaluating it; i.e., projecting its ramifications, focusing on a particular set of ramifications, and selecting actions appropriate for that plan. The Decision phase concludes when the individual progresses from reflection to action. Finally, in the *Action* phase, the individual begins acting on the plan derived from the previous stage. The fourth phase concludes when the action is either completed or interrupted, and the individual begins observing the altered state of the environment. Because the OODA Loop is inherently iterative, the results of the Action phase modify the individual's situation and affect his/her ongoing ability to observe.




Figure 1. The OODA Loop (Boyd, 1987).

Though it has only recently been thrust fully into the spotlight, the OODA concept is not new. In fact, the OODA model bears many similarities to the Stimulus-Hypothesis-Options-Response (SHOR) model developed by Wohl, Entin, and Eterno (1983) to describe decision tasks within command and control. Specifically, “the overall intent of this [SHOR] model is to represent the human decisionmaker as a controller working in an uncertain environment with multiple hypotheses about what is going on in the battle” (Wohl, Entin, & Eterno, 1983, p. 25).

Like the OODA model, SHOR breaks the decision-action cycle down into four parts. In SHOR, these are referred to as the Stimulus, Hypothesis, Options, and Response stages; which bear a remarkable correspondence to the Observe, Orient, Decide, and Act phases of the OODA Loop. The *Stimulus* phase of the SHOR model involves the processing of sensory data in the environment. In the *Hypothesis* phase, as in the Orient phase of the OODA Loop, the data are integrated with prior knowledge and transformed into information. Subsequently, hypotheses about the current state of the situation, given the information that has been perceived and integrated, are generated and evaluated. That is, the individual attempts to form a coherent picture of the situation. The *Options* phase consists of the generation and evaluation of potential courses of action based on the hypotheses that were generated in the previous stage. Finally, the *Response* phase, like the Act phase in the OODA Loop, involves executing the plans that were made. The similarity between the two models is depicted in Figure 2. As described by Whitaker and Kuperman (1996) and by Wohl, Entin, and Eterno (1983), models such as the OODA Loop and SHOR can be used to represent decision-making in the battlefield. Whitaker and Kuperman (1996), for example, demonstrated how the OODA Loop can be used to analyze tasks and missions involved in theater missile defense attack operations (e.g., “scud hunting”). Similarly, Wohl and his colleagues showed how the SHOR model could be used to depict a commander’s decision-making process in an antisubmarine warfare situation.

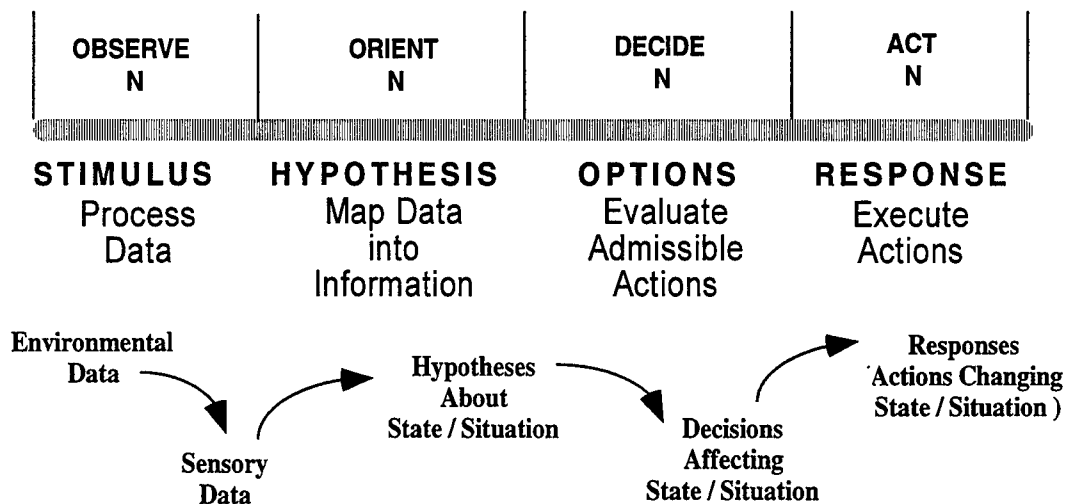


Figure 2. Correspondences between the OODA and SHOR models (SHOR descriptions adapted from Wohl, Entin, & Eterno, 1983).

Information Warfare: Summary and Implications

In essence, the focus of IW is information and how it can be used to overcome an adversary. Engaging in IW means that we must not only protect our own ability to gather, use, and disseminate information but also impede the enemy's ability to do the same. The ultimate goal is to achieve information dominance: the ability to use informational superiority in order to act and react faster than the enemy. This practical payoff of information dominance is illustrated through the OODA Loop, a cyclical model of the perceptual, cognitive, and enactive factors involved in the decision-making process. To achieve information dominance, we must be able to Observe, Orient, Decide, and Act in an informed and knowledgeable manner more quickly than the enemy.

Because information is the focus of IW, movement into the arena of the Third Wave battlespace will necessarily be accompanied by a growing emphasis on the human capacity to understand and use that information effectively. In the Third Wave battlespace, human operators will need to figure out ways to effectively degrade the enemy's information while preserving their own. They will have to process incoming information, analyze it, determine its validity, integrate it with other information, and decide what to do with it. The information that comes in may be degraded; it may be "false" information that the enemy has altered in some way; it may conflict with prior information. The human beings who must make sense of the incoming information will therefore be functioning under considerable uncertainty. Given the increased tempo of events in the Third Wave battlespace (i.e., the need to stay one step ahead of the enemy), they will also be burdened by the need to perform more quickly than ever before. The very nature of the Third Wave battlespace implies a need to discern how human operators will cope with such information-intensive tasks. That is, in order to understand human behavior in the Third Wave battlespace, we need to develop a better understanding of how humans process information.

MODELS OF HUMAN INFORMATION PROCESSING

As noted by Sanders and McCormick (1987), the study of how humans process information about their environment has been an active area of research in psychology for over a century now. This area of study rose to prominence, however, during the 1960s, in part as a

result of the information revolution and associated advances in computer technology and in part due to work being done in the communication sciences. One particularly influential force came directly from the field of communications in the form of information theory, a model concerned primarily with the ways in which information can be measured (Pierce, 1980; Shannon & Weaver, 1949). Given the growing interest in studying the human ability to process “information,” the theory received considerable attention in the field of psychology.

Information Theory

Information theory was largely originated with the publication of Claude Shannon’s two-part paper entitled “A Mathematical Theory of Communication” in the *Bell System Technical Journal* in 1948 (later reprinted in monograph form in 1949 by Shannon and Weaver). This document has come to be regarded as the foundation of modern information theory (Pierce, 1980). As the title of Shannon’s paper implies, information theory is a mathematical theory of communication. As such, it provides a universal measure of the amount of information in a message in terms of choice or uncertainty. Information theory further specifies how to determine the quantity of information that can be transmitted over both perfect and noisy communication channels. Shannon and Weaver’s theory also indicates how to measure the rate at which a message source generates information as well as how to encode messages for efficient and accurate transmission.

The beginnings of information theory itself can be traced to the study of electrical communications. In fact, some of the ideas that are critical to information theory date back to the very origins of electrical communication; i.e., work on the electrical telegraph. For example, it was noted that the spaces (the absence of an electric current), dots (an electric current of short duration), and dashes (an electric current of longer duration) comprising a telegraphic message did not always transmit precisely. Dots and dashes sent out over a submarine cable tended to spread out and overlap, losing their distinctiveness. During magnetic storms, extraneous signals tended to appear on telegraph lines and submarine cables. It was further noted that small, extraneous electric currents (i.e., noise) were invariably present in any message, interfering with the interpretability of the actual signals that were transmitted.

Further advancements in communication that ultimately influenced information theory came about during World War II. During the war, radar operators needed to be able to predict the flight paths of airplanes, using noisy and inaccurate radar data, so that the planes could be shot down. The signal representing the current position of the aircraft was a combination of desirable signals (the electric current representing data concerning the present position of the airplane) and undesirable signals (meaningless erratic currents, or noise). Researchers soon realized that, if the frequencies most strongly present in the signal differed from those most strongly present in the noise, it would be advantageous to filter out the undesirable currents by passing the signal-plus-noise through an electric circuit that would attenuate the frequencies strongly present in the noise but not those in the signal.

One of the basic concepts of information theory is that of the flow of information through a communication system, as depicted in Figure 3. According to this view, the information source selects a desired message from a set of possible messages. The encoder changes the message into some type of signal and transmits it to the decoder over the communication channel, which may or may not be subject to the effects of noise. The decoder changes the signal back into a message and transmits it to its destination. For example, in telegraphy, a written message is encoded into a sequence of interrupted currents of varying lengths (i.e., dots, dashes, and spaces) and transmitted via cable to a receiver, which decodes the signals that are received back into a written message. In oral communication, the information source is the speaker's brain, and the encoder is his/her voice. The speaker's voice produces the varying sound pressure that serves as the signal transmitted through air (the channel) to the receiver--the listener whose brain decodes the signals back into meaningful speech. Within an individual perceiver, the information source is some stimulus that is encoded by the individual's sensory receptors. The central nervous system is the communication channel that conveys the message to cortical centers for decoding. The destination is the organism's response to the message.

Another critical aspect of information theory is the concept of uncertainty and how it relates to information. As can be seen in the model of communication, the recipient of the message is at the *receiving* end of the communication process and is therefore unaware of what message the information source will choose to send. That is, the receiver possesses a certain amount of uncertainty regarding the message. First, there is uncertainty due to ignorance of what

message will be sent. This uncertainty is resolved upon receipt of the message. That is, delivery of the message reduces the receiver's uncertainty. The amount of information conveyed by a message is directly proportional to the amount of uncertainty as to what message will actually be produced. For example, a message transmitted by a system that can send one of ten different messages conveys more information than a message sent by a system that is capable of sending only one type of signal. This measure of uncertainty, or the amount of information conveyed by a message from a given source, is referred to as entropy.




Figure 3. The flow of information through a communication system.

The entropy of information theory is measured in terms of the average number of bits necessary to encode the messages produced by the source. A bit represents a binary choice, or a decision among two possibilities having equal probabilities. At the message source, a bit represents a certain amount of choice as to which message will be generated. At the destination, a bit of information resolves a certain amount of uncertainty. More specifically, the amount of information conveyed by a message is the logarithm to the base two of the number of choices that are available. Entropy increases as the number of messages among which the source may choose increases. It also increases as the uncertainty of the recipient increases.

A system characterized by low entropy or low information is one that is highly organized and possesses a low degree of randomness or choice. Conversely, high entropy implies a system with a high degree of randomness or choice. The entropy of a system can be used to obtain an estimate not only of the amount of information in a message but also of the amount of

redundancy. A ratio of the actual entropy of a source to its maximum possible entropy produces a measure of relative entropy. Redundancy is equal to one minus the relative entropy. Thus, redundancy represents that part of the message that is determined not by the free choice of the source but by the rules governing the combination of the symbols being used. For example, English text messages have considerable redundancy because of the constraints placed on letter combinations (e.g., a word beginning with the letter *j* cannot have as its second letter *b, c, d, f, g, j, k, l, q, r, t, v, w, x, or z*). Thus, the second letter cannot be freely chosen by the source; it is determined first by the rules of usage governing the English language and second by the source's intentions. Redundancy therefore implies that certain parts of the message are unnecessary and repetitive; if such parts were missing, the message would still be essentially complete.

As we have seen, the recipient has some degree of uncertainty as to which message the source will transmit. The second type of uncertainty at the destination is due to the fact that the message may have been altered during its transmission through a noisy and imperfect channel. Hence, the receiver may be uncertain as to whether the message received matches the version that was actually sent. Communication channels and circuits do not always relay information perfectly. During transmission, the signal sent by the source may fall prey to noise: undesirable additions in the form of distortions of sound (telephony), static (radio), distortions in shape or shading (television), or errors in transmission (telegraphy or facsimile). Thus, the receiver's uncertainty as to which message the sender selected may not be completely resolved even upon receipt of the message. The remaining uncertainty depends on the probability that a received symbol will be other than the symbol that was transmitted. The distinction between the two forms of uncertainty is important. Uncertainty due to freedom of choice on the part of the message source is desirable. Uncertainty that arises because of noise or errors is undesirable because it distorts the message. The uncertainty as to which symbol was transmitted when a given symbol is received provides a measure of the amount of information that was lost during transmission. The ratio of the entropy of the message to the entropy of the transmitted signal provides a measure of equivocation: the average uncertainty associated with a message when the signal is known. Any residual uncertainty that remains when the signal is known represents undesirable uncertainty due to noise. It must be subtracted out in order to derive the useful information in the received signal.

As errors in transmission become more probable, the capacity of the channel decreases; i.e., the number of bits of information that can be sent per binary digit transmitted decreases. In order to achieve a transmission with few errors, the rate of transmission must be reduced so that it is less than the channel capacity. In effect, one must add redundancy to the message by adding in sequences of unnecessary or repetitive symbols. As Pierce (1980) notes, the task of achieving efficient, error-free transmission turns out to be a problem of removing the inefficient redundancy from messages that they possess inherently and adding in redundancy of the right sort in order to allow subsequent correction of errors made during transmission.

The concepts of information, uncertainty, redundancy, and entropy that are fundamental to information theory appealed greatly to researchers in the field of psychology, who were becoming increasingly interested in determining how humans process information. In fact, information theory was introduced to psychology by Miller and Frick (1949) shortly after the publication of Shannon and Weaver's monograph. Following initial research endeavors with information theory, work devoted to the information processing approach to human cognition grew exponentially. In an attempt to understand human information processing, a number of different models have been created. In general, a model can be defined as an abstract representation of a system or process (Sanders & McCormick, 1987). A cognitive model is one that attempts to represent and describe the mental processes by which humans perform some task in order to advance our understanding of human behavior (Card, Moran, & Newell, 1983). Even more specifically, an *information processing* model of human cognition assumes that it can be subdivided into a series of stages during which certain unique operations are carried out on incoming information. The information processing model raises two critical questions: (1) What are the stages through which information is processed? and (2) How is information represented in the human mind? Basically, the question is "What is happening inside the human head to produce human cognition?" Numerous information processing models have been developed in an attempt to answer these questions.

Models of Memory and Attention

The Modal Model of Memory

One of the earliest models of human information processing was Atkinson and Shiffrin's model of human memory. During the 1960s, a major controversy in psychology centered around the

issue of whether memory consisted of a single system versus multiple systems. Some theorists adhered to a dichotomous or duplex approach, believing that short- and long-term memory involved separate underlying systems. Other theorists argued that the workings of the two types of memory reflected the operation of a single unitary system. In the ensuing years, the outcomes of a plethora of experimental and neuropsychological studies led to the conclusion that memory can be subdivided into separate short-term and long-term storage systems.

Specifically, four major pieces of evidence favored the two-system view. First, tasks such as free recall appeared to consist of separate short- and long-term components. When presented with a list of unrelated words for immediate recall in any order, individuals tend to recall the last few items particularly well. If the list is followed by a short filled delay, this *recency effect* disappears while performance on earlier items remains relatively unchanged. Thus, it seemed that the recency items were held in some type of temporary short-term store whereas the earlier items had had time to be encoded in long-term memory. Second, the short-term store appeared to have a very limited storage capacity but allowed rapid input and retrieval from storage. The long-term store, on the other hand, appeared to have virtually unlimited storage capacity but slower input and retrieval. Third, the short-term store appeared to rely on an acoustic or phonological form of encoding whereas the long-term store seemed to involve semantic coding. For example, when given a short list of words for immediate serial recall, individuals make more mistakes if the words are phonologically similar (e.g., cap, can, mad, map) rather than dissimilar (e.g., cap, late, old, big). If the list is followed by a filled delay so that long-term memory is involved, similarity of meaning becomes more important than similarity of sound. Finally, neuropsychological studies of brain-damaged individuals suggested that short- and long-term stores could be separately and differentially impaired. In some cases of amnesia, for instance, a patient might retain normal short-term memory but have difficulty retrieving long-term memories.

Although a number of different duplex models of memory were developed, the version proposed by Atkinson and Shiffrin (1968) came to be the modal or representative model of human memory. According to the Atkinson and Shiffrin model, which is depicted in Figure 4, memory is divided into three structural components: the sensory register, the short-term store, and the long-term store. Information arriving from the environment is first stored for a matter of milliseconds in visual, auditory, or haptic sensory buffer stores. These storage systems are

responsible for prolonging physical representations of briefly presented stimuli long enough to enable transfer to more durable forms of storage. Of the sensory buffer stores, the visual sensory register has been most widely studied. The results of numerous investigations have demonstrated that a highly accurate visual image persists for a very short period of time and then decays within approximately several hundred milliseconds. In addition, subsequent visual stimulation can alter or erase prior stimulation from the sensory register. A large amount of information enters the sensory register and then decays very rapidly; our senses are constantly bombarded with stimuli, some of which are relevant and many of which are irrelevant. Hence, it is up to the individual to select particular portions of the incoming information for further processing. The individual must decide which sensory register to attend to (e.g., visual, auditory, haptic) as well as where and what to scan within the system.

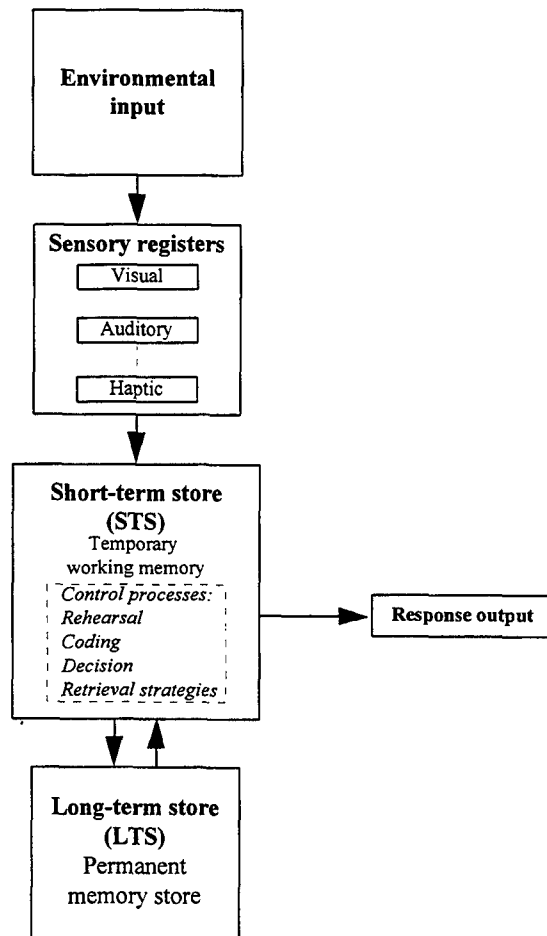


Figure 4. The flow of information through the memory system as conceived by Atkinson and Shiffrin (1968).

From the peripheral sensory stores, the attended information next enters a more durable but limited capacity short-term store that can retain information for a matter of seconds. This component of memory may be thought of as the individual's current state of consciousness. It functions as a working memory since it receives information from both the sensory register and the long-term store. In turn, the short-term store transmits information to the long-term store. Thus, the short-term store is a critical component of the model. Without it, information cannot get into or out of the long-term store. In general, information in the short-term store is lost within a period of about 15 to 30 seconds unless it is maintained by control processes that the individual may use, depending upon such factors as the task or instructions. Control processes govern such activities as rehearsal, informational flow, imagery, memory search, and output of responses. For example, in order to remember a telephone number that we have just looked up in the directory, we tend to repeat it rapidly over and over until we have completed the call. The repetition of the number serves to retain it in the short-term store. Thus, the primary purpose of rehearsal is to increase the length of time the information remains in the short-term store. A second purpose of rehearsal is to increase the strength of a trace in the long-term store, both by increasing the time the information remains in the short-term store and by allowing ample time for coding and storage processes to function. In order to remember a new telephone number on a more permanent basis, we may repeat it again and again to imbed it firmly in the long-term store. In essence, rehearsal serves to regenerate the trace in the short-term store, thereby prolonging its decay. Further studies have revealed that approximately five to nine items can be maintained in the short-term store via rehearsal. Thus, the short-term store is limited not only in duration but also in capacity.

The short-term store serves several useful functions. First, it decouples the memory system from the external environment and relieves the system from the responsibility of moment-to-moment attention to environmental changes. The sensory register, not the short-term store, is responsible for receiving environmental input from the sensory system and briefly retaining it so it can be transferred to one of the memory stores. Second, the short-term store provides a working memory in which information can be manipulated on a temporary basis. Third, it is often the primary memory device by which many tasks are completed since information can be maintained long enough to finish the task if desired.

Finally, in the Atkinson and Shiffrin (1968) model, the long-term store serves as a relatively permanent storehouse for information that has been transferred from the short-term store. Information here does not decay and become lost in the same manner as it does in the sensory register and the short-term store. It is hypothesized that information stored in the long-term store is never thereafter destroyed or eliminated. However, the ability to retrieve the information can vary considerably with time and interfering material. Hence, locating desired information in the long-term store becomes a matter of effective search and retrieval. Because long-term memory is extremely large, the search must always be made along some dimension, or on the basis of particular cues. Once the desired memory trace is located in the long-term store, it must also be successfully retrieved. If only a partial trace can be recovered, retrieval depends upon filling in the missing information by guessing or performing another search based on the partial trace that is available. Shiffrin and Atkinson (1969) proposed that the long-term store is a self-addressable memory, one in which information is stored according to the locations specified by its contents. As they point out, such a system is comparable to a library shelving system based upon the contents of books. Books are arranged according to the information they contain; e.g., history in one section, psychology in another, literature in yet another area. In order to retrieve a book, a user follows the same procedure used to store it in the first place. In terms of human memory, the information itself will define a number of areas in which it is likely to be stored; consequently, the memory search will have certain natural starting points.

In summary, Atkinson and Shiffrin (1968) viewed human memory as an organized system in which information proceeded sequentially from one structure to another. Information entered the sensory register, where it either decayed rapidly or progressed to the short-term store. In the short-term store, the information lasted for a longer duration, but again either decayed or progressed to another memory component--the long-term store. The long-term store was said to be capable of retaining information indefinitely and providing inputs to the short-term store to assist in the processing of incoming information.

Baddeley's Model of Working Memory

While the Atkinson and Shiffrin model was accepted for many years, problems soon became apparent. In particular, four major problems surfaced. First, given that the modal model proposes that the short-term store is crucial for encoding information into long-term memory, brain-damaged patients suffering from short-term memory deficits should also exhibit problems

with long-term learning. This does not appear to be the case, however. Second, the model proposes that maintaining items in short-term memory through rehearsal ensures transfer to long-term memory. Empirical evidence indicates that simple repetition does not enhance accessibility. Third, the existence of long-term recency effects and other anomalies related to the recency effects were inconsistent with the modal model's account of recency. Finally, the types of encoding in short- and long-term stores do not appear to be as clear-cut as a simple phonological/semantic split.

These inconsistencies led to the development of a number of revised approaches. Nevertheless, it is important to point out that, despite the proliferation of alternative approaches, many of the basic concepts in Atkinson and Shiffrin's (1968) original model continued to survive in one form or another. In particular, as Shiffrin (1993) notes, some model of short-term memory has been incorporated into nearly every domain of cognition. Further, three dimensions of the concept, which were espoused by the modal model in 1968, continue to be widely accepted. Namely, in every case, short-term memory is said to be characterized by temporary activation, control processes, and capacity limitations. Hence, subsequent theoretical advances have attempted not to supplant short-term memory but to further clarify its nature in greater detail.

One of these alternative approaches is Baddeley's model of working memory (Baddeley, 1986, 1990; Baddeley & Hitch, 1974). In essence, Baddeley replaced Atkinson and Shiffrin's (1968) concept of a unitary short-term store with a multi-component working memory model. His model was an attempt to account both for the evidence that fit Atkinson and Shiffrin's view of short-term memory as well as those features that were problematic. In addition, Baddeley sought to illustrate the role of working memory as a temporary storage device necessary for cognitive tasks such as reasoning, comprehension, and learning. His tripartite model proposed that working memory consists of a central executive, an articulatory loop, and a visuo-spatial sketchpad. The central executive serves as a controlling attentional system that supervises and coordinates the activities of the two slave systems: the articulatory loop and the visuo-spatial sketchpad. The articulatory loop was assumed to be responsible for the manipulation of speech-based information, whereas the visuo-spatial sketchpad was assumed to establish and manipulate visual images. The model is portrayed in Figure 5.




Figure 5. A simplified representation of the working memory model (Baddeley, 1990, p. 71).

Articulatory loop. The articulatory loop was proposed in order to account for considerable evidence indicating the importance of speech-based coding in short-term memory. This portion of working memory consists of (1) a phonological store capable of holding speech-based information for about two seconds and (2) an articulatory control process underlying subvocal rehearsal that can refresh the memory trace and feed it back into the store. The operation of the articulatory loop can best be understood by describing the manner in which it can account for a number of factors that influence memory span, including acoustic similarity, unattended speech, word length, and articulatory suppression.

The acoustic similarity effect refers to impaired immediate serial recall when items sound similar. Thus, the letter sequence *PGTVCD* will be more difficult to recall in order than *RHXKWY*. According to Baddeley's model, the acoustic similarity effect occurs because the system in which the information is stored is based on a phonological code. Items that sound similar will therefore have similar codes. Recall of the items requires discriminating among their memory traces, an act that will be harder to accomplish when the traces are similar. Thus, the level of recall will be lower than if the items are dissimilar in sound.

The unattended speech effect refers to the disruption in immediate serial recall that occurs whenever the presentation of the material is accompanied by speech sounds which are to be ignored. The effect occurs regardless of whether the unattended speech is meaningful or not.

As Baddeley points out, this phenomenon can be explained by inferring that the unattended material nevertheless gains access to the phonological store, thereby interfering with subsequent recall of attended information. Whether or not the unattended speech is meaningful is irrelevant since the phonological store retains phonological information but not semantic information (i.e., it holds the sounds not the meaning). Further study indicated that the deterioration in recall could be affected by music but not by white noise, implying that the effect is confined to speech-based sounds.

The word length effect is a phenomenon that identifies the importance of the spoken duration of words for immediate memory span. Specifically, memory span represents the number of items that can be spoken in about two seconds. Thus, word length rather than the absolute number of words is the critical factor in determining the number of items that can be recalled. Further, the essential feature appears to be the duration of the spoken word as opposed merely to the number of syllables it possesses. Baddeley has explained this effect by suggesting that the process of subvocal articulation of presented material sets up speech motor programs that run in real time, with the result that longer words take longer to run. Subvocal rehearsal is assumed to maintain items in the phonological store by refreshing the memory trace; hence, it will be able to maintain more items if it can run faster (i.e., with shorter words), thereby increasing the memory span. Thus, memory span will be directly dependent on the duration of the words to be retained.

Another effect that illustrates the functioning of the articulatory loop is articulatory suppression. Articulatory suppression refers to the disruption of the phonological loop that occurs whenever covert or overt articulation of an irrelevant is required. For example, if an individual is required to repeat the letter *A* continuously while simultaneously attempting to complete a digit span task, memory span will be reduced. This effect has been attributed to the simultaneous demands on the phonological loop. Articulation of irrelevant sounds prevents the articulatory loop from maintaining material already in the phonological store from the digit task (i.e., it prevents subvocal rehearsal and refresh of the memory trace for relevant items). In addition, irrelevant items may themselves be fed into the phonological store.

In summary, the essence of Baddeley's phonological loop hypothesis is that memory span depends upon the rate of rehearsal, which is approximately equivalent to the number of

items that can be vocalized in a two-second span. Thus, the number of items that can be recalled will depend on how long the items take to articulate. The importance of this hypothesis in everyday cognition can be seen in its role in learning to read, comprehending language, and acquiring vocabulary. First, evidence for involvement of the phonological loop in the ability to read comes from investigation of both normal and problem readers. Extensive study of children who have trouble learning how to read has revealed that such youngsters generally have an impaired memory span as well as difficulty with tasks involving some form of phonological manipulation. For example, they find it harder than their reading contemporaries to judge whether or not two words rhyme. On the flip side, once children do learn how to read, they tend to have both an enhanced memory span and phonological awareness. Second, the articulatory loop seems to be critical for speech production and comprehension. Specifically, the articulatory loop appears to be important for holding words in memory during sentence processing. In order to understand the second half of a sentence, one must be able to retain the order of the first half. Particularly telling evidence for involvement of the articulatory loop and memory span in language comprehension comes from the case of a patient who suffered memory problems following an epileptic seizure. He could understand short sentences but became lost with wordy discourse. Thus, because of an impaired memory span, he could retain only a few words at a time. Often, he could grasp only the first phrase or two of a conversation. Finally, the articulatory loop has been shown to be involved in vocabulary acquisition. As vocabulary size increases, children perform better on a variety of phonological tasks.

Visuo-spatial sketchpad. The second slave system in Baddeley's model of working memory is the visuo-spatial sketchpad responsible for setting up and manipulating visuo-spatial images. This system has not been studied as extensively as the articulatory loop. Evidence to date indicates that the visuo-spatial sketchpad can be fed directly through visual perception or indirectly through the generation of visual imagery. The system appears to be used in setting up and using visual imagery mnemonics, or memory aids. For example, in one study, individuals were asked to learn a list of ten words by associating each item with a particular location on a university campus. To retrieve the items, they imagined themselves walking through the campus from one location to the next, scanning each location for the item to be recalled. This mnemonic enhanced performance unless individuals were also required to perform a spatial tracking task during recall. The pursuit rotor tracking task required the individual to keep a stylus in constant contact with a spot of light that followed a circular path. The visuo-spatial nature of this task

interfered with individuals' ability to visualize the campus locations and therefore reduced the effectiveness of the imagery mnemonic. That is, because they could not visualize the locations, individuals could not recall the items associated with the locations. Presumably, both activities were competing for the visuo-spatial sketchpad.

Although it is used in visual imagery mnemonics, the visuo-spatial sketchpad does not appear to be responsible for the effects of imageability on long-term memory. It has been shown that concrete and readily imageable word pairs such as *house-blue* lead to greater recall than abstract words such as *truth-gratitude*. If the visuo-spatial sketchpad is necessary for setting up the image for imageable word pairs, then a concurrent visuo-spatial task such as tracking should interfere with the process. On the contrary, tracking was shown not to disrupt the recall of imageable pairs more than abstract pairs. Thus, the visuo-spatial does not appear to have a role in the facilitating effect of imagery in long-term memory.

As suggested by its name, the visuo-spatial sketchpad appears to possess both visual and spatial characteristics. In an effort to determine whether the sketchpad is primarily visual or spatial, an imagery task was combined with a task that was visual but not spatial (brightness judgment) or with a task that was spatial but not visual (an ingenious auditory tracking task that required blindfolded participants to maintain a beam of light on a swinging pendulum that emitted a particular sound only when illuminated). The results of this study indicated that the tracking task interfered with the imagery task more than the brightness judgment task, implying that the sketchpad is primarily spatial. Additional investigations, however, revealed that the visuo-spatial sketchpad also possesses visual characteristics. Namely, the unattended picture effect indicated that the presentation of unattended color patches (which are visual but not spatial) interferes with subsequent recall when the material was learned through the use of visual imagery techniques as opposed to verbal rehearsal.

Unlike the articulatory loop, the role of the visuo-spatial sketchpad in everyday cognition has not been widely explored. It does appear to be important for geographical orientation and for planning spatial tasks. With respect to spatial planning, for example, it has been shown that abacus experts who are able to perform mathematical calculations using only a mental representation of the device rely on a visuo-spatial representation held in the sketchpad. These experts had extremely long digit spans on the order of fourteen to sixteen digits but normal

memory spans for letters and other items. When asked to perform either a digit or letter span task concurrently with either a verbal or visual task, the experts performed most poorly on the digit span task when it was coupled with a visual task. These outcomes suggested that the abacus experts were using a visuo-spatial system such as the sketchpad to remember the digits; i.e., that they were mentally scanning an abacus.

Central executive. Although relatively little is known about the visuo-spatial sketchpad in comparison to the articulatory loop, even less is known about the central executive. According to Baddeley (1990), this component of working memory “has tended to become something of a ragbag for consigning such important but difficult problems as how information from the various slave systems is combined, and how strategies are selected and operated” (p. 117). Baddeley further points out that the central executive often functions more like an attentional system than a memory store. In an effort to describe the functioning of the central executive, Baddeley has drawn upon Norman and Shallice’s (1986) attention to action model of attentional control (described in detail in a later section of this document).

According to the attention to action model, ongoing actions can be controlled in one of two ways. First, in the case of well-learned skills, prior learning allows the activity to executive fairly automatically. For example, most people are quite capable of driving without actively attending to each maneuver. Such relatively automated skills can generally be completed concurrently with other activities with little or no interference. Second, ongoing actions can be controlled via the supervisory activating system (SAS). The SAS is capable of interrupting and modifying ongoing behavior. It is assumed to do so by systematically biasing existing probabilities to make one course of action more likely than another. Thus, the SAS is responsible for conscious selection of actions.

Using Norman and Shallice’s (1986) model as a guide, Baddeley proposed that the central executive of working memory is responsible for the selection, initiation, and termination of processing routines such as encoding, storing, and retrieving. Some evidence for the existence of the central executive comes from studies of reading comprehension in children. Work here has shown that the crucial difference between low and high scorers on tests of reading comprehension is working memory capacity; however, it is not due to either the articulatory loop or the visuo-spatial sketchpad. For example, high comprehenders are not better at verbatim

memory (i.e., recognizing whether a particular sentence occurred in a previously read passage) than low comprehenders, but they are better at making inferences from what they read. Further, both high and low comprehenders exhibit the typical word-length effect when remembering words and when remembering picture names that vary in length. Both of these outcomes suggest that the two groups are using the articulatory loop in the normal fashion. Further, there is little reason to believe that the difference stems from the visuo-spatial sketchpad given the nature of reading comprehension tasks. Thus, it is assumed that the differences exhibited by high and low comprehenders are due to the one remaining component of working memory, the central executive.

Summary. As noted earlier, many of the models of human memory that were developed as alternatives to Atkinson and Shiffrin's (1968) modal model retained its essential features but sought to modify them so as to conform more closely to empirical findings and explain some of the anomalies in the literature. Baddeley's model of working memory is a prime example of just such an alternative approach. Rather than eliminate Atkinson and Shiffrin's concept of short-term memory, he analyzed it into three separate components: the articulatory loop, the visuo-spatial sketchpad, and the central executive. As a whole, his model retains the basic characteristics of a short-term or working memory: temporary activation, control processes, and capacity limitations. The articulatory loop and the visuo-spatial sketchpad are limited in both duration and capacity. The central executive is the component of working memory that handles the control processes governing the activities of the two slave systems. Baddeley's major contribution was to provide a more detailed depiction of working memory and its potential subsystems.

Controlled and Automatic Information Processing

Yet another alternative to the modal model of memory is Shiffrin and Schneider's (1977) theory of controlled and automatic human information processing. It marks a rather pronounced shift away from the views that dominated the 1960s and 1970s. In particular, in the Shiffrin-Schneider model, short-term and long-term memory are no longer seen as separate storage systems. Instead, the short-term store represents that part of the long-term store that is currently activated. The short-term store serves as a "workspace" for decision-making, thinking, and control processes. The model was developed on the basis of a series of laboratory experiments

which suggested that information processing could be subdivided into two different modes: controlled and automatic.

The basic procedure used in the initial set of studies involved asking participants to remember a set of target items and then determine whether any of the items was present in each of a series of stimulus presentations. During each stimulus presentation, four elements appeared simultaneously in a square for a brief period of time. The presentation of 20 such frames in immediate succession constituted a trial. The elements comprising the frames were either characters or random dot masks. Prior to each trial, participants were presented with a varying number of items called the memory set. During the trial, they were required to detect any items from the memory set that appeared in subsequent frames. The independent variables that were manipulated in the first set of studies included frame size, memory set size, and type of mapping. Both frame size (the number of characters presented in each frame) and memory set size (the number of characters to be remembered and detected) varied from one to four. The most important manipulation was the mapping, or the relation of memory-set items to distracters. In the consistent mapping (CM) procedure, memory-set items and distracters came from distinct sets so that memory-set items were never distracters (and vice versa). Further, memory-set items were from one category (e.g., digits), whereas distracters were from another (e.g., consonants). In the varied mapping (VM) procedure, memory-set items and distracters were randomly intermixed over trials and were from a single category. Thus, on VM trials, a memory-set item on one trial might later be a distracter on a subsequent trial.

The initial studies were analyzed in terms of either the accuracy of the detection response or the reaction time. The results indicated substantial differences in the VM and CM conditions, suggesting that qualitatively different processes were operating in the two conditions. Specifically, VM conditions were exceedingly difficult and were degraded by task load (i.e., increases in frame size or memory-set size). By comparison, CM conditions were relatively easy and were virtually unaffected by load. Shiffrin and Schneider (1977) hypothesized that the observed differences were due to the consistency of the mapping over trials of the memory-set items and distracters to responses. In the VM conditions, the mapping was not consistent. A memory-set item requiring a response on one trial might later become a distracter requiring no response. Hence, Shiffrin and Schneider suggested that a controlled, serial search would be required on these trials, making them more difficult and more time-consuming than CM trials.

Additionally, they argued that the consistent mapping on CM trials led to the development of automatic detection, which enabled automatic-attention responses to become associated with memory-set items. As a result of the automatic-attention response, the controlled serial search was not required. Instead, observers could operate via a parallel detection process unaffected by task load.

Shiffrin and Schneider's (1977) initial studies were augmented with additional experiments designed to clarify their findings and to test their hypotheses further. Four experiments were conducted using variations of the procedure from the initial studies. Experiment 1 examined the development of automatic processing. Observers were required to detect targets from one set of arbitrary letters and to ignore distracters from a second set of letters over the course of several thousand trials. During this period, controlled processing gave way to the development of automatic detection, which led to considerable improvement in performance. At the start of the experiment, performance was poor since automatic detection had not been learned and controlled search had to be used. All observers reported extensive, attention-demanding rehearsal of the memory set during about the first 600 trials; but they subsequently became unaware of rehearsal or other controlled processing. When the two sets of letters were reversed so that distracters became targets and vice versa, performance dropped drastically and recovered only gradually. Immediately after the reversal, the hit rate fell to a level far below that observed at the very start of the experiment. It did not reach 90% until 2100 trials had been completed, a level that was attained after only 900 trials of original training. These outcomes suggested that the automatic-attention response is a long-term phenomenon that is highly resistant to change. It can eventually be "unlearned," but only after considerable amounts of retraining.

Experiment 2 duplicated the first experiment with the exception of using pre-experimentally categorized memory and distracter sets, either digits or consonants. That is, observers who initially searched for digits in a background of consonant distracters were subsequently switched to searching for consonants in a background of digits, and vice versa. Following the reversal in a VM condition, performance was identical to normal VM controlled search. In the CM condition, however, the reversal of memory and distracter sets reduced performance to near-chance levels. This particular result was important in demonstrating that the phenomenon of automatic detection is not due to categorical differences between the memory

and distracter sets since the digit/consonant distinction still existed following the reversal. Rather, it is the consistency of the mapping that is the key to automatic detection. When the mapping is no longer consistent, automatic processing must give way to controlled search. The results further indicated that observers in the CM condition were using a different type of controlled search after the reversal, one that compared the category of each input to the memory-set category. Thus, categories can facilitate performance, but only by assisting controlled search and not automatic search.

Experiment 3 clarified the role of categories in controlled search and in the development of automatic detection. One experimental condition was designed to generate controlled search in a situation where categorization could not develop. This was achieved by using eight consonants from which memory-set items and distracters were randomly drawn on each trial. The second condition also used controlled search, but in a situation where categorization of the memory sets was possible. This was achieved by using eight consonants subdivided into two sets of four visually confusable letters that remained disjoint throughout the experiment. What varied from trial to trial was which particular set constituted the memory-set and which the distracter-set. Thus, in both conditions, a varied mapping procedure was used to prevent automatic detection. The key difference was the fact that the consonants in the second condition were never intermixed and could eventually become well-learned. In both conditions, the VM trials were followed by CM trials, which permitted an examination of the course of the development of automatic detection when a categorization was or was not present at the start of the CM trials. The results indicated that category learning for arbitrary, visually confusable sets of letters occurred in VM conditions after 25 sessions of training. Controlled search was still necessary at this point, but the comparison had switched from individual items to categories. When training was switched to a CM condition, performance improved considerably and automatic detection was learned. These outcomes implied that the role of categories is not only to improve controlled search but also to enhance the speed of acquisition of automatic detection.

Finally, Experiment 4 tested observers' ability to focus attention on particular portions of the display despite distractions from (1) neutral characters, (2) current targets in to-be-ignored portions of the display, or (3) items in to-be-ignored locations that observers had previously been trained to respond to with automatic-attention responses. Experiment 4 demonstrated that controlled search can be directed to locations that the observer intends to search but that

automatic-attention responses can cause attention to be allocated to positions that should be ignored. That is, once observers have been trained to respond to CM targets with an automatic-attention response, the targets are virtually impossible to ignore, even when they occur in locations of the display that observers have been explicitly instructed to ignore. The automatic-attention response interrupts and redirects ongoing controlled processing.

Their experimental findings led Shiffrin and Schneider (1977) to develop a theory of memory based on the observed differences in controlled and automatic processing. In their theory, memory is viewed as a large and permanent collection of nodes that can become increasingly interrelated through learning. An individual node in memory often consists of a complex set of elements, including associative connections to other nodes, programs for responses or other actions, and directions for other forms of information processing. A node can be distinguished from other nodes because it is unitized; i.e., when any one of the node's elements is activated, all of them are activated. The nodes comprising the memory stores are said to be arranged in levels, implying that certain nodes may activate other nodes, but not vice versa. Most of the nodes in memory are normally inactive and comprise what is referred to as a long-term store. The long-term store is a permanent, passive repository for information. The set of currently activated nodes is termed the short-term store. The short-term store is temporary since information is lost or forgotten once it reverts to an inactive state. In addition to providing a temporary storage area for current information, the short-term store also provides a work space for decision-making, thinking, and control processes in general. Control of the information processing system is accomplished by manipulating the flow of information into and out of the short-term store. Such control processes include decision-making, rehearsal, coding, and searching both the long- and short-term stores for information.

Within this view, an automatic process can be defined as a sequence of nodes that is nearly always activated in response to some input without the need for active control or attention. Because automatic processes rely upon a relatively permanent set of associative connections in the long-term store, they will require considerable training to develop and will be difficult to suppress or alter once learned. A controlled process, on the other hand, uses a temporary sequence of nodes activated under the observer's control and attention. The sequence is temporary since each and every activation requires the observer's attention. Further, because the observer's active attention is needed, only one such sequence at a time can be controlled without

interference. Controlled processes are therefore highly limited by the capacity of the short-term store, but they are also easy to establish and alter. They can also be applied in novel situations where automatic sequences have not yet been learned.

Learning itself occurs via the transfer of information from the short-term store to the long-term store. During transfer, information not previously present in the long-term store is formed. This process occurs by associating the new information with information structures already present in the long-term store. Thus, transfer implies the formation of new associations between nodes that have not previously been associated in the long-term store. Most new associative structures will include the context in the short-term store at the time of the transfer as one of their components. According to Shiffrin and Schneider (1977), storage of new information is achieved through attention and controlled processing, including rote rehearsal, maintenance rehearsal, and coding rehearsal. Some degree of attention or controlled processing is a prerequisite for storage. Thus, controlled processing will underlie the development of automatic processing, a phenomenon observed in many of Shiffrin and Schneider's (1977) studies where automatic detection developed following considerable amounts of consistent training.

Given these theoretical underpinnings, the nature of controlled and automatic processed can now be described in more detail. Controlled search and detection is highly demanding of attentional capacity. The limitations of control processes are based on those of the short-term store (e.g., the limited amount of information that can be maintained without loss). Because these limitations prevent multiple control processes from occurring concurrently, these processes often consist of strings of single controlled operations. Thus, controlled processing is serial in nature with a limited comparison rate. It is also heavily dependent on cognitive load. Control processes are easily established and altered without excessive training. Control processes can be used to control the flow of information within and between levels and between the short- and long-term stores. Finally, control processes exhibit a rapid development of asymptotic performance; i.e., performance levels stabilize very quickly at an asymptotic value in the absence of automatic processing.

In contrast to controlled search and detection, automatic search and detection is relatively well-learned in long-term memory; hence, automatic processes are not hindered by the

capacity limitations of the short-term store and do not require attention for their completion. As a consequence, they are also virtually unaffected by cognitive load. Automatic processes require considerable training to develop and are exceedingly difficult to alter or suppress once learned. Performance improves only gradually as the automatic sequence is learned; asymptotic performance levels may not be reached for thousands of trials. According to Shiffrin and Schneider (1977), three factors contribute to the process of automatic detection. First, there is an automatic-attention response to the encoded features from the input target. Second, an automatic "target" response is learned that tells the observer when a target is present among the inputs. Third, in some instances, an automatic overt motor response is learned in response to a target. Of these factors, the first two were the primary contributors to automatic detection in Shiffrin and Schneider's (1977) studies.

As Shiffrin and Schneider (1977) point out, a system based on two different processing modes is advantageous in many respects. In novel situations or in situations requiring moment-to-moment decision-making, controlled processing may be used to perform accurately (albeit slowly). As the situation becomes familiar, automatic processing will gradually develop--demands on attention will be reduced, performance will improve, other activities can be completed in parallel. This type of system enables the individual to make efficient use of its limited-capacity processing system. Once automatic processing develops, the short-term store can be devoted to new tasks. Thus, even though some activities may have become automatic, the system still allows the individual to deal with novel situations for which automatic sequences have not yet been learned by means of controlled processing.

Norman and Shallice's Attention to Action Model

With their attention to action model, Norman and Shallice (1986) have attempted to clarify the role of attention in the control of action. As they point out, the role of attention in perception has been widely examined, but its purpose in the control of action has not. Norman and Shallice's model represents an attempt to account for the role of attention in action, both when performance is automatic and when it is under deliberate conscious control. Their model is organized around the notion that a set of active schemas awaits the appropriate set of conditions so that they can be selected for the control of action. Their concern is primarily with observable actions; however, as they note, the same principles apply to internal actions of cognitive processing. They further point out that their examination of deliberate versus automatic action is complementary in many

ways to Shiffrin and Schneider's (1977) examination of controlled and automatic detection. Some actions can be performed without the need for active, directed attention; others require deliberate conscious control. Tasks requiring deliberate attentional resources include: (1) planning and decision-making, (2) troubleshooting, (3) ill-learned or novel tasks, (4) dangerous or technically difficult work, and (5) actions requiring the suppression of a strong habitual response or intense resistance to temptation.

In order to account for the control of both automatic actions and those actions requiring deliberate conscious control, Norman and Shallice proposed that two complementary processes operate in the selection and control of action. One is sufficient for relatively simple or well-learned actions that can be completed automatically. The other process permits conscious, attentional control to modulate performance. A simple, well-learned action sequence can be represented by a set of schemas. When triggered by relevant perceptual input, the set of schemas results in the selection of the appropriate body, limb, or finger movements. The representation of the simple, well-learned action sequence by means of action schemas constitutes what Norman and Shallice refer to as a "horizontal thread." More specifically, the horizontal thread refers to an autonomous, self-sufficient strand of processing structures and procedures that can complete required activities without the need for conscious or attentional control. These structures underlie the performance of well-learned, habitual tasks. The sequence itself can often be depicted by a relatively linear flow of information among the various psychological processing structures; hence, the name "horizontal thread."

According to Norman and Shallice (1986), the individual schemas of the horizontal threads each have an activation value. A particular schema is selected whenever its activation level exceeds a given threshold. Once selected it continues to operate until it has satisfied its goal or completed its operations, unless it is actively switched off or blocked by the absence of some critical resource or information that is currently being used by another more highly activated schema. If numerous schemas have been activated simultaneously by a given perceptual input, the one that is most highly activated is selected. This procedure ensures that scheduling of actions is simple and direct, as need be for routine or habitual activities. No direct attentional control of selection is required.

The basic mechanism for avoiding conflicts in performance is contention scheduling. It allows simultaneous action of cooperative acts and prevents concurrent action of conflicting ones. Thus, contention scheduling resolves competition for selection by preventing competitive use of common or related structures and negotiating cooperative, shared use of common structures or operations when possible. It acts through activation and inhibition of supporting and conflicting schemas. First, the sets of potential schemas compete with one another in the determination of their activation value. Activation value is determined in part by the degree to which the existing environmental conditions match the trigger specifications for a given schema. Second, the selection takes place on the basis of activation value alone. Schemas that require the use of any common processing structures will inhibit one another. Schemas that rely on one another for the completion of a given activity will activate one another. In particular, any well-learned action sequence is represented by a set of schemas, one of which serves as the source schema. When the source schema is activated, the others in the set will be as well. Thus, activation can be determined by influences from contention scheduling, from the satisfaction of trigger conditions, and from the selection of other schemas.

In addition, a schema's activation level can come from "vertical thread" influences. The vertical thread represents an additional control structure required for novel or complex tasks that do not have schemas available for their control. This additional system is the Supervisory Attentional System (SAS). It provides one source of control upon the selection of schemas, but it operates entirely through the application of extra activation and inhibition to schemas in order to bias their selection by the contention scheduling mechanisms (i.e., in order to make it more or less probable that they will be selected). Thus, when attention to the task is required, it can increase the activation values for desired schemas and decrease the values for undesired schemas. Motivation functions in a similar manner, but more slowly, operating over longer periods of time. The overall system proposed by Norman and Shallice (1986) is depicted in Figure 6. The basic premise of the model is that two levels of control are possible for well-learned action sequences: deliberate conscious control and automatic contention scheduling of the horizontal threads.

It is important to point out that attentional control, when needed, is used only to bias the selection of particular schemas. In general, attentional control is too slow to provide the high precision of accuracy and timing needed to perform skilled acts. Deliberate conscious control generally involves serial processing steps, each of which requires 100 ms or more. Thus,

conscious control would be too slow to account for skilled human behavior that requires action sequences to be initiated exactly when conditions call for them (in some cases they must be accurate to the nearest 20 ms). Evidence that conscious attentional control is not necessary for the initiation or execution of action sequences comes from the study of slips of action--the intrusion of unwanted behavior, often during the performance of routine tasks. One class of errors known as capture errors refers to the performance of some action without conscious control or knowledge. For example, while going to the garage to get his car for work, one individual stopped to put on his boots and coat as if going to work in the garden behind the garage.

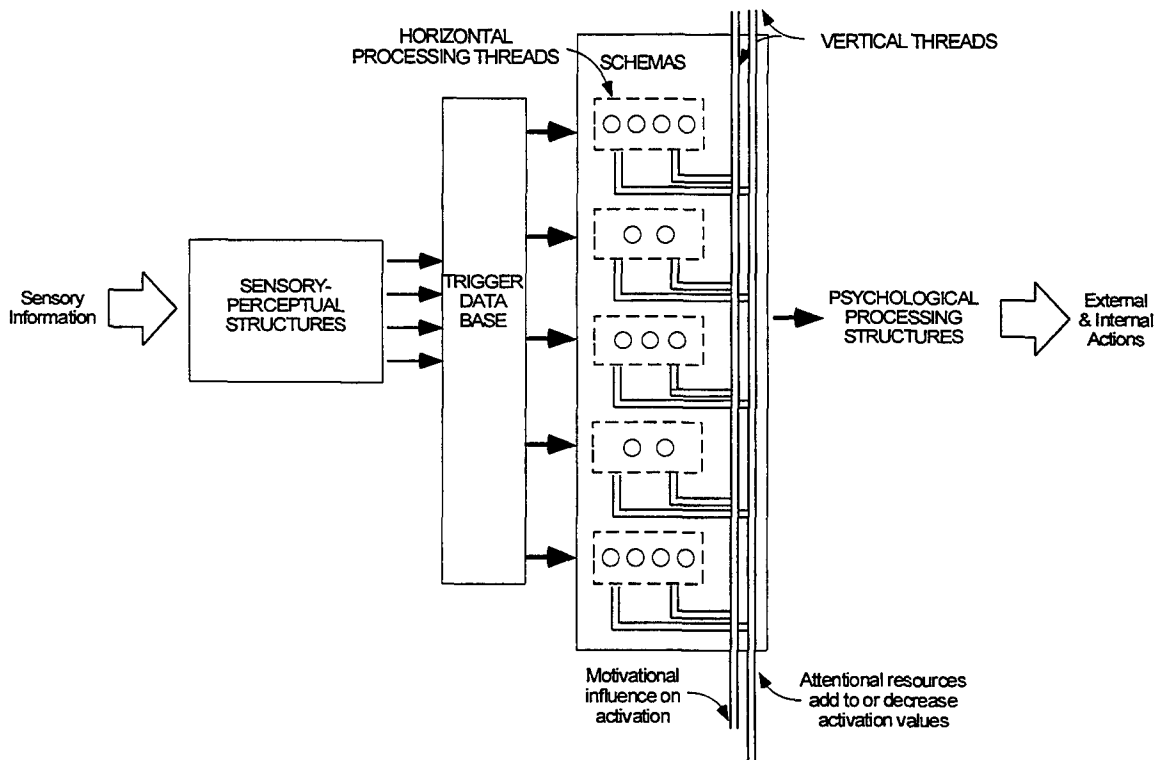


Figure 6. The overall system comprising the attention to action model proposed by Norman and Shallice (1986, p. 7).

Capture errors and the like can easily be explained by the attention to action model. If a routine task is being completed and does not require continuous monitoring and activation from the SAS, its component schemas can be selected using contention scheduling alone. Hence, the SAS can be directed toward activating some other non-competing schema. Normally, the

component schemas in the routine action would still be satisfactorily selected by contention scheduling alone. Occasionally, however, while one is "thinking about something else," a schema that controls an incorrect action can be more strongly activated than the correct schema and be executed. Because the SAS is directed elsewhere, it would not immediately catch the mistake, and a capture error would result.

As Norman and Shallice (1986) point out, perhaps the strongest evidence for their model comes from neuropsychology. Namely, the attention to action model has also been able to account for a particular type of brain damage known as frontal lobe syndrome, which is characterized by disturbed attention, distractibility, and difficulty not only with mastering new tasks but also with planning, organizing, and controlling action. The model would imply that such patients suffer from a deficit to the supervisory attentional system. There is considerable evidence to support this contention. For example, frontal lobe patients tend to perseverate when completing a task; i.e., they tend to become fixed in a certain routine and find it difficult to break out. When asked to draw circles, they may readily comply; however, when asked to switch from circles to squares, they continue drawing circles. According to the model, the initial activity or strategy continues to run because the patients have lost the capacity to interrupt and change ongoing activity as a result of impairment to the SAS. As another example, such individuals have considerable difficulty with verbal fluency tasks that require the production of words fitting a particular category (e.g., words beginning with the letter *B* or words belonging to the category of furniture). Presumably, this task is difficult because there is no standard overlearned program for generating sequences of items from a category that can be executed. The SAS model can also account for the apparent contradiction between increased perseveration and increased distractibility. If behavior is left under the control of the horizontal threads and one schema is more strongly activated than the others, it will be difficult to prevent it from controlling behavior. On the other hand, when several schemas have similar activations, it will be difficult to select among them since the vertical thread has been impaired. The end result in this case will be increased distractibility.

In summary, Norman and Shallice (1986) have specified a control system that not only allows for the relative autonomy of well-learned actions but also acknowledges that most of our actions still go according to plan. Their model can account not only for correct performance but also for the errors and slips of action that can occur. They achieve this by providing two types of

control structure: (1) horizontal threads, each of which comprises a self-sufficient strand of specialized processing structures called schemas; and (2) vertical threads, which interact with the horizontal threads to provide the means by which attentional or motivational factors can modulate the schema activation values. Horizontal threads control habitual activities without the need for moment-to-moment attentional control, receiving their triggering conditions from environmental input or from previously active schemas. Higher level attentional processes come into play via the vertical threads in novel or critical conditions when currently active schemas are insufficient to achieve the goal. They augment or decrease schema activation levels in order to modify ongoing action. Motivational variables can also influence schema activation along the vertical threads, but they are assumed to function over longer time periods than the attentional resources.

Global Workspace Theory

The Global Workspace Theory represents an attempt to provide a unified theoretical approach to explain a large set of phenomena associated with the cognitive processes that occur when we become conscious of something (Baars, 1983). According to the Global Workspace Theory, the nervous system can be described as a parallel distributed information processing system in which specialized processors perform complex and efficient processing more or less autonomously. Each processor is capable of performing some task on the symbolic representations that it receives as input. The processors are also able to combine to form new processors capable of performing novel tasks. Unlike many other models of information processing, the Global Workspace Theory does not require the presence of a central executive to control the functioning of the processors; rather, they are able to decide what should be processed by their own criteria. However, the processors do require some mechanism for information exchange in order to interact with one another. This central interchange comprises the global data base, which resembles a form of short-term or limited capacity working memory. In this model, consciousness is the result of the global workspace in the brain distributing information to the vast number of parallel unconscious processors comprising the rest of the brain. More specifically, conscious contents are said to reflect a special operating mode of the global data base, one in which there is a stable and coherent global representation that provides information to the entire nervous system.

Baars' (1983) theory was developed largely on the basis of a contrastive analysis comparing conscious versus unconscious processes across numerous domains. These contrasts can be subdivided into capability constraints, which contrast the abilities of conscious versus unconscious processes, and boundary constraints, which demonstrate the limits of our experience of some conscious content. The capability constraints include three types. First, conscious processes are computationally inefficient in comparison to unconscious processes, which have the ability to respond quickly and without error. For example, when completing a novel task that we have not yet learned, our performance is usually laborious and error-prone. We often have to mentally rehearse the steps involved in completing the task as we perform them, talking ourselves through each one. Once the task has become well-learned, we can complete it routinely or automatically, whereupon our performance becomes smooth, rapid, and virtually error-free. Thus, as conscious processes become more proficient, they also become less consciously available. We can no longer verbalize how we completed them or what processes were involved.

Second, conscious processes have great range and relational capacity, whereas unconscious processes are limited in both domain and autonomy. A huge variety of phenomena can be experienced consciously. In fact, conscious processes seem to participate in all known mental processes at some time or another. Further, conscious contents can be related to one another almost without limit. The best example of this relational capacity is classical conditioning, where virtually any stimulus is able to serve as a signal for virtually any other event. In contrast, unconscious processors by themselves are relatively limited and autonomous. For example, the visual pathway is essentially limited to processing information from the retina and little else. The autonomy of unconscious processors can be observed in the case of slips of speech or action--involuntary phenomena that violate conscious control. They occur unintentionally without the individual's awareness and are seemingly unrelated to other ongoing activities. Thus, the slip is a surprising act that is inconsistent with the individual's intentions. The individual becomes aware of the slip only when the action is related to its proper context and recognized as an error. The autonomy of unconscious processes can further be seen in the inability to control undesirable habits. They tend to creep into our everyday behavior despite the best of intentions.

Third, conscious processes are characterized by unity, seriality, and limited capacity. Unconscious processes, on the other hand, are highly diverse and can operate in parallel with seemingly unlimited capacity. The unity of conscious processes can be seen in our inability to experience two mutually exclusive organizations of input simultaneously. For example, we cannot simultaneously see both images in an ambiguous stimulus. We see first one, then the other, but never both at once. Further, we are able to have only one meaning of a word in mind at a time. Even though we know that alternative meanings exist, they remain unconscious in the current context. Only one process can be conscious at any given time. Hence, conscious experience is not only unified but also serial. Because only one conscious process can occur at a time, it follows that the system must necessarily be characterized by a limited capacity.

In addition to the capability constraints, the functioning of conscious and unconscious processes can be understood by examining their boundary constraints. Boundary constraints demonstrate under what conditions conscious events become unconscious, and vice versa. Two types of boundary conditions are called synchronic constraints since they occur concurrently with a conscious experience but are not themselves conscious. First, there must be some internal representation of the context within which a percept occurs, but this contextual representation is not itself conscious. Contextual assumptions are used to make sense of the world, but we are not aware of them. For example, we do not realize that we interpret trapezoids as rectangles inside a building until we encounter the oddities of the Ames trapezoidal room. Second, sensory input that is not interpretable within the present context is also not conscious. For example, when listening to a speaker, we are normally conscious of the meaning of a particular word, given the context in which the speaker uses it. Generally, we remain unconscious of the many other meanings of that particular word (e.g., "bank"). Thus, context by itself is unconscious; and input by itself without a context is unconscious. Consciousness arises only when some input occurs in a context.

Two other types of boundary constraints are diachronic since they occur before and after a conscious representation. First, preperceptual processes are not conscious. Visual input is preprocessed for less than a second before it becomes conscious. During this time, various hypotheses are brought to bear on the problem of representing the input. As Baars (1983) points out, preperceptual processes involve a set of hypotheses that are undefined within the current context because they are unstable and mutually competitive. Because input without context is

itself unconscious, it is not surprising that preprocessing of input is not conscious. By the time the processors cooperate to establish a coherent context, they become conscious. Second, conscious percepts habituate rapidly if the input remains predictable. When some stimulus is repeated or continued past a certain point, it is no longer experienced. Thus, the processing that occurs before stimulus presentation and after prolonged stimulation is unconscious.

Baars' (1983) Global Workspace Theory is an attempt to explain these differences between conscious and unconscious processes. In order to do so, Baars drew upon a popular model in artificial intelligence distributed-processing systems in which a globally accessible block of working memory orchestrates communication and novel interaction among the individual processors. Baars proposed that a similar structure exists in the human brain in the form of the global workspace that supports conscious experience. The global workspace is accessible to all of the specialized processors, signifying that they can potentially have their contents occupy working memory. The global workspace can also "broadcast" its contents globally so that every processor receives or has access to conscious content. However, consistent with the notion that consciousness is serial and of limited capacity, only one processor's representations can be broadcast at any given time.

The Global Workspace Theory holds that consciousness is the entity that unites specialized and non-specialized processes through distributed information processing and permits them to interact. The global workspace serves as a central information exchange, similar to that used by artificial intelligence workers to permit any set of specialized processors to cooperate or compete in order to solve some central problem. In essence, a global data base is a memory to which all processors have potential access and from which all can potentially receive input. Any representation in the global data base is distributed to all of the specialized processors, but only some of them are able to act on the global data base in return, to propose hypotheses that can be distributed to any of the others. Only the global information is conscious; the operation of the specialists is not normally conscious.

Each specialist can decide on the relevance of the global representation for its own domain. Specialists are assumed to be triggered by a disparity between the global representation and their own internal representation of their domain. For example, spatial specialists are sensitive to visual input, and syntactic specialists are sensitive to linguistic input. In this manner,

the specialist can decide whether to process the global representation. All specialists are potentially responsive to global input, but they do not necessarily accept all global information. As Baars (1983) points out, the ability to distribute information globally is useful when it cannot be determined *a priori* which one of the processors needs the information. If the global representation is neither redundant nor irrelevant to some specialist, the specialized processor will attempt to adapt to the global information; i.e., it will attempt to reduce the mismatch that activated it. Baars likens this process to that of neural habituation, wherein neurons cease to fire with continuous input unless there is some change (i.e., a mismatch with the previous adaptation). At this point, the neuron will activate again until the new input has become redundant, equilibrium is restored, and it ceases firing again.

A number of different processors may cooperate or compete in sending hypotheses to the global data base by acting to confirm or disconfirm global hypotheses until all competition is resolved. In order to establish a stable global representation, a number of processors must cooperate to create what Baars refers to as a context--a set of stable constraints on a global representation provided by a set of cooperating processors. Those aspects of a global representation that are entirely stable can be called a context because they will influence other components to organize themselves so as to fit their constraints. A stable, global representation becomes conscious when and if it provides global information to the system as a whole. In other words, consciousness reflects a stable and coherent global data base that broadcasts information to the entire nervous system

As Baars (1983) points out, the concept of a global workspace has a number of advantages. First, since global information is distributed to all processors, a processor that is able to act on it can do so immediately. Second, under conditions of uncertainty, a global data base can combine information from many different sources to produce greater certainty than any one specialist could produce alone. Third, a distributed processing system with a global data base is inherently equipped for learning since it is an adapting system. Global information is distributed to many different specialists, which adapt to the new aspects of the global representation that are relevant for their domain. Fourth, a global data base can optimize the trade-off between structure and flexibility--the need for specialized, structured systems to handle standard problems versus the need for flexibility to cope with novel situations. A global data base allows one to change from a highly structured approach to a highly flexible one. One can

have the advantage of structure if the problem is relevant to some specialized processor as well as the advantage of flexibility in choosing among alternative processors or having a number of specialized processors cooperate to solve a problem. Fifth, the same processors can be used in different tasks. For example, since speaking and listening have many components in common, they may involve many of the same processors, which are simply organized differently for speech output and speech input. Sixth, the global workspace supports a highly adaptive allocation of processing resources. Many unconscious processes compete for access to consciousness, but the processing resources of the central nervous system are focused on the single most relevant stimulus at any given time. Hence, input is prioritized so that only the most dangerous, most attractive, most beneficial, or most interesting experience gains our attention.

The concept of a global workspace is not without its disadvantages, however. It requires a large number of processing resources to function since all specialists must continually monitor the central information relevant to their domain. Further, global problem-solving is relatively slow in comparison to the fast and efficient processing of a specialist that can handle a particular problem. Many different processors must learn to cooperate in order to produce a solution to a global problem. As Baars (1983) notes, these disadvantages of a global workspace bear many similarities to the functioning of consciousness, which seems to demand many resources, is slow in comparison to unconscious information processing, and is circumvented to speed up processing once a conscious solution to a problem becomes habitual or automatic. Indeed, as Baars argues, conscious processes are closely associated with a system that acts very much like a distributed system with a global data base.

In summary, Baars' (1983) Global Workspace Theory is able to explain a number of psychological phenomena and is consistent with many facts about the nature of consciousness. It posits a distributed information processing system composed of specialized processors covering all aspects of mental function. In lieu of a central executive control mechanism is a global workspace that permits the processors to interact and exchange information. Consciousness reflects the current contents of the global workspace and is closely identified with short-term memory and the limited-capacity components of the cognitive system. We are conscious of something when there is an interaction between input and context, resulting in a stable and coherent global representation that provides information to the nervous system as a whole. In other words, when we are conscious of something, we are adapting to it in a global way.

Multiple Resource Theory

Multiple resource theory is a model of human information processing designed primarily to account for variations in the efficiency of dual-task performance (i.e., the extent to which two tasks can be completed simultaneously as well as each can be performed alone). The theory centers around two major concepts: processing resources and structure. Processing resources refer to an "underlying commodity, of limited availability, that enables performance of a task" (Wickens, 1984, p. 67). During the performance of any task, many different mental operations may need to be completed (e.g., perceiving, rehearsing, responding), and each will require some of the operator's processing resources. Thus, resources can be thought of as the mental effort supplied for the performance of some mental operation. The allocation of resources is assumed to be under voluntary control, but the supply is scarce. Consequently, when the resources demanded by a task exceed the available supply, dual-task performance efficiency will suffer. Such degradations are likely to occur as the difficulty of either task increases, requiring more and more processing resources. In addition to the concept of resources, multiple resource theory further uses the concept of structure to explain variations in timesharing efficiency. Structure refers to such things as stages of processing, modalities of input, and requirements for manual response. Two tasks that demand resources from the same structure (e.g., both demanding visual resources) will interfere more than two heterogeneous tasks that demand resources from separate structures (e.g., visual and auditory).

The multiple resource theory of information processing was developed in large part as a consequence of the problems associated with the assumption of a single central pool of resources available to all information processes. Namely, single resource theory could not account for a growing body of empirical data from dual-task interference studies, which seemed to indicate that interference between tasks was dependent not solely on their difficulty but also on their structure (i.e., the stages, modalities, and codes of processing required). Presumably, a difficult task will demand more resources for its completion than a simple task and should therefore interfere to a greater extent with the performance of a secondary task. To the contrary, however, Wickens (1976) discovered that a manual tracking task was disrupted more by another motor response task than by an auditory signal detection task, *even though* the detection task was judged to be more difficult. Thus, it appeared that the two manual tasks were competing for

resources from a different pool than that required by the detection task on the basis of their stage of processing (response selection/execution vs. perceptual processing).

As another example, Wickens, Sandry, and Vidulich (1983) demonstrated that increases in the difficulty of one task do not always degrade the performance of a second task, as would be expected according to single resource theory. They paired a manual tracking task with an RT task wherein the stimuli were presented auditorially and required verbal yes-no responses from the participants. Increased difficulty of the tracking task had no effect on the RT task. However, when the RT task was altered so that the stimuli were presented visually and required manual responses, any increase in the difficulty of the tracking task degraded the concurrent performance of the RT task. Thus, in this instance, interference appeared to be a result of the codes of processing and output (spatial/manual vs. verbal/vocal) rather than difficulty. Interference occurred only when both tasks required spatial processing and manual output.

Consequently, multiple resource theory holds that the human information processing system consists of several independent capacities, each with its own limited capability to process information, as opposed to one single supply of undifferentiated resources (Navon & Gopher, 1979; Wickens, 1984, 1992). These separate capacities can be defined on the basis of three dimensions, as illustrated in Figure 7: (1) stage of processing (perceptual/central processing vs. response selection/execution); (2) modality of input (visual vs. auditory); and (3) codes of processing and output (spatial/manual vs. verbal/vocal). In essence, multiple resource theory holds that when two tasks demand separate rather than common resources on any of the three dimensions, timesharing will be more efficient and alterations in the difficulty of one task will be less likely to interfere with the performance of the second task. To the extent that two tasks impose similar demands, they will compete for common resources and disrupt dual-task performance.

First, with respect to stages of processing, the theory maintains that perceptual and central processing depend upon a common resource pool and that this reservoir is functionally separate from the resources underlying response selection and execution. Along these lines, Wickens and Kessel (1980) found that a tracking task that demanded response execution resources was disrupted by another tracking task but not by a math task that required central processing resources for its completion. Second, the argument that modalities of input define

resource pools holds that timesharing will be more effective between modalities than within a modality. Indeed, Isreal (1980) demonstrated that two tasks could be timeshared effectively when the information was presented to different sensory modalities. On the other hand, performance was disrupted when both tasks were presented within the same modality (e.g., both visual or both auditory). Finally, the third dimension of the model, codes of processing and output, implies that spatial and verbal processes each rely on functionally separate resources. This notion corresponds to anatomical evidence which indicates that spatial processing occurs chiefly in the right hemisphere of the brain, whereas verbal processing resides primarily in the left hemisphere. It is further supported by empirical data, as in a study which indicated that recognition performance was degraded when two spatial targets were presented simultaneously but not when a spatial and a verbal target were presented concurrently (Moscovitch & Klein, 1980).

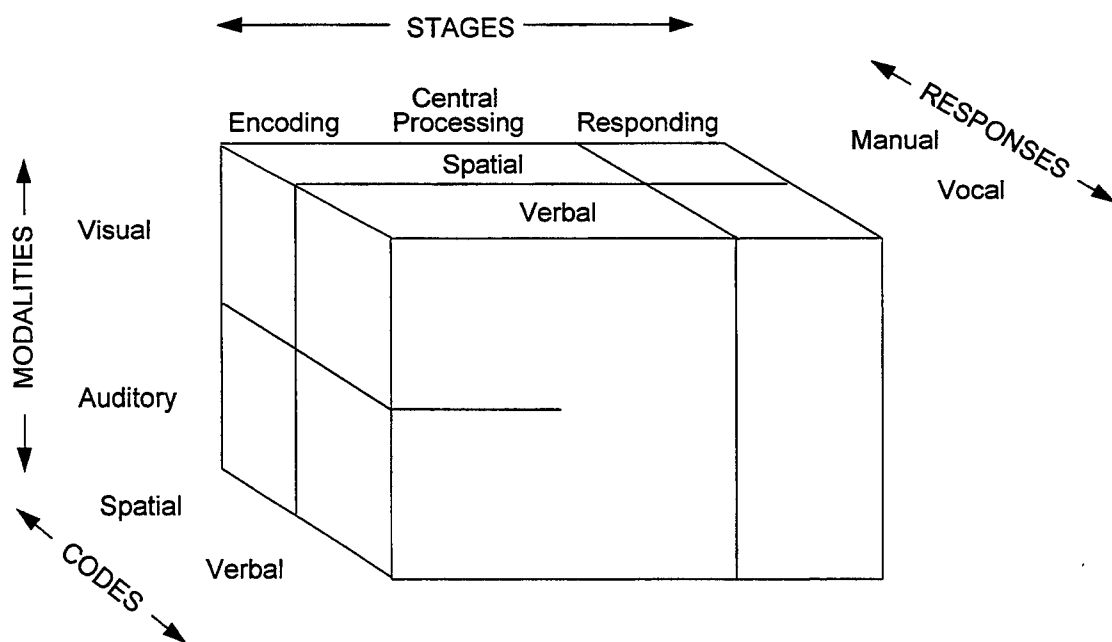


Figure 7. A proposed dimensional structure of human processing resources (Wickens, 1992, p. 375).

Although multiple resource theory has received support from a number of empirical studies, it has also received considerable criticism. One of the chief complaints is its lack of

precision in defining exactly what resources are (Kantowitz, 1985; Navon, 1984). For example, Norman and Bobrow (1975) refer to resources as "such things as processing effort, the various forms of memory capacity, and communication channels" (p. 45). This definition is not only broad but is also further hampered by the fact that it includes terms often regarded as synonymous with resources (i.e., effort and capacity). As Kantowitz (1985) points out, the definition refers to three major concepts (processing effort, memory capacity, and communication channels), each of which can be operationally defined in a number of ways. Thus, it provides examples of resources, but it never specifies precisely what they are and how they can be identified. Kantowitz further notes that it is difficult to find an explicit definition of the term in Kahneman's (1973) book, which is devoted entirely to attention and effort, even though the word appears on nearly every page. In essence, Kahneman defines a resource as a non-specific limited input, a definition that again fails to specify what a resource is. As Kantowitz argues, these vague definitions are even more damaging for multiple resource theory because they are the best attempts thus far.

Navon (1984) not only laments this lack of definitional clarity but also goes so far as to question the necessity of the theory altogether. He argues that many of the effects cited in support of multiple resource theory (e.g., task difficulty and dual-task degradation) can just as easily be explained by means of other approaches. That is, they are not exclusively predicted by a resource hypothesis and can often be expected regardless of whether limited resources are involved. "The issue is how necessary is the resource terminology for dealing with these phenomena" (Navon, 1984, p. 219). In a similar vein, Kantowitz (1985) points out that multiple resource theory introduces so many parameters that it is almost impossible to falsify the model. "One can always add another resource to explain results, or, if fewer rather than more resources are required to accommodate data, one can add a mysterious concurrency cost to 'explain' why certain combinations of task require more capacity than might be predicted at first blush" (p. 162).

Another troubling aspect of multiple resource theory is its circularity. First it is hypothesized that resources are necessary for task completion and that more resources are needed as the difficulty increases. Next an experiment is conducted to determine whether the hypothesis is supported by empirical data. The results indicate that performance deteriorated as the task difficulty increased, and it is concluded that performance faltered because of a paucity of

resources. However, in order to arrive at that conclusion, it was imperative to make the assumption that resources were needed in the first place. As Navon (1984) puts it,

...the hypothesis asserts something about a hypothetical variable, *amount of resources*, whose very existence (or its relevance for the performance of the task being studied) is the issue in question. To operationalize it one must implicitly *assume* interpretation in terms of resources of the empirical effects of dual-task deficit, priority, or difficulty of the concurrent task. For example, Prediction 1 presupposes that the amount of resources available to the target task can be constrained by the presence of a concurrent task, but this, in turn, requires the assumption that the concurrent task does indeed consume resources out of the same limited pool, which is just what is to be proved. (p. 223).

Despite these criticisms of multiple resource theory, Kantowitz (1985) nevertheless feels that it should not be abandoned altogether. As he points out, the ultimate criterion for the utility of any theoretical concept is its ability to predict behavior. And multiple resource theory does enjoy this advantage to some extent. It is able to explain much of the empirical data that single resource theory could not, and it is able to characterize the nature of many dual-task situations. As such, it can be useful during system design to make predictions regarding task interference. That is, using the guidelines provided by multiple resource theory, the system designer can strive for a design criterion that minimizes the overlap of demands on common resources in an attempt to select the configuration that will produce the best multiple-task performance.

Artificial Intelligence Models

General Problem Solver

The General Problem Solver (GPS) is an artificial intelligence (AI) model designed to simulate the processes humans go through during problem solving (Ernst & Newell, 1969; Newell & Simon, 1963, 1972). Its goal is not merely to solve problems efficiently but to emulate the processes that normal humans use when they attempt the same problems. The GPS was the first computer program able to simulate a variety of human symbolic behavior. More specifically, it is a heuristic program, one that solves problems and accomplishes complex tasks via intelligence (Ernst & Newell, 1969). The GPS was designed to be a *general* problem solver; i.e., it was intended to have the capability to solve a variety of problems and not just a single type. This

generality was accomplished through the problem-solving techniques made available to the GPS: heuristic search. The heuristic search techniques of the GPS involve the interaction of states, operators, and goals within a task environment. That is, given an initial situation, a desired situation, and a set of operators; the goal is to find a set of operators that will transform the initial situation into the desired situation.

Thus, during problem solving in the GPS, states or objects are transformed by various operators. Operators refer to actions that change the problem from one state to another. Information about the task environment is organized into subgoals, and the accomplishment of each subgoal leads to progress toward the goal state. The basic strategy used by the program to guide its search during problem solving is means-end analysis, which involves determining the desired “ends” and then figuring out the “means” by which they can be achieved. In means-end analysis, the focus is on the *difference* between the current problem (State A) and the goal state (State B). To achieve the goal state, the initial state must undergo certain transformations to make it identical to State B. The problem-solving process involves analyzing the features of A and B and detecting the difference between them by a matching process. Those features of A that do not match B undergo a series of transformations. These transformed features are then checked against B’s features, and the cycle repeats until a match is found. The problem is solved when the features of the current state are identical to those of the goal state.

In essence, the GPS program provides a way to achieve a goal by establishing subgoals whose attainment leads to the accomplishment of the initial goal. The GPS has three types of goals that it works to achieve. (1) The *transform goal* seeks to change State A into State B. (2) The *apply-operator goal* refers to the goal of applying some operator to State A. (3) The *reduce goal* attempt to reduce the difference between State A and State B. Similarly, the GPS also has three different methods, which provide various ways of accomplishing these goals to achieve the final goal of changing the original state into the goal state. (1) The *transform method* involves creating a new transition state via three steps. First, the original state is matched to the goal state to find the difference between them. Second, a new state is produced to reduce the difference. Finally, the new state is transformed into the goal state. (2) The *apply-operator method* involves finding a state to which an operator can be applied. If the apply-operator method determines that an operator can be applied to the original state, it applies it. If the operator cannot be applied to the original state, the state is changed to a new state to which the operator is then applied. (3)

The *reduce method* entails a search for the right operator for the situation. Specifically, it consists of searching for and applying an operator that will help reduce the difference between the current state and the goal state.

Thus, in order to achieve the goal of transforming State A into State B, the GPS first matches the states element by element. If the match reveals a difference, a subgoal is established to reduce the difference. The first step in difference reduction is to locate an operator that is relevant to the difference. The two basic criteria for selecting operators are desirability and feasibility: the operator should produce a state that is similar to the desired state, and the operator should be applicable to its input state. If an appropriate operator is found, a subgoal is set up to apply the operator to State A. When this subgoal is achieved, a new state, A', is produced that represents a modification of the original state in the direction of reducing the difference between it and the desired state. A new subgoal is then created to transform State A' into State B. A successful transformation here leads to final goal attainment.

To assess the similarity between the problem solving procedures used by the GPS and those used by humans, Newell and Simon (1963) compared the GPS trace with a human protocol from an attempt to solve a problem in elementary symbolic logic. During the procedure, the human participant was asked to verbalize his thought processes at each step. These were then used to obtain the protocols; i.e., verbatim records of everything the participant said during the experiment. While both the GPS and the human arrived at the correct solution, there were some important differences in the methods they used. First, the human coped with some items in parallel, whereas GPS always proceeded sequentially. Thus, while the human was able to handle two applications of the same rule simultaneously, the GPS could not. Second, GPS was unable to distinguish between internal and external applications. The human apparently applied some rules covertly without writing them out, but the GPS was not able to make distinctions between overt and covert actions. As Newell and Simon (1963) point out, this discrepancy can be problematic since written expressions can have very different memory characteristics from internalized expressions that remain in the head. Finally, at one point the human participant realized that he had misapplied a rule and proceeded to correct. Unlike the human, the GPS was not equipped with hindsight into its past actions. Despite these differences, Newell and Simon (1963) concluded that the GPS provided an adequate simulation of human problem solving behavior.

Subsequently, the GPS was used to simulate a variety of human problem solving behaviors in addition to symbolic logic. It has been used to solve transport problems such as the now-famous Missionaries and Cannibals task: three missionaries and three cannibals must cross a river. They have only one boat that holds two people. The number of cannibals on any side of the river must never exceed the number of missionaries, or the cannibals will eat the missionaries. The problem is to find the most efficient method for transporting all six people across the river without having any missionaries eaten. In addition to the Missionaries and Cannibals problem the GPS has been used to solve cryptarithmic problems, in which the problem solver must find the appropriate numbers to substitute for letters in addition and subtraction problems. The GPS has also been used to solve grammatical analyses of sentences, proofs in logic, and trigonometry problems.

Although the GPS has been able to solve several different types of problems, it is still quite limited. As Ernst and Newell (1969) point out, many of the problems that it can handle are simple by human standards. Further, there are numerous types of problems that the GPS cannot solve (e.g., tasks that involve extended use of arithmetic, tasks that involve a large data base, and tasks that require expertise). In fact, Newell and Simon eventually ceased working on the GPS because its generality was not as great as they had hoped. Nevertheless, the model is important because it represents the initial attempt to model human intelligence via a computer program. The GPS provided valuable information in the form of "a series of lessons that give a more perfect view of the nature of problem solving and what is required to construct processes that accomplish it" (Ernst & Newell, 1969, p. 2). It set the stage for the development of other AI models capable of greater generality.

ACT-R

The latest in a series of adaptive control of thought (ACT) models developed by John R. Anderson (1993) is the ACT-R (*R* for rational) model. It serves not only as a theory of human cognition but also as a computer program that processes information according to the tenets of the theory, thus providing information that can be used to test its adequacy. The foundation of the ACT-R theory is the production system. A production system is a cognitive architecture, or a relatively complete proposal about the structure of human cognition. Rather than simply trying to explain only a small aspect of cognition, it attempts to provide a complete specification of the

system. The primary driving force that structured the theory was Anderson's belief in the unity of human cognition; i.e., the belief that higher cognitive processes such as memory, language, and problem solving are all manifestations of the same underlying system. Thus, a single set of principles will suffice to explain all of these processes.

As summarized by Anderson (1993), the ACT-R theory emerged from the combination of four basic constraints. (1) It should be consistent with a wide variety of data. (2) It should be expressed as a production-system architecture in order to serve as a complete description of human cognition. (3) Because human cognition occurs in the human brain, the implementation of ACT-R should be in terms of neural-like computations. Therefore, in choosing among alternative approaches that seem equally viable, the one that more closely adheres to current knowledge of neural processing should be selected. (4) On the premise that human cognition is rational, or adapted to the structure of the environment, it should yield optimal behavior (thus the *R* in ACT-R).

The ACT-R theory makes the fundamental claim that a cognitive skill is composed of production rules and that these provide the correct architecture for achieving a unitary mental system. Thus, according to the ACT-R model, production rules are the answer to the question of what occurs in the human head to produce human cognition. Production rules are condition-action pairs, or pairs of IF-THEN clauses. The IF (condition) part specifies the conditions under which the rule will apply, and the THEN (action) part specifies what should be done under those circumstances. If the elements of the present situation match the condition, then the production rule can be applied, with the action dictating what to do.

In addition to production rules, the ACT-R model is characterized by three types of memory: working, declarative, and procedural. Working memory refers to active memory containing information that the system can currently access. It includes information retrieved from internal long-term memory stores as well as temporary information from the outside world deposited during encoding processes. Declarative memory refers to "what is" knowledge about the world that people can describe or report, including both episodic and semantic information. Declarative memory contains a network of knowledge represented in chunks or working memory elements (WMEs), which provide a means of organizing a set of elements into a unit for efficient storage in long-term memory. New information is stored in declarative memory by retrieving

declarative information that is already stored and temporarily holding it in working memory to process incoming information. Procedural memory refers to “how to” memory (i.e., knowledge required to engage in activities such as riding a bike or sending an e-mail). It can only be demonstrated through performance. The basic unit of knowledge in procedural memory is the production rule.

The ACT-R theory holds that complex cognitive processes are achieved through the completion of three repeating stages: (1) matching the conditions of various production rules with information in working memory; (2) selecting the rule that provides the best match based on ACT-R’s conflict resolution procedures; and (3) firing the selected rule. The first stage, pattern matching, is the process by which the system determines if a particular production rule’s conditions match the contents of working memory. In other words, the goal of this stage is to determine which conditions match the current problem to be solved. This is achieved in part through a process of spreading activation (the most prominent neural-like feature of ACT-R). A structure’s level of activation controls both the rate at which it is processed by the pattern matcher for production conditions and its probability of successful matching. Since information can have an impact on behavior only by being matched to the condition of some production rule, activation therefore controls the rate of information processing. As Anderson (1983) puts it, “it is the ‘energy’ that runs the ‘cognitive machinery’” (p. 86). Activation spreads from original sources to other related items that bear some association to current sources of activation. Thus, spreading activation favors information that is most similar to the immediate context. As might be expected, pattern matching is the most computationally demanding portion of executing productions. The process is assumed to occur in parallel, but some candidates will be completed before others because they are less complex or receive more resources. Hence, the ordering of production rules is serial, depending on when their computation is completed.

Any rule whose condition matches the information in working memory will work; however, because several different production rules may match, some system for determining which provides the best match is required. In ACT-R, this process of selecting the best production rule occurs in the second stage: conflict resolution. The approach in ACT-R is to design a conflict resolution system that minimizes computational cost while still retrieving the production rule that will produce the best result. Computational resources are devoted to production rules according to their likelihood of success. The production rule that is ultimately

selected is the one with the greatest expected payoff; i.e., the one that has the highest likelihood of leading successfully to the goal while simultaneously minimizing the cost of computational resources. The system stops when it determines that the computation time for further investigation of alternatives is not worth the expected improvement. At this point (Stage 3), the system fires the best production it has found *thus far* (not necessarily the best overall).

In summary, ACT-R maintains that human cognition can be explained by the condition-action pairs comprising production rules. To complete a cognitive process (e.g., adding two three-digit numbers), the conditions stipulated in the IF part of the production rule are matched to chunks in working memory representing the current problem to be solved. Several different production rules may apply to the same situation. The one that provides the optimal match while minimizing computational resources is ultimately selected, and the THEN part of the rule is executed. This action adds information to working memory that will be used in the process of achieving the goals that remain for the current problem.

There is considerable evidence to support many of the tenets of ACT-R, in part because the theory was developed on the basis of detailed phenomena in memory, learning, and control (Anderson, 1983, 1990). As Anderson (1993) points out, one line of evidence is the intuitive nature of using production rules to describe the cognitive processes involved in tasks like addition. The production rules seem to adequately capture the nature of the process. There is also abundant support for the existence of both declarative and procedural long-term memory stores. Specifically, the two memory stores appear to possess a number of different properties. First, they have been shown to differ in terms of their reportability (i.e., declarative knowledge is reportable but procedural knowledge is not). Further, declarative memory is subject to associative priming but procedural knowledge is not. For example, seeing the word *tennis* primes for the word *racquet* (i.e., the word can be read more quickly than if *tennis* had not appeared) but not for one's tennis skills (i.e., one cannot play tennis better or more rapidly). As a third example of their unique properties, the two memory stores differ in retention. People generally grow more skilled with procedural knowledge over time but worse at recalling declarative knowledge. ACT-R's conflict resolution scheme is supported by the observation that people tend to set some sort of acceptance threshold and then select the first action that exceeds the threshold. Empirically, Anderson, Kushmerick, and Lebiere (1993) showed that the amount of time to select a rule was not correlated with the mere quantity of alternatives, implying that

humans do not evaluate all possible moves, but rather stop once they have discovered one that is satisfactory.

Nevertheless, work remains to be done to further enhance the theory. Anderson and his colleagues plan to continue fine-tuning the production rule analysis of skill acquisition and are concurrently attempting to improve their understanding of the origins of these rules. In addition, they have begun more in-depth study of the time course by which knowledge progresses from declarative to procedural form. They intend to continue applying the ACT-R theory in real-world situations to assist students in skill acquisition, obtaining valuable information in the process that can be used to further refine the theory. Currently, the theory does not represent certain aspects of rational analysis; in particular, categorical and causal inference. A goal for the future is to incorporate these elements into ACT-R in order to enhance its treatment of the initial structuring of declarative knowledge. Further, as Anderson (1993) himself points out, the theory cannot cope with situations in which productions might misapply. That is, ACT-R can handle errors of omission (a production rule does not fire in time) but not errors of commission (the wrong rule applies), a feature that is required if the theory is to provide an accurate depiction of human cognitive skill acquisition.

Perhaps the greatest challenge for ACT-R is providing direct evidence for productions in human memory. Although the concept is intuitive and the theory is able to predict human performance and skill acquisition to a large extent, no direct evidence for productions exists. Further, as Barsalou (1995) points out, the production system language is so powerful that a set of productions could probably be written to explain any finding. Thus, the fact that a production system accounts for data is not necessarily evidence for productions *per se*.

Soar

Like the ACT-R model, Soar is a cognitive architecture, or a unified theory of cognition that attempts to provide a relatively complete proposal about the structure of human cognition (Laird, Newell, & Rosenbloom, 1987; Newell, 1990; Rosenbloom, Laird, Newell, & McCarl, 1991). In essence, the Soar architecture is an attempt to represent general intelligence, i.e., the diversity of intelligent behavior characteristic of human cognition. The ultimate goal is to provide the structure that would permit the system to perform the full range of cognitive tasks, use the full range of problem solving methods appropriate for such tasks, and learn as a consequence of task

completion. Soar proposes that the problem space and its components is the vehicle for achieving this goal. Like ACT-R, Soar is simultaneously a theory of cognition and a programming language for artificial language. Historically, Soar stood for *State, Operator And Result* to represent the notion that all problem solving in Soar involves a search through a problem space in which an Operator is applied to a State to obtain a desired Result. Over time, the community no longer regarded Soar as an acronym; consequently, it is not written in the upper case any more.

Soar was developed on the basis of four methodological assumptions. The first assumption was that it would be more useful to focus on the cognitive band as opposed to the neural or rational bands, not only because an understanding of the cognitive band can constrain neural and rational models but also because of the plethora of data about the cognitive band that could be incorporated into the Soar model. The second assumption was that general intelligence could best be studied without making distinctions between human and artificial intelligence. In fact, the ultimate goal is for Soar to serve as a basis for both human and artificial cognition. The third assumption was that the Soar architecture should be developed on the basis of simplicity and uniformity; thus, intelligent behavior should be described by a small set of independent mechanisms. Finally, it was assumed that Soar's adequacy could be evaluated only by a rigorous and long-term process of testing its limits and modifying the architecture accordingly.

According to Rosenbloom et al. (1991), Soar can be understood in terms of its ability to fulfill three critical requirements of a general intelligence. First, a general intelligence requires a memory with a large capacity for the storage of a variety of types of knowledge. This stored knowledge must be accessible for use in task performance. Second, a general intelligence must have the ability to generate or select a course of action that is appropriate for the current situation. Third, a general intelligence must be able to direct this behavior towards some end; i.e., it must be able to set and work towards goals.

Accordingly, Soar can be described as a sequence of three cognitive levels: a memory level, a decision level, and a goal level. First, with respect to the memory level, Soar contains both a long-term memory and a working memory. In Soar, all long-term knowledge is stored in a single production memory. Regardless of whether a particular item represents procedural, declarative, or episodic knowledge, it is stored in production memory in the form of one or more

productions. As in the ACT-R model, a production in Soar represents a condition-action or IF-THEN structure that executes whenever its conditions are met. Whereas the production memory is a long-term memory store, working memory is a temporary memory that contains all of Soar's short-term processing context. One important structure of working memory is the preference, which is responsible for determining the acceptability and desirability of actions. Acceptability preferences determine which actions are viable candidates (acceptability vs. rejection), and desirability preferences provide a partial ordering of potential actions (better vs. indifferent).

The second level of the Soar architecture, the decision level, is the level at which decision-making occurs. The decision level proceeds in a two-phase elaborate-decide cycle. During the elaboration phase, production memory is accessed repeatedly to retrieve into working memory all productions relevant to the current situation. The process occurs in parallel until quiescence is reached; i.e., until no more productions can execute. Following quiescence, the decision procedure selects one of the retrieved actions based on the preferences from working memory. The end result of the decision procedure is the selection of a single production that is new, acceptable, not rejected, and more desirable than other alternatives that are also new, acceptable, and not rejected.

The third and final level of Soar is the goal level. In Soar, goals are established whenever a decision cannot be made; i.e., when the decision procedure reaches an impasse. An impasse occurs when the decision procedure is unable to select an action due to incomplete or inconsistent information. Specifically, an impasse arises from one of four circumstances: when there are no alternatives that can be selected (no-change and rejection impasses) or when there are multiple alternatives but insufficient preferences to permit a choice to be made among them (tie and conflict impasses). When an impasse occurs, the Soar architecture generates the subgoal of resolving the impasse and creates a new performance context for doing so. Since nothing more can be accomplished in the original context because of the impasse, the creation of a new context allows decision-making to continue in pursuit of the goal of resolving the impasse. If an impasse occurs in this new subgoal, yet another subgoal and performance context are created. By responding to an impasse with the creation of a subgoal, Soar is able to search for additional information that can lead to resolution of the impasse. The subgoal is terminated when the impasse is resolved.

In Soar, all symbolic goal-oriented tasks are formulated in what are called problem spaces. A problem space consists of a set of states and a set of operators. The states represent situations, and the operators represent actions that yield other states once applied. Every task of attaining a goal can be understood in terms of finding a desired state in a problem space through the application of operators to a current state to yield a new state. That is, problem solving during goal attainment is driven by decisions that result in the selection of problem spaces, states, and operators. Given a goal, a problem space is selected in which achievement of the desired state can be pursued. An initial state is selected that represents the initial situation, and an operator is selected for application to the initial state to achieve progress toward the desired state. This process continues until a sequence of operators has been discovered that transforms the initial state into the desired state in which the goal has been achieved.

An example of Soar's problem-space architecture is portrayed in Figure 8. In this example, the task is to re-arrange a set of blocks from its initial state to a desired pattern. Any of the operators in the problem space relevant to this particular situation can be applied to the current state to attain the desired state. In this example, three operators are possible: *o1* places block *A* on top of block *C*; *o2* moves *A* to the table; and *o3* places block *C* atop *A*. Operating within the problem space requires knowledge to implement the operators and to guide the search. This knowledge is provided by the long-term production memory and brought to bear on the current state. Search control is knowledge that guides the selection of problem spaces, states, and operators through productions and preferences.

In Soar, all learning occurs by the acquisition of chunks. Chunks refer to productions that characterize the problem solving that occurs when accomplishing subgoals. The action portion of the chunk represents the knowledge that is generated during the subgoal; i.e., the results of the subgoal. The condition portion of the chunk represents an access path to this knowledge. The condition contains the elements of the situation that led to the creation of the subgoal and its eventual resolution. Chunking produces implicit generalization, which signifies that chunks can transfer to situations other than the one in which they were learned. The production will function not only in the exact same situation but in many others as well. Thus, the consequence of chunking is the avoidance of an impasse in a similar situation in the future. Decision-making will not stop because the architecture will henceforth be able to retrieve

information directly from production memory. It should be noted that this type of learning is experience-based since chunking occurs as a consequence of what Soar actually experiences.

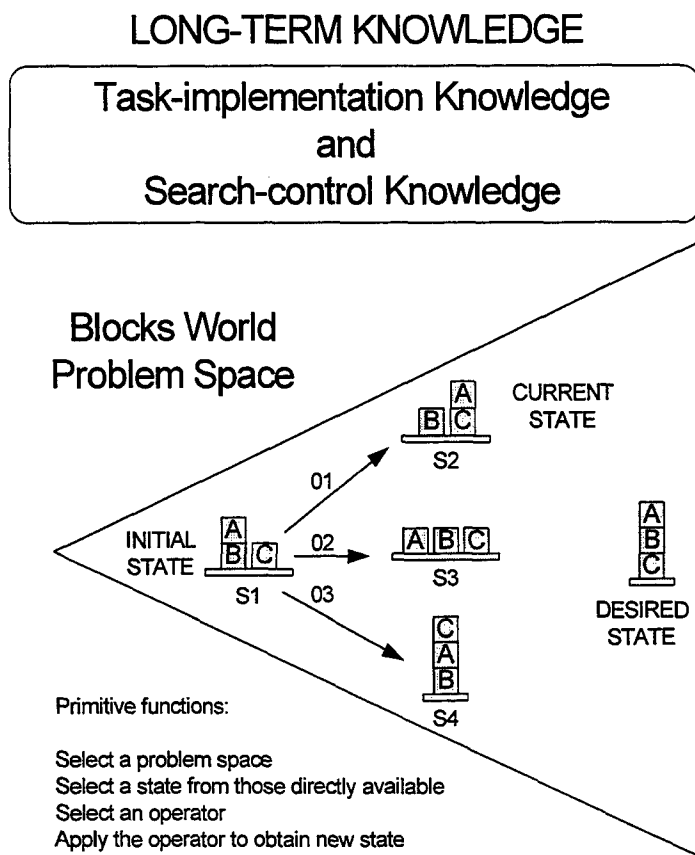


Figure 8. The Soar problem space architecture (Newell, 1990, p. 161).

Soar has been applied successfully in a wide variety of cognitive tasks, including search-based tasks, knowledge-based tasks, learning tasks, and robotic tasks (Rosenbloom et al., 1991). With respect to search-based tasks, Soar can perform over 30 different search methods during tasks completion. It is also capable of creating “hybrid” methods by combining various types. At least five varieties of knowledge-based construction and classification tasks have been implemented in Soar. For example, R1-Soar represents a construction task in which the goal is to construct a computer configuration. Neomycin-Soar is an example of a classification task in which the goal is to diagnose an illness by selecting among alternatives. As described earlier,

learning tasks in Soar occur through chunking. The precise type of learning that occurs depends on the subgoals that are created; however, what is learned can be transferred to other tasks within the same problem, other instances of the same problem, and other problems altogether. Thus, learning by chunking can occur in search-based, knowledge-based, and robotic tasks. One example of a robotic task is Robo-Soar, which is equipped with a Puma arm that enables it to solve block manipulation problems.

According to Rosenbloom et al. (1991), Soar's power and flexibility can be traced to at least four sources. First, its architecture is universal, providing Soar with the capability to complete any computable task. Second, the architecture is uniform. Soar has only a single type of long-term memory, a single type of task representation (i.e., the problem space), and a single type of decision procedure. Such features keep complexity to a minimum. Nevertheless, power and efficiency are still attainable through the chunking procedure, which enables the acquisition of new knowledge. Third, the specific mechanisms incorporated into the architecture provide another source of power. For example, the production memory provides access to large amounts of knowledge, whereas working memory allows global access to processing state. The decision procedure enables immediate reaction to new situations and provides the basis for the generation of new knowledge through impasse resolution. Fourth, power arises from the interaction effects that result from the integration of all of the capabilities within a single system. One example can be seen in the combination of strong methods, weak methods, and learning during task completion. Strong methods refer to possessing knowledge of how to proceed at each step. They tend to be efficient methods that produce high-quality results. Weak methods are based on searching to make up for a lack of knowledge as to what should be done. They make the system more robust by providing it with mechanisms for situations in which the strong methods are insufficient. Learning results in the addition of knowledge, transforming weak methods into strong ones.

Perhaps a more telling assessment of Soar's capabilities is an examination of how representative it is of human cognition. In response, Newell (1990) has provided a list of twelve features that Soar has in common with basic human cognition. A few examples will be presented here. First, Soar behaves intelligently; its behavior is predictable on the basis of what it knows and what it is attempting to achieve. Second, like humans, Soar is goal-oriented; it constructs its own goals whenever it is unable to proceed. Third, Soar's recognition system is highly

associative. It does not have deliberate access to its entire store of knowledge. Instead, retrieval cues in working memory (i.e., the characteristics of the current situation) function to elicit relevant knowledge. Finally, like humans, Soar does not often know how it does things since the learned procedures in the chunks are not articulable.

Although Soar does indeed possess many features of a general intelligence, it cannot yet accomplish all that intelligent humans can. Soar is not yet a nonbrittle system; i.e., one that does abruptly fail when it moves beyond some predicted task scope. Humans show some degree of nonbrittleness, though it is not complete. For example, teachers may function just as well in their own classrooms as in a new classroom; however, they may fail completely when transported to a steel mill). Thus, humans can cope outside a predefined task scope to some extent but not in all situations. According to Newell (1990), "Soar has enough attributes of general intelligence that it might permit a significant step in that direction, but no direct demonstrations of that have been attempted yet" (pp. 231-232).

GOMS

GOMS is an acronym formed from the words *Goals*, *Operators*, *Methods*, and *Selection Rules*. The GOMS model is a description of the procedural knowledge users must have in order to accomplish intended tasks on a given device or system (Card, Moran, & Newell, 1983). In essence, a GOMS model consists of descriptions for the *Methods* needed to complete specified *Goals*. The *Methods* are a series of steps consisting of *Operators* the user performs. If multiple *Methods* for accomplishing a *Goal* exist, then the GOMS model also includes *Selection Rules* that choose the appropriate *Method* for the context. In developing the GOMS model, Card, Moran, and Newell (1983) sought to fulfill two purposes. First, they wanted to construct a model consistent with the current state of knowledge regarding the various forms of human information processing: perception, memory, learning, problem solving, etc. Hence, the GOMS model is based theoretically on information-processing psychology. Second, they intended to bring this knowledge to bear on practical problems; i.e., to use their theoretical foundation to develop an applied psychology. Thus, they translated the theory into information that could readily be applied to real-world problems during system design. In so doing, they chose to focus on human-computer interaction. In particular, GOMS has been applied most extensively to the task of manuscript editing.

With respect to the theoretical underpinnings of GOMS, Card, Moran, and Newell (1983) drew upon a model of human information processing that is both justified by current research and suitable for an applied psychology of human-computer interaction. The resulting Model Human Processor is an attempt to articulate the mechanisms underlying user performance during human-computer interaction. The Model Human Processor can be described by (1) a set of memories and processors and (2) a set of principles of operation. Memories and processors are characterized by four important parameters: processor cycle time, memory storage capacity, memory decay time, and memory code type (acoustic, visual, physical, semantic). In addition, the memories and processors are grouped into three interacting subsystems: the perceptual system, the cognitive system, and the motor system. The basic Model Human Processor is depicted in Figure 9. As portrayed in the figure, the perceptual system consists of sensors such as the eyes and ears and their associated buffer memories, which are responsible for retaining output from the sensory system so that it can eventually be symbolically coded by the cognitive system. Thus, the perceptual system receives sensations of the physical world and stores them in sensory memory in a physical code; i.e., a non-symbolic analogue to the external stimulus that is affected by its physical properties. The primary buffer memories are the visual and auditory image stores.

The cognitive system receives information from sensory image stores in working memory and uses previously stored information in long-term memory to make decisions as to how to respond. Shortly after a physical representation of a stimulus appears in sensory memory, a symbolic representation (acoustic or visual) occurs in working memory. For very simple tasks, the cognitive system serves only as a connector between inputs such as these from the perceptual system and the appropriate outputs of the motor system. For more complex tasks, the cognitive system is further responsible for learning, retrieving information from long-term memory, and problem solving. Material from long-term memory is stored in a semantic code and becomes available to working memory through spreading activation. Activated elements of long-term memory consist of groups of related symbols called chunks. When a chunk is activated, the activation spreads to related chunks, which can in turn activate other related chunks. The basic principle of operation of the cognitive processor is the recognize-act cycle, which is comparable to the fetch-execute cycle of standard computers. On each cycle, the contents of working memory initiate associatively-linked actions in long-term memory (recognize), which in turn

modify the contents of working memory (act), setting up the next cycle. Finally, the motor system is responsible for carrying out the response dictated by the cognitive processor.

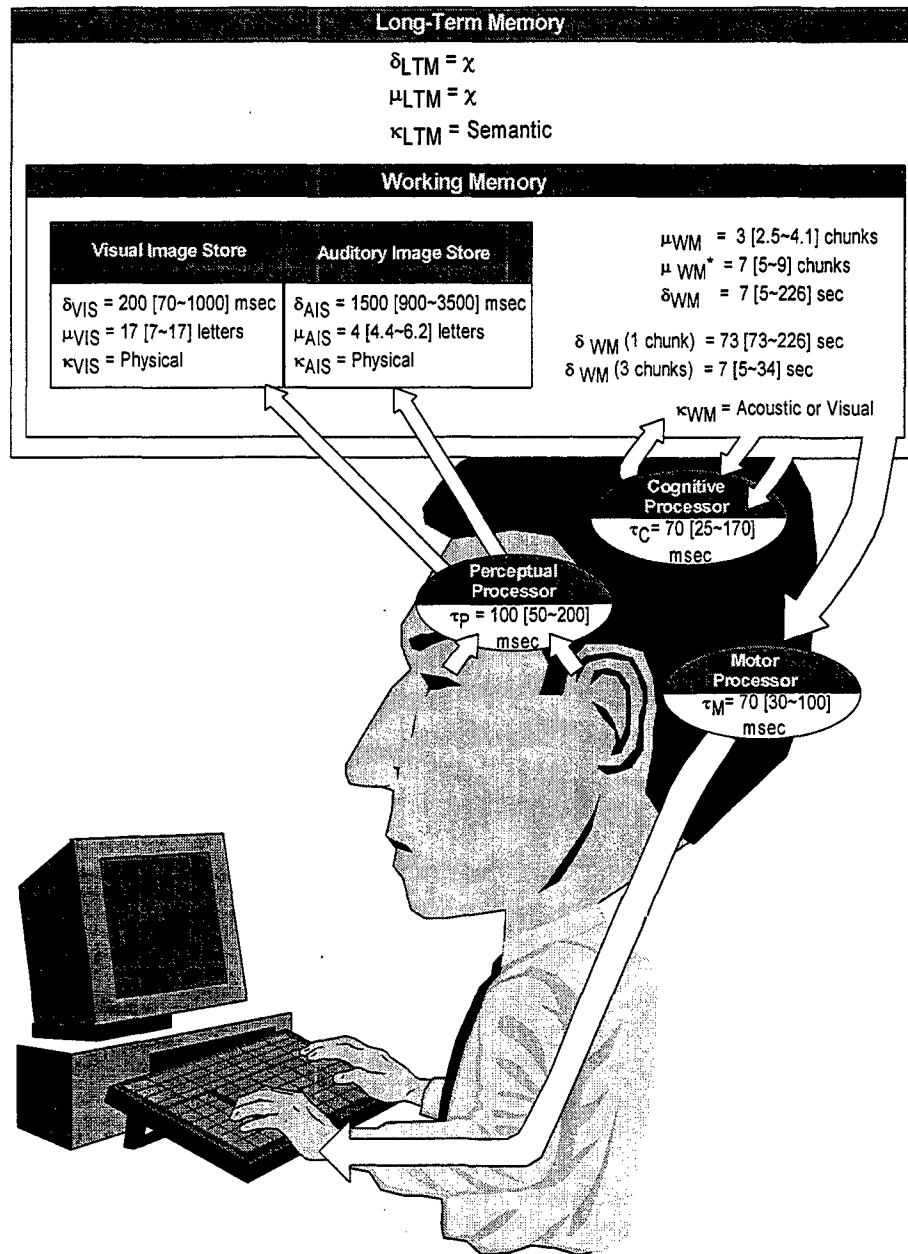


Figure 9. The Model Human Processor of GOMS (Card, Moran, & Newell, 1983, p. 26).

As Card, Moran, and Newell (1983) point out, "a model so simple does not, of course, do justice to the richness and subtlety of the human mind. But it does help us to understand, predict, and even to calculate human performance relevant to human-computer interaction" (p. 44). Thus, their main focus is on the application of the GOMS model. The Model Human Processor is simply a means for understanding what lies behind the goals, operators, methods, and selection rules that manifest themselves in human behavior.

In keeping with this philosophy, GOMS is based on the Rationality Principle of task analysis, which states that users will behave rationally to attain their goals. In order to predict a user's behavior, the task must first be analyzed to determine what the user's goals are. Consequently, it is no surprise that GOMS begins with the user's Goals. Goals are symbolic representations that not only define a state of affairs to be achieved but also determine possible methods by which they may be accomplished. Examples of goals include correcting a file in a word processor and editing a cell in a spreadsheet. Goals may be represented on several different levels. For instance, the higher-level goal of correcting a file can be decomposed into lower level goals called unit tasks, which represent the task of correcting each mistake in the file. The unit tasks can be further broken down into smaller sub-goals (e.g., locating the line that contains an error, modifying the error, etc.).

Operators are the perceptual, cognitive, and motor acts that change the user's mental state or the task environment. Examples of perceptual operators would be looking at a computer display and visually locating an error on a page. A cognitive operator might entail determining whether to proceed to the next page of a manuscript that the user is editing (if all errors on the current page have been corrected). Examples of motor operators include pressing keys on a keyboard and moving a computer mouse. Operators are defined by specific effects (output) and specific durations. For example, the output of the typing operator is a sequence of keystrokes, and the duration is a function of the number of characters to be typed.

Methods are exact sequences of operators that may be performed to achieve a goal. They are described as sequences of goals and operators, with conditional tests on the contents of working memory and on the state of the task environment. The following lines represent the method for achieving the goal of acquiring the unit task during manuscript editing:

Goal: Acquire Unit-Task

- . **Get Next Page** *if at end of manuscript page*
- . **Get Next Task**

The goal of acquiring the unit task will be followed by either **Get Next Page** or **Get Next Task**, depending on the state of the task environment when the conditional test "*if at end of manuscript page*" is performed. If the user is at the end of a manuscript page, he/she will go to the next page of the document and then get the next task (i.e., locate the next correction to be made). If the user is not finished with the current page, he/she will go immediately to retrieval of the next task on that page. Methods such as these in GOMS represent learned procedures that the user already has available; they are not plans created during task performance.

Finally, Selection Rules become important when there is more than one method for accomplishing a goal. Selection Rules are simple IF-THEN decisions that are used to determine which of several alternative Methods should be performed for a given task. For example, during the task of manuscript editing, either the mouse or individual keys on the keyboard may be used to move the cursor from one line to another. Some users may prefer the mouse the majority of the time, while others may prefer to use the arrow keys on the keyboard. However, these preferences may change further depending on the state of the environment. For example, IF the next correction to be made is many lines away from the change that has just been made, THEN the user may opt to scroll through the document quickly with the mouse rather than use the slower method of pressing the down arrow key. Selection Rules come into play in these situations.

A GOMS task analysis consists of describing the Goals, Operators, Methods, and Selection Rules for a set of tasks. One important feature of a GOMS task analysis is that the knowledge required to complete the tasks is described in such a way that it can actually be executed by either a computer or a human operator. A critical aspect of this analysis is deciding what should and should not be described. Those mental processes that are critical for interface design are also critical elements of the task analysis. Other mental processes that have nothing to do with interface design need not be analyzed.

GOMS models themselves can be presented at four levels of detail: (1) the unit task level, (2) the functional level, (3) the argument level, and (4) the keystroke level. The unit task level provides the highest level of abstraction and the lowest level of detail. This level is useful for structuring and exploring the tasks that a new computer-based system should support early in the design process. The functional level represents the decomposition of the unit tasks into the functional cycle of determining the unit task to be performed and executing it. At the argument level, methods are broken down into the individual steps of specifying commands and arguments that must be supplied to perform the task (e.g., locating the line with the error, modifying the error). Finally, the keystroke level provides the lowest level of abstraction and the highest level of detail. At this level, individual keypresses and mouse movements are represented, and basic perceptual, cognitive, and motor operations are introduced.

In order to assess the utility of GOMS for application to human-computer interaction, Card, Moran, and Newell (1983) conducted three experiments that focused on (1) Selection Rules in GOMS, (2) time predictions, and (3) grain of analysis. In all three experiments, participants were given a manuscript marked with corrections and asked to use a text-editor to make the corrections. The purpose of the first experiment was to determine if users' method choices for accomplishing various goals could be accurately described in terms of the Selection Rules of a GOMS model. In completing the task, users could choose from four alternative methods to locate the next line of text to be edited and two alternative methods for modifying the text. The results indicated first that there were clearly individual differences in how users decided which method to use. More importantly, however, the results showed that each individual user's behavior was highly structured and could be captured accurately by a GOMS model about 90% of the time.

The second experiment was designed to examine the sequencing and duration of the operators used to accomplish a task in order to test time predictions computed from the GOMS model of the manuscript editing task. First, the sequence of operators predicted by the model was matched against the sequence actually observed when the users completed the task. The percentage of matches varied from 79% to 98% with an average of 88%. Second, the times to perform unit tasks predicted by the GOMS model were compared with the observed times for the users to complete the unit tasks. This comparison indicated that the average model error was 33%. However, when the prediction was based on the time to edit the entire manuscript rather

than the time for each unit task, the error rate dropped considerably to only 4%. Thus, these outcomes indicated that GOMS could readily be used to predict task durations and determine which of several design options might be optimal in terms of the speed of task completion.

The purpose of the final experiment was to examine the effects of grain of analysis on the accuracy of the GOMS model. GOMS models were constructed at the unit task level (coarsest), the functional level, the argument level, and the keystroke level (finest). In terms of the accuracy for predicting sequences of operators, the match between predicted and observed sequences was 96% for functional level models. However, as the grain of analysis became finer, the accuracy declined (to a low of about 50% at the keystroke level). With respect to the accuracy of time predictions, the results indicated that accuracy at the functional level and finer levels was essentially independent of the grain of analysis. The average model error at the functional level was 29%. Overall, taking into account the various models, the average error ranged from 20% to 40%.

In addition to these three experiments which demonstrated the basic utility of GOMS, Card, Moran, and Newell (1983) undertook two extensions of the GOMS analysis. In the first rather straightforward extension, they demonstrated that it was possible to construct a GOMS model for a different text-editor from that used in their previous experiments. Second, based on the observation of widely varying individual differences in the three studies just described, they constructed GOMS models to predict user's actions. In these models, operator times had to be expressed as probability distributions rather than single numbers. In addition, the models contained probabilistic selection rules and conditionalities for predicting which method the user will employ. To date, however, the predictive validity of these models has not been examined.

As demonstrated by Card, Moran, and Newell (1983), GOMS models are useful for predicting learning and performance; for characterizing design decisions from the user's perspective; for user training; and for reference documentation. For example, a task can be broken down into its Goals, Operators, Methods, and Selection Rules for the purpose of determining task duration and comparing that estimate to an alternative task with a different design. In this way, GOMS models can be used during system design to choose among alternative design options. The task analysis can also be used to instruct users during training on the various steps that need to be completed to accomplish a goal as well as the alternative

methods that may be available to carry out a particular step. Finally, a GOMS task analysis provides a detailed account of the procedures involved in task completion that can be referenced as needed (e.g., during a later re-design of the system).

Although GOMS is capable of providing a complete and accurate description of error-free behavior, it is not appropriate if errors occur. This is a serious limitation since errors do exist in routine cognitive skilled behavior. In recognition of this fact, Card, Moran, and Newell (1983) have begun to devise methods for enabling GOMS to cope with errors. When conducting their study of the grain of analysis, they discovered that approximately 26% of the total time was spent on errors. Errors occurred on 36% of the tasks and doubled the amount of time needed to perform the tasks in which they occurred. Further, users typically committed one of two radically different sorts of errors: (1) small, frequent, routine errors that could be corrected quickly (e.g., typing errors) or (2) large, infrequent, but enormously time-consuming errors that required considerable problem-solving to correct (e.g., getting lost in a large document and attempting to find one's place).

According to Card, Moran, and Newell (1983), GOMS can be modified to handle errors with the addition of another type of goal--the correction goal--which would be accomplished by selecting a correction method. A user who makes an error during manuscript editing proceeds through five stages: (1) committing the error, (2) detecting the error, (3) resetting the editor to allow the correction, (4) correcting the error, and (5) resuming error-free activity. Thus, when a simple typing error occurs and the user becomes aware of the mistake, a Goal to correct the error is established. The Method for accomplishing the Goal might consist of the following steps:

Goal: Correct (Bad Character)

- . Look at display
- . Compare
- . Type (Delete character)
- . Type (Insert correct character)

That is, the user would observe the display and compare the character with the intended one, delete the faulty character, and type the new one before picking up where he/she left off.

Although these procedures have good face validity, their predictive validity has not yet been studied empirically.

Two other criticisms of GOMS may be made on the basis of its modeling approach. First, GOMS is not strictly a model of human information processing. Instead, it is a system for characterizing human behavior during task completion that draws upon existing knowledge of information processing. Therefore, GOMS in and of itself does not make any new or unique propositions about the structure of the human mental system. It simply proposes that we can use current information processing principles as a basis for understanding what mental actions might underlie observable behavior. Second, the GOMS model was developed with an initial focus on human-computer interaction; more specifically, on the task of manuscript editing. Although many of its devices may well be applicable to other domains, the developers of the model have not yet provided any additional examples. If GOMS is to be versatile and accessible, its utility in other areas (e.g., target acquisition) needs to be established.

EPIC

EPIC, which stands for Executive Process-Interactive Control, is an architecture for modeling human information-processing performance, with an emphasis on multiple-task performance (Kieras & Meyer, 1994; Meyer & Kieras, 1997). According to Kieras and Meyer, the goal of the EPIC project is to develop a comprehensive computational theory of multiple-task performance that (1) is based on current cognitive psychological theory as well as the results of empirical human performance studies; (2) will support the quantitative prediction of mental workload; and (3) is useful in practical system design. The final goal is particularly important so that the theoretical model can also serve as an engineering model; i.e., one that is both simple and quantitatively accurate enough for application to system design.

As described by Meyer and Kieras (1997), the EPIC architecture was designed on the basis of five guiding principles. (1) The model must represent an integrated information-processing architecture that incorporates known features of human cognitive processing and performance. (2) The model must incorporate a production-system formalism that permits specification of the procedural knowledge required to perform various tasks separately and in combination. (3) The limited processing capacity assumption, which holds that there is an upper bound on the number of tasks for which information can be processed concurrently, is not

necessary for a complete model of human cognition. (4) The model must emphasize the task strategies and executive processes that are critical for multiple-task performance. (5) The model must explicitly account for perceptual-motor constraints on multiple-task performance.

The EPIC model that emerged from application of these principles has many features in common with Card, Moran, and Newell's (1983) Model Human Processor described in the previous section. Specifically, EPIC was designed to combine the basic information processing and perceptual-motor mechanisms represented in the Model Human Processor with a cognitive analysis of procedural skill. Accordingly, EPIC consists of a collection of processors and memories. At the core is a production-rule cognitive processor surrounded by perceptual-motor peripherals. Thus, the processors can be subdivided into three classes. (1) The *cognitive* processor consists of a working memory, a long-term memory, a production memory, and a production rule interpreter. (2) The *perceptual* processors include visual, auditory, and tactile processors. (3) Finally, the *motor* processors include ocular, vocal, and manual processors.

The structure of the processors and memories in EPIC is portrayed in Figure 10. As shown in the figure, information flows from the sense organs, through visual, auditory, or tactile perceptual processors, to the cognitive processor, and finally back out to ocular, vocal, or manual motor processors. The cognitive processor consists of a production rule interpreter and a working memory. As in many other AI models, tasks are completed through the execution of production rules, or IF-THEN rules specifying the conditions that must be present in order for an action to be relevant. The production rule interpreter serves to determine the applicability of various production rules to the current situation. In EPIC, working memory contains all of the temporary information that is tested for and manipulated by these production rules. Working memory also contains control information such as task goals and sequencing information and provides a short-term storage area for sensory inputs. Two other forms of memory in EPIC are a long-term memory store for declarative information and a production memory for the storage of procedural knowledge. They provide critical information to the production rule interpreter during the evaluation of alternative production rules.

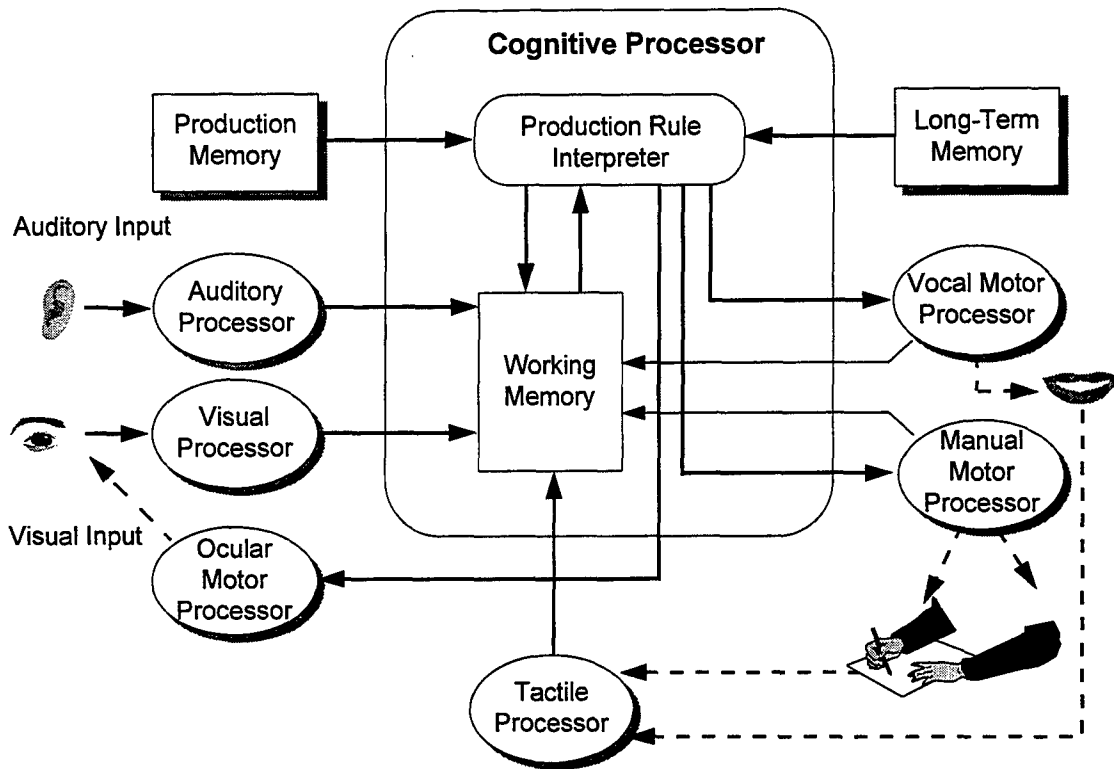


Figure 10. Overall structure of the EPIC architecture showing information flow paths as solid lines and mechanical control or connections as dotted lines (Kieras & Meyer, 1994, p. 3).

The cognitive processor in EPIC is assumed to operate cyclically without pausing between the end of one cycle and the beginning of the next. During each cognitive processor cycle, three types of activities occur. First, the contents of working memory are updated to reflect the results of actions completed by the perceptual, cognitive, and motor processors during the immediately preceding cycle. Second, the conditions of production rules are tested to determine which ones match the current contents of working memory. Finally, the actions of all rules whose conditions are satisfied are executed. Procedures such as conflict resolution and spreading activation mechanisms are not used to control which production rules are applied at a given time. Instead, the execution of a rule depends solely on whether its conditions are satisfied by the contents of working memory. Further, unlike the Model Human Processor, the cognitive processor in EPIC imposes no upper limit on the number of rules that can be tested or executed concurrently; therefore, multiple rules whose conditions are satisfied can fire and execute their actions in parallel. Because of this capability, EPIC is not characterized by a feature common to

many other information-processing theories, a central-processing bottleneck that limits response selection or other decision making for concurrent tasks.

During the execution of an EPIC model, a simulated human with general procedural knowledge of the task interacts with a simulated task environment. The inputs to EPIC's perceptual processors are assumed to be physical stimuli presented through simulated display devices for each sensory modality. The output from the model is the sequence of serial and parallel processes that take place in the course of task completion, the total time to perform the task, and various indices of mental workload (e.g., the amount of information that must be maintained in working memory). The construction of an EPIC model begins with an analysis of the information-processing requirements for a selected task or task combination. This analysis results in specification of what production rules are to be used in EPIC's cognitive processor, what the initial contents of working memory should be, and what stimulus inputs from the environment are required to start the task.

Next, if two or more tasks must be coordinated simultaneously, the model must specify how the actions performed by the separate sets of production rules for each task are to be coordinated. EPIC is capable of testing and executing production rules in parallel; however, some sort of supervisory control is required to ensure that all tasks are completed simultaneously without conflict (e.g., two tasks cannot use the same physical sensors such as the eyes at the same time). In EPIC, such executive control processes are handled by incorporating additional sets of production rules separate from those required for the individual tasks. The executive processes maintain task priorities and coordinate progress on concurrent tasks by means of various forms of supervisory control; e.g., inserting and deleting goals in working memory; directing the eyes to look at one place or another in visual space. Thus, the supervisory control mechanism in EPIC is not a structurally separate entity; rather, it takes the form of production rules whose format and application parallel the rule sets used to perform individual tasks.

In the context of multiple-task performance, EPIC assumes that capacity limitations arise as a result of limited structural resources as opposed to a limited cognitive processor. Since the cognitive processor fires rules in parallel, limitations must come from peripheral sources. According to Kieras and Meyer (1994), the limitations come from the peripheral sense organs and effectors. For example, the eyes are constrained in that they can fixate on only one location

at a time; a hand cannot be in two positions at once; etc. They tested this aspect of the EPIC model by simulating a dual-task situation in which an operator must perform two simple stimulus-response tasks in succession. In each case, the operator must make one of two responses depending on which type of stimulus occurred. Although the stimuli for the two tasks appear in succession, the time delay between them can be manipulated from very short (0 s) to long (1 s). Empirical studies with human operators have established that the reaction time for the second task increases substantially as the delay between the two stimuli decreases, but drops to a relatively fast baseline level if the delay is long enough.

The traditional explanation for this effect is that the central cognitive processor can complete only a single action at a time; hence, the process of selecting a response for the second task will be delayed until the response for the first task has been selected or executed. Conversely, because EPIC's cognitive processor is assumed to fire response selection rules in parallel, Kieras and Meyer (1994) have argued that the effect on reaction time for the second task stems not from the postponement of response selection within the cognitive processor but merely from a delay in response production. In order to test the validity of their claim, they constructed two EPIC models to represent both types of explanations: (1) an EPIC model based on the conventional assumption that the cognitive processor is able to select only a single response at a time, and (2) an EPIC model whose cognitive processor can select responses in parallel. In the first model, response selection for Task 2 must wait until the response for Task 1 has executed. In the second model, on the other hand, response selection for Task 2 can occur concurrently with the execution of Task 1. Hence, the second task can be executed as soon as the first is completed since the appropriate response has already been chosen. The results of the task simulation clearly indicated that the reaction times predicted by the parallel cognitive processor provided a better fit to the reaction times produced by human operators.

The success of the model in this respect was later shown to be generalizable to situations involving various types of stimulus-response combinations (Meyer & Kieras, 1997). For example, it was demonstrated that the EPIC model provided a good fit to the reaction time data from studies in which the stimulus modality was either visual or auditory, and the response modality was either manual or vocal. Overall, the fit between the empirical data and EPIC's predictions was 98% (Meyer & Kieras, 1997). Thus, one important contribution of EPIC was the demonstration that a central limitation is not a necessary assumption for an information

processing theory; rather, the data can be explained by competition for the same peripheral sources given the requirements of the task.

In addition to its application to dual-task paradigms, EPIC has also been used to simulate telephone operator tasks. The EPIC architecture was programmed with a set of production rules representing all possible instances of the tasks and responses required of a telephone operator. The perceptual and motor processors generated the times required to move the eyes around, perceive stimuli on the operator's workstation screen, and reach for and press the appropriate keys. Preliminary results indicated that the EPIC models were able to generate useful and accurate predictions for task completion times more easily than models that had previously been developed under the assumptions of the Model Human Processor (Kieras & Meyer, 1994).

Although EPIC has proven useful in predicting reaction times and task completion times, it is limited in one important respect. Namely, as Chong (1995) points out, EPIC is an architecture for modeling human performance that provides a thorough computational theory of human perceptual and motor processes but lacks an equally rigorous theory of cognition. For example, EPIC's cognitive processor has no learning capabilities and thus can never be entirely representative of human information processing. It may predict such things as task completion time perfectly well, but it cannot accurately represent what occurs during human information processing without a complete theory of cognition (e.g., it can model novice behavior and it can model expert behavior, but it cannot simulate the *transition* from novice to expert behavior). Thus, in an attempt to remedy this drawback of EPIC, Chong has proposed merging it with a cognitive architecture that does provide a thorough computational theory of human cognition: Soar.

In essence, EPIC has a theory of perceptual and motor processing but no theory of cognition. Soar has a theory of cognition but no theory of sensor or motor processes. A merging of the two systems would presumably produce a unified theory that combines the best of both worlds. To complete the merger, Chong (1995) replaced EPIC's cognitive processor with Soar's. In the new EPIC-Soar system, EPIC's cognitive processor serves merely to receive perceptual and motor processor messages as input and return motor commands as output. Everything in between is handled by Soar's cognitive processor. Specifically, EPIC sends perceptual and

motor messages to Soar and then waits while Soar processes the information and returns motor commands; EPIC then receives the motor commands and executes them.

To determine whether the merger would function as intended, Chong used EPIC-Soar to model a situation in which an operator must complete a tracking task and a choice reaction time task concurrently. The results indicated that the modified EPIC-Soar model provided more accurate predictions of reaction time for the choice task than did the EPIC model alone. Chong further demonstrated that EPIC-Soar was capable of modeling learning behavior on this same task. The "learning" model, which simulated the transition from novice to expert behavior in EPIC-Soar, provided more accurate estimates of the expert's choice reaction times than did the "expert" model, which simulated only expert behavior in EPIC-Soar.

Thus, it was shown that EPIC could benefit from the addition of a strong theory of cognition. EPIC and Soar were merged successfully, and the unified model provided more accurate predictions of reaction time than did EPIC alone. However, according to Chong (1995), the combined system will never reach the level of an architecture that is unified to begin with, primarily because EPIC is implemented in Common Lisp for Unix whereas Soar is implemented in C on SGI workstations. Hence, the procedures required to connect the two systems may allow them to function together, but they also slow the system down. Nevertheless, given Chong's demonstrations, additional work on the cognitive processor of EPIC itself may lead to a unified EPIC model that simultaneously provides strong theories of perceptual, motor, *and* cognitive processing.

Models of Visual Attention

Uniform Connectedness

The theory of uniform connectedness represents an attempt to specify the factors that account for the distribution of attention in space (Kramer & Watson, 1996). Traditionally, explanations have fallen into one of two categories: (1) Space-based attentional theories hold that attention selects regions of space independent of the objects they contain. Attention is viewed as a spotlight that illuminates a region of space. Objects that fall within the spotlight are processed; those lying outside the spotlight are not. (2) Object-based theories maintain that attention selects objects rather than regions of space. Selection will necessarily be spatial in nature since objects do

occupy regions of space; however, it is the objects themselves that are selected and not the regions.

There is psychophysical and neuropsychological support for both views, and it is now generally recognized that both space-based and object-based modes play an important role in visual selection (Kramer & Watson, 1996). The question that remains to be answered is, "Under what circumstances does each mode apply?" Several hypotheses have been developed. These include the shape judgment hypothesis, the mandatory processing hypothesis, and Kramer and Watson's (1996) principle of uniform connectedness (UC). According to the shape judgment hypothesis, visual selection will be object-based when the task requires judgments about geometric features such as shape that can be most easily computed within object-centered representations. Conversely, visual spatial attention will be space-based when the task involves judgments about other features such as hue, saturation, and brightness. In support of the shape judgment hypothesis, several studies have shown that performance is enhanced when individuals are able to reference shape judgments to a single object as opposed to multiple items. That is, observers are better at making shape judgments when they can focus their attention on a single object.

A second hypothesis, the mandatory processing hypothesis, holds that once an object is selected, all of its properties are automatically processed. Thus, according to this view, performance will be enhanced whenever individuals judge two properties of a single object, as compared with situations in which one property is located on each of two objects, regardless of whether the judgments involve shape, color, or luminance. Support for this view can be seen in the Stroop effect, in which individuals are asked to name the color of the ink in which a word is written. Participants are able to respond more quickly when the color of the ink and the word are congruent (e.g., the word RED printed in red ink) as opposed to when the ink and the word are mismatched (e.g., the word RED printed in blue ink). That is, all of the features of the object are being processed at once; facilitation occurs when the features match, and performance degradation occurs when they are incongruent.

The third hypothesis is the UC principle. According to this view, early visual processes of edge and line detection and figure-ground segmentation define an entry-level unit or percept on the basis of the principle of UC. This principle holds that regions of the visual field that have

relatively homogeneous surface characteristics such as lightness, color, motion, and texture tend to be perceived initially as single units or percepts. Subsequently, the entry-level percepts defined by UC can be combined into larger units on the basis of similarity, proximity, closure, etc. The entry-level percepts can also be subdivided into smaller components. However, the UC regions continue to have a strongly perceived identity even after they have been grouped or further subdivided. Thus, the UC hypothesis holds that object-based performance will be enhanced whenever a task requires the processing of multiple properties of a UC region. Conversely, object-based costs will occur when properties of several different UC regions are to be judged.

In an effort to assess the relative effectiveness of these three hypotheses, Kramer and Watson (1996) conducted an empirical study in which individuals were asked to perform a conjunction judgment task; i.e., they were asked to determine whether two predefined properties were present on each trial. A trial consisted of the presentation of a set of two wrenches on a computer display. The two predefined properties either occurred on a single wrench or were distributed across the two wrenches in the set. There were four conjunction judgment conditions: (1) open-end/bent-end, (2) texture/color, (3) color/gap, and (4) open-end/bent-end non-UC. In the first condition, participants were required to make shape judgments by deciding whether the set of wrenches possessed an open-end and a bent-end. In the second condition, they were asked to judge the orientation (horizontal, clockwise, or counterclockwise) and the color (red, green, or yellow) of the texture pattern comprising the wrenches. Each participant was assigned a target conjunction to look for; i.e., a combination of orientation and color that represented the presence of a target. In the third condition, individuals were required to judge color (red or blue) and gap size (large or small). As in the second condition, each individual was assigned a target conjunction that included one color and one gap size. Finally, in the fourth condition, they were asked to decide whether the wrenches possessed an open-end and a bent-end. This condition differed from the first, however, in that the shafts of the wrenches were colored with a blue and red checkerboard pattern.

On each trial, participants were asked to determine whether the two target properties were present, regardless of whether they occurred on a single wrench or were distributed across the two wrenches. Their response consisted of pressing one key on the computer keyboard if both properties were present and a different key if only one property was present. Both accuracy

and reaction time were assessed. Evidence for the shape judgment hypothesis would be provided if object-based effects were obtained in the two shape judgment conditions (1 and 4) but not in the texture/color and color/gap conditions. Evidence in favor of the mandatory processing hypothesis would consist of object-based effects in each of the conjunction judgment conditions (i.e., if individuals were faster or more accurate when both properties occurred on a single object as opposed to being distributed across the two objects). Evidence for the UC hypothesis would consist of object-based effects for the open-end/bent-end and texture/color conditions but not for the remaining two conditions. Presumably, the first two conditions would involve judgments of UC regions but the remaining two would not.

The results indicated that reaction time was significantly faster when both properties were contained in the same object *only* in the open-end/bent-end and texture/color conditions. This outcome is consistent with the UC hypothesis but not with the shape judgment and mandatory processing hypotheses. Consequently, Kramer and Watson (1996) concluded that object-based visual selection may be driven by the UC principle rather than by shape or mandatory processing. Their results suggest that shape judgments do not necessarily imply object-based visual selection. Further, object-based visual selection does not appear to entail mandatory processing; once an object is selected, it is not guaranteed that all of its features will be processed. Object-based advantages seem to hinge on homogeneity of the visual field. When the visual field has relatively homogeneous surface characteristics such as lightness, color, motion, and texture, the regions of the visual field will tend to be perceived as a single unit, leading to object-based visual selection advantages.

CODE Theory of Visual Attention

Like Kramer and Watson's (1996) principle of uniform connectedness, the CODE theory of visual attention (CTVA) is not a comprehensive model of human information processing; rather, it too focuses on a single aspect of information processing: visual spatial attention. The CTVA is a computational model that attempts to specify what it is that attention selects by integrating space-based and object-based theories of visual attention (Logan, 1996). Specifically, it was formed by merging the COntour DEtector (CODE) theory of perceptual grouping by proximity (Compton & Logan, 1993; van Oeffelen & Vos, 1982, 1983) with Bundesen's (1990) theory of visual attention.

CODE theory of perceptual grouping by proximity. The CODE theory of perceptual grouping by proximity provides two representations of space. (1) An analog representation of the locations of items in space is produced by bottom-up or data-driven processes that depend solely on the proximities of the various items in the display. (2) A quasi-analog, quasi-discrete representation of objects and groups of objects is produced by an interaction between top-down processes that apply a threshold to the analog representation of locations and the bottom-up processes that generated the analog representation itself. With respect to the locations of items in space, the CODE theory of perceptual grouping assumes that the representation of location is distributed across space. Thus, locations are not points but distributions in 1-D, 2-D, and 3-D space. CODE further assumes that the location of each item in space is represented by its own distribution. The separate distributions are then summed to produce what is referred to as a CODE surface that represents the locations of the items in space.

The representation of groups depends upon the application of a threshold to this CODE surface. The threshold provides a cutoff such that items residing in the same above-threshold region of the CODE surface belong to the same perceptual group. Items lying in different above-threshold regions are part of different perceptual groups. These concepts are depicted in Figure 11. The first panel of the figure shows a two-dimensional arrangement of five dots. The second panel demonstrates the 2-D CODE surface produced by the distributions representing the locations of the dots in space. Finally, the third panel shows the application of a threshold to the CODE surface. As can be seen in the figure, the threshold yields three perceptual groups. If the threshold is raised, the items will be separated into a greater number of perceptual groups. Ultimately, if the threshold is high enough, the items will be separated into five groups, each group containing a single dot. If the threshold is lowered from its present position, the items will be grouped into fewer categories. Ultimately, if the threshold is low enough, all five dots will be contained in a single perceptual group.

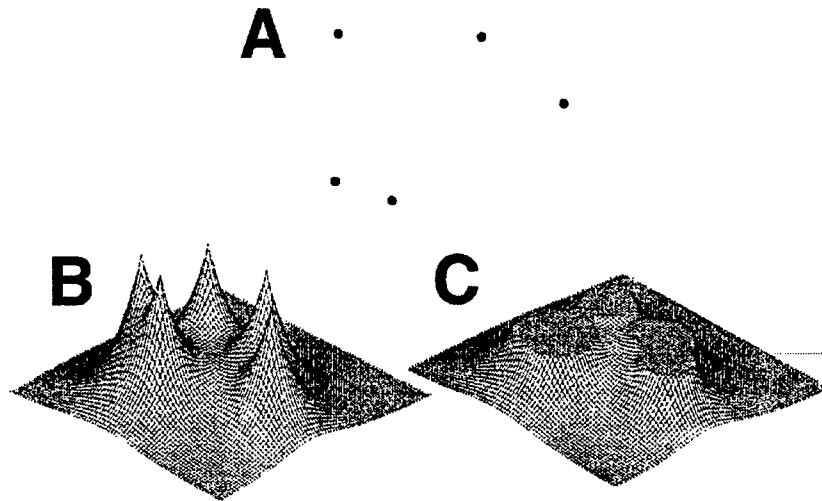


Figure 11. The representation of groups according to the CODE theory of perceptual grouping by proximity. Panel A shows a two-dimensional arrangement of five dots. Panel B shows the 2-D CODE surface for the five dots. Panel C shows the application of a threshold to the CODE surface (Logan, 1996, p. 607).

When the CODE theory is applied to attention, it is assumed that the distributions that comprise the CODE surface represent distributions of the features of items. Information about features of items is assumed to be distributed over space. The height of the distribution at any point in space represents the probability of sampling the features of the item to which it corresponds. This probability is typically highest near the center of the item and decreases as distance from the center of the item increases. The theory further assumes that attention selects among perceptual objects by choosing among above-threshold regions of the CODE surface. That is, attention samples the features of items that are available within the above-threshold region. The probability of sampling these features, which is equal to the area of the distribution lying within the above-threshold region, is referred to as feature catch.

The feature catch for a set of items depends on (1) the proximities of the items in the display, (2) the variability of the feature distributions, and (3) the threshold applied to the CODE surface. Whereas the proximities of the items is determined by the experimenter or the external world, distributional variability and the location of the threshold are treated as variable

parameters of the model. With respect to variability, an increase in the variability (λ) of the distributions impacts the feature catch in two ways. First, it reduces the contribution of items within the same group to the feature catch by reducing the area of their distributions lying within the above-threshold region. Second, an increase in variability increases the contributions of proximal items outside the group by increasing the area of their distributions lying within the above-threshold region. With respect to the threshold, an increase in the threshold reduces the magnitude of the feature catch, thereby decreasing the contribution of items inside and outside the above-threshold region. Thus, at high thresholds, attention will be focused on the central target item to the exclusion of nearby items.

Bundesen's theory of visual attention. The second component of the CTVA is Bundesen's (1990) theory of visual attention, which conducts further processing on the input that it receives from CODE (i.e., the sum of the feature catches from all items whose distributions lie in the above-threshold region). Bundesen's theory of visual attention was developed to explain the processes by which people choose among available inputs. As such, it is a necessary addition to the CODE theory because it permits selection from CODE's inputs. According to Bundesen's theory, choices are made among categorizations of perceptual inputs. Two levels of representation are assumed: (1) a perceptual level that consists of features of display items; and (2) a conceptual level that consists of categorizations of display items and display features. The two representations are associated by a parameter, $\eta(x, i)$, which represents the amount of sensory evidence for membership in category i that comes from item x . The variable x is an index for a display item, representing one member of a set of display items. The variable i represents a particular categorization for the item x (e.g., green or round).

The theory of visual attention holds that selection is made among perceptual items and categorizations by choosing a particular categorization for a given item. The final choice is determined by a "race" among the alternative categorizations. The first categorization to finish is selected, resulting in the selection of both a perceptual item and a categorization for the item. The η values are important determinants of the outcome of the race. These values, which represent the strength of evidence for the applicability of the categorizations that correspond to them, determine the rate at which those categorizations are processed. Larger η values signify greater evidence for a categorization and correspond to faster processing. All other things being equal, those categorizations with the largest η values are the most likely to "win" the race.

However, as Bundesen (1990) points out, the magnitudes of the η values can be further modified by the individual's personal bias (β_i) to apply a particular category to a given item and by the individual's priority (π_j) for attending to those items belonging to some category j . Thus, although the η value might provide strong evidence for a particular categorization, the ultimate decision will also depend on the person's biases. He/she may be biased toward applying some category to items or attending more to items that might belong to that category. As demonstrated by Bundesen (1990), the theory of visual attention is quite good at predicting both the accuracy and reaction time of categorization responses.

CODE theory of visual attention (CTVA). Merging the CODE theory of perceptual grouping by proximity with Bundesen's theory of visual attention produced the CODE theory of visual attention. First, CODE provides the input in the form of the feature catch, which represents the sensory data to define the η values. That is, the feature catch modifies the strength of sensory evidence from the various items in the display. Items falling within the perceptual group from which the feature catch is sampled contribute a great deal of sensory evidence. Nearby items in different perceptual groups contribute less information, and items far from the group contribute very little sensory evidence. Mathematically, the η values are multiplied by a number, c_x , between 0 and 1.0 that depends on the area of the distribution of x within the above-threshold region. Second, Bundesen's theory provides the β and π values that permit the selection of an appropriate response. Thus, it provides CODE with the capability for within-object selection and for response generation. The output from the resulting CTVA model includes predictions of reaction time and accuracy.

One other aspect of the CTVA is its interface with Logan's (1995) theory of attention, which attempts to account for selection between perceptual objects. Logan's theory is useful in completing the CTVA because it provides the mechanism for selecting which perceptual groups to sample. At any given time, feature catches from several different perceptual groups may be available for processing. The CTVA by itself cannot account for selection among the perceptual groups. Logan's theory involves two representations of the items in a display. A perceptual representation corresponds to the layout of objects and surfaces, and a conceptual representation consists of propositions that express the spatial relations among the objects. Directing attention from one perceptual object to another involves coordinating the two representations in order to

comprehend the spatial relations among the objects. Inputs to Logan's theory are schematic representations of objects as points, lines, and volumes. These inputs can be provided by CODE; i.e., by the perceptual objects defined by application of a threshold to a CODE surface. Thus, CODE provides information to determine which perceptual object to focus on as well as information to select a category for the perceptual object and make an appropriate response.

Relationships among the various components of the CTVA are depicted in Figure 12. In the early visual processes, location and identity are combined in the feature distributions and the CODE surface. The locations of the items in a display are determined by the environment, and the spread of the features from the items is determined by the CODE parameter for variability, λ . Application of a threshold to the CODE surface divides the display into perceptual groups that serve as inputs to the late visual processes, where identity and location are separate. The late identity system, Bundesen's (1990) theory of visual attention, is depicted in the lower branch of the figure. It takes the feature catch from each item in the display and computes the strength of sensory evidence (η values) for the categories relevant to the response alternatives. The η values, which are modified by bias (β) and relevance (π), determine the probability and latency with which different categories are selected. The late location system, Logan's (1995) theory, is shown in the upper branch of the figure. It takes as input the perceptual organization of the display provided by application of the threshold to the CODE surface. The late location system takes two different perceptual objects provided by CODE and outputs a relation between them.

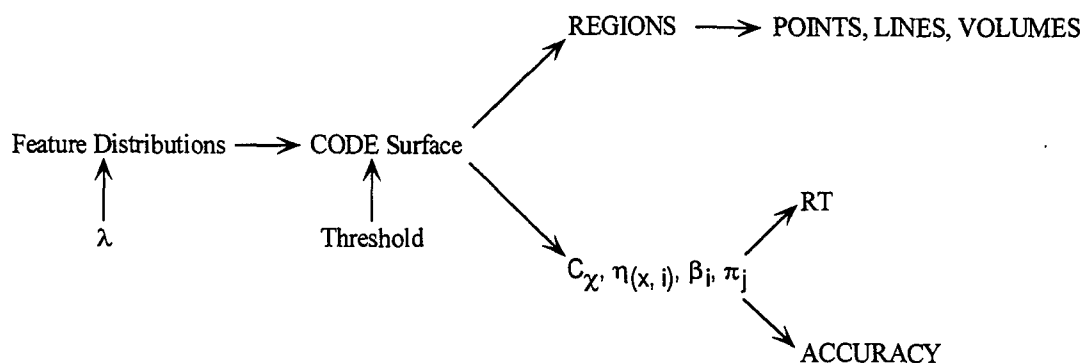


Figure 12. Architecture of the CODE theory of visual attention indicating the parameters and representations associated with the early identity and location system, the late identity system, and the late location system (Logan, 1996, p. 616).

Logan (1996) demonstrated that the CTVA was able to predict reaction time and accuracy data from seven empirical situations involving grouping by proximity and distance between items in a display. In one demonstration, Logan (1996) attempted to replicate the distance effects in illusory conjunctions from a study conducted by Cohen and Ivry (1989). An illusory conjunction is an erroneous combination of the features of different objects that generally occurs under conditions of stress or attentional overload. For example, the presentation of a green *T* and a red *L* might be misinterpreted as a red *T* (i.e., the observer mistakenly combined the identity of the first letter with the color of the second). Cohen and Ivry demonstrated that illusory conjunctions were less likely to occur as the distance between the objects decreased. In their first two experiments, they briefly displayed either a central digit (Experiment 1) or a pair of digits (Experiment 2) along with a pair of peripheral letters. The participants' task was to first name the digit (Experiment 1) or the smaller or larger of the two digits (Experiment 2) and then name the color or identity of one of the letters. One letter was always an *O*. The other was either an *F* or an *X*. The colors were pink, yellow, green, and blue. The *O* served as a distracter; the task was to name the color and identity of the letter that was not an *O*. The primary independent variable was the distance between the letters (near versus far). In both experiments, the probability of reporting combinations of letter identities and colors increased as the distance between the letters decreased.

This outcome was easily predicted using the CTVA. According to the CTVA, illusory conjunctions occur when the feature catch from a given above-threshold region contains features from different items and the first relevant features to finish the "race" come from different items. The probability of an illusory conjunction will depend on the overlap of the feature distributions from the different items in the feature catch. The greater the distance between the items, the smaller the overlap, and the less likely an illusory conjunction. The output from the CTVA accurately captured Cohen and Ivry's (1989) main finding: illusory conjunctions were more prevalent in the near condition than in the far condition.

In the remaining six demonstrations, Logan (1996) was able to show that the CTVA predicted data from six other similar studies reasonably well. (1) The CTVA replicated the finding that illusory conjunctions are more likely if the features that are combined belong to the same perceptual group than if they belong to different groups. (2) It predicted the improvement in reaction time that occurs when additional items are placed in a display in such a way that the

distracters cluster together and isolate the target item. (3) The CTVA replicated the result that the difficulty of searching for targets that are conjunctions of separable features (conjunction search; e.g., searching for a red *T* in a field of red *L*s and green *T*s) can be reduced by increasing the distance between adjacent items. (4) It predicted the ease of triple conjunction search, in which targets are conjunctions of three features and distracters contain only one target feature, as compared to double conjunction search, in which targets are combinations of two features. (5) The CTVA provided a mathematical account of the relationship between the difficulty of a discrimination and the search rate; i.e., search rate decreases as the difficulty increases. (6) Finally, the CTVA replicated the finding that nearby items associated with the same response as the target improve reaction time and accuracy, whereas flanking items associated with the opposite response degrade performance. These effects diminish as the distance between the target and the flanking items increases.

Evaluation of the CTVA. According to Logan (1996), the CTVA has several advantages. First, it is capable of providing reasonably accurate quantitative accounts of seven phenomena critical to visual spatial attention. As Logan points out, this advantage is important because the accounts of many competing theories of visual attention have been primarily qualitative. Second, the CTVA provides a formal representation of space in the attention literature by combining the best features of both space-based and object-based approaches to visual spatial attention. In Logan's (1996) words, "the CTVA model is strong primarily because it was built from strong components" (p. 641).

At the same time, however, the CTVA suffers from a number of limitations. First, the CTVA is abstract. For example, it does not deal with the nature of the features that comprise the feature distributions. It says nothing about the effects of motion. Second, according to Logan (1996), a more serious limitation derives from the CTVA's assumption that objects can be represented by points in space if the threshold is high enough. This assumption prevents it from dealing with objects that extend in space, with structured objects, and with interconnected or overlapping objects. In other words, the CTVA would be incapable of handling many real-world situations. Finally, the CTVA defines objects only in terms of location. While location may be an important defining characteristic of an object, it is certainly not the only one. Items may be grouped not only by proximity but also by similarity or common fate or connectedness. In fact,

one of Logan's (1996) goals for future research is to extend the CTVA to handle other grouping principles.

The effectiveness of the CTVA model can also be evaluated in terms of its practical utility. As a theory of visual spatial attention, it should have something to say about human behavior in the context of visual spatial tasks. Further, it should minimally produce some guidelines for system design. With respect to human attentional behavior, the CTVA would hold that the sensory characteristics of the objects in a display would be primary contributors to observers' identification decisions. However, observer bias and relevance can modify the strength of sensory evidence from a given object and influence observers' responses. Thus, an image analyst scanning a display for a target vehicle would base decisions primarily on the strength of his/her perceptions. The category to which a given object is assigned (e.g., target versus nontarget) will depend not only on these perceptions but also on the analyst's bias toward applying the "target" category. With respect to the design of the workstation itself, the CTVA would say that the proximity of items in a display is critical. Thus, objects (e.g., buttons, indicators, etc.) that belong to the same category might be placed close together. Items that should never be confused with one another should be placed far apart so they are not categorized in the same perceptual group.

Models of Language Comprehension

Construction-Integration Model

The construction-integration model is a cognitive architecture for comprehension that attempts to account for a wide range of language comprehension tasks (Kintsch, 1988, 1992a, 1992b, 1994a, 1994b). In essence, it attempts to clarify the processes involved in understanding material read from a text. It focuses primarily on the manner in which text-based material activates the comprehender's existing knowledge base and uses it to achieve an integrated representation of knowledge and text. Traditional views of knowledge use in discourse comprehension hold that comprehension is dominated by top-down effects and expectation-driven processing. That is, we understand much of what we read because we expect to see certain words and phrases, based on prior experience and knowledge. Our knowledge base itself provides part of the context within which the text is interpreted, serving as a filter that admits only the appropriate meaning of an ambiguous word and suppresses inappropriate ones. In other words, according to traditional

views of discourse processing, "people understand correctly because they sort of know what is going to come" (Kintsch, 1988, p. 164). Accordingly, analysis is assumed to proceed in a top-down predictive manner unless those expectations are proven wrong; it is only at this point that bottom-up or data-driven processing takes over.

In keeping with these views, the traditional approach to modeling knowledge use in comprehension has been to design powerful rules to ensure that the correct elements are generated in the right context. However, it is difficult to design a system powerful enough to produce correct results but at the same time flexible enough to function in the variable and ambiguous world of language comprehension. In an attempt to circumvent these difficulties, Kintsch developed a model of discourse processing with a much weaker rule-based production system that generates many elements, as opposed to attempting to produce a single correct element. The rules are powerful enough so that the correct element is likely to be among those generated, even though many inappropriate or irrelevant items will also be produced. Subsequent processing is used to strengthen the contextually appropriate elements and inhibit unrelated ones. The weak production system is advantageous because it equips the model with the flexibility needed to operate in a wide range of contexts.

More specifically, Kintsch's proposed *construction-integration* model combines (1) a construction process in which a text base is constructed from the text input as well as from the comprehender's knowledge base, with (2) an integration phase, in which this text base is integrated into a coherent whole. The knowledge base is represented as an associative network and, during the construction process, is assumed to be activated without guidance from top-down control structures. The construction process itself is modeled as a weak rule-based production system. Because both contextually relevant and irrelevant knowledge will be activated as a result of the manner in which the construction phase operates, the subsequent integration process is needed to weed out any irrelevant or contradictory material. In sum, as Kintsch (1992a) points out, the construction-integration model is a hybrid theory that effectively combines production systems and connectionist approaches.

Knowledge representation. Under the construction-integration approach, knowledge is represented as an associative net. Each node, or location in the network, represents a concept or proposition that can be linked with other nodes. The network then is a pattern of interconnected

propositions and concepts. Propositions are abstractions resembling sentences; they tie together concepts and ideas. A propositional network then is a pattern of interconnected propositions that make statements or assertions about the nature of the world. Connections among the nodes have a strength value ranging from -1 to 1. Nodes further consist of a head plus a number of slots for arguments, which may represent attributes, parts, cases of verbs, or arguments of functions. The immediate associates and the semantic neighbors of a node constitute its core meaning. Its full and complete meaning can only be created by examining its relations to all other nodes in the network. However, since it is impossible to deal with the entire net at once, only those propositions that can actually be activated at a given moment in time can affect the meaning of a concept. Consequently, the meaning of a concept is always situation specific and context dependent.

Construction process. According to the construction-integration model, a text base is constructed in four stages that involve: (a) forming the concepts and propositions directly corresponding to the text input; (b) elaborating each element by selecting a small number of its most closely associated neighbors from the general knowledge net; (c) inferring certain additional propositions; and (d) assigning connection strengths to all pairs of elements that have been created. The end result of the construction process is an "initial, enriched, but incoherent and possibly contradictory text base, which is then subjected to an integration process to form a coherent structure" (Kintsch, 1988, p. 166).

During Step A, a propositional representation of the text is constructed from the linguistic input and from a knowledge system as described earlier. For example, a propositional representation of the simple sentence, "Jane baked a cake," would show *Jane* and *cake* as the agent and object, respectively, in the *bake* proposition. Because the *bake* proposition requires a person as the agent, a test would be made of whether Jane is a person (e.g., by searching through the knowledge net for the *Jane* proposition). At this point, the model does not require that the correct proposition always be formed. Construction rules for building propositions are weakened, allowing for the construction of incorrect or incomplete propositions, which are then dealt with at a later stage.

In Step B of the construction process, each concept or proposition that has been formed in the first step serves as a cue for the retrieval of associated nodes in the knowledge net. For

example, the bake proposition would involve retrieval of those propositions closely associated with baking a cake. Thus, propositions such as eating cake, putting a cake in the oven, preparing dinner, and enjoying cake might be retrieved. At this stage, the construction process still lacks guidance and intelligence; items are simply produced in the hope that some of them might be useful.

In Step C, additional inferences are generated. The random elaboration mechanism of Step B will generally not be sufficient to produce all of the inferences necessary for comprehension. Thus, Step C involves a more controlled generation of specific inferences. Finally, in Step D, interconnections between all of the elements are specified. Elements are interconnected in one of two ways. First, the propositions derived directly from the text are positively interconnected with strength values proportional to their proximity in the text base. Second, propositions in the text base can "inherit" interconnections from the general knowledge net. That is, if two propositions relevant to the text are connected in the knowledge net with a particular strength value, s , they will have the same strength of connection in the text base itself.

Integration. During the final phase of the discourse comprehension process, integration is needed to clear up what is still an incoherent and inconsistent network. Specifically, the integration process removes unwanted and inappropriate elements from the text representation. Text comprehension is assumed to be organized in cycles that correspond roughly to short sentences or phrases. A new net is constructed in each cycle, with essential items from the previous cycle being carried over into the short-term buffer. Once the net is constructed, integration occurs; i.e., activation is spread until the system stabilizes. Stabilization generally occurs rapidly. If the integration process fails, new constructions are added to the net, and integration is attempted again. Normally, clusters of highly interconnected propositions attract most of the activation in the network, thereby deactivating sparsely interconnected portions of the network as well as nodes with negative links. Thus, the integration process produces a new activation vector with high activation values for some of the nodes and low or zero values for many others. Those nodes that are highly activated constitute the discourse representation formed on each processing cycle.

Applications. Kintsch (1988, 1992a) has described a number of domains in which the construction-integration model has been successfully applied. First, it has been used to

understand how knowledge is used during word identification in discourse. According to the construction-integration model, word identification occurs in a number of "stages." These stages are merely convenient labels, however; processing is continuous in reality. In the first stage, sense activation, the number of word candidates consistent with the perceptual input is progressively reduced through perceptual feature analysis. Once the number of candidates has been reduced to a manageable number, the semantic context becomes important. Thus, during the sense-selection stage, a small number of nodes is selected, each of which activates its strongest semantic or associative neighbors in the knowledge net. If there is a node whose associates fit the context, it will be taken as the meaning of the to-be-identified word. This association check is particularly critical for homonyms since perceptual analysis alone cannot determine which meaning is appropriate (e.g., *bank* can refer to either the financial institution or a river bank). Finally, during the sense elaboration phase, the meaning of the word is contextually explored and elaborated as more information about the context becomes available and the meaning of the discourse begins to emerge.

As Kintsch (1988) points out, this model is consistent with experimental data on the priming effect. The priming effect is the enhanced response speed that occurs when a target word is preceded by a closely related word. For example, observers are faster at determining that *apple* is a word when it is preceded by *tree* but not by *cat*. Studies have shown that a homonym can be primed by words related to any of its meanings (e.g., *bank* might be primed by both *money* and *river*). However, if sufficient time has passed to allow complete processing of the priming word in its context, only context-appropriate associates are primed. Thus, *river* will prime *bank* if the sentence reads, "Jerry slipped at the bank and got soaking wet," but *money* will no longer serve as a prime. This process of contextually-appropriate priming occurs by about 400 ms. Prior to that time, both appropriate and inappropriate words serve as primes.

In addition to word recognition in discourse, the construction-integration model has also been used to describe solving and understanding word arithmetic problems. It has been used to account for data from experiments on sentence recognition. Its applicability to poetic language has also been explored. As Kintsch (1988) notes, the model is essentially explanatory or qualitative in that it facilitates interpretation of various phenomena. It can be used to explain the subprocesses involved during different types of discourse comprehension. However, it does not

readily yield itself to quantitative predictions. At best, it can be used to predict that a particular problem will be difficult or easy, but it does not say how it will manifest itself quantitatively.

Summary. In summary, according to the construction-integration model, comprehension consists of constructing a mental representation of the information provided by a text and integrating it with internal knowledge, beliefs, and goals. This representation consists of concepts and propositions that form an interrelated network. Thus, comprehension consists of the construction of a propositional network. In the early phases of discourse comprehension, all knowledge that might potentially be related to the text input is activated. Consequently, a great deal of inappropriate or irrelevant information may initially be included. In later phases, the inconsistencies and irrelevant information are eliminated so that only the context-appropriate information remains. Unlike many traditional theories of discourse comprehension, the construction-integration model emphasizes the bottom-up, perception-like aspects of comprehension as opposed to the controlled, conscious, problem-solving processes.

Sanford and Garrod's Model of Written Discourse Comprehension

Like Kintsch's construction-integration model, Sanford and Garrod's (1981) model of written discourse comprehension is intended to explain the processes by which readers come to understand written words. As Sanford and Garrod point out, the comprehension of written discourse is more than just understanding the meaning of each sentence comprising the discourse. In most instances, the meaning of a text is highly dependent upon readers bringing additional knowledge to bear on the words on the page before them. For example, readers typically make a number of inferences when reading text in order to derive a coherent interpretation of the whole passage. They may make lexical inferences to solve problems involving lexical ambiguity (e.g., they infer that the word "bank" refers to the financial institution, given the context in which it is used). They may also make spatial and temporal inferences in order to organize the events and episodes that occur in the passage. Extrapolative inferences occur when readers extrapolate beyond the words that are actually printed in order to establish sensible links. It is rare for a writer to present all of the details surrounding an event; it is much more likely that the writer will assume readers are capable of understanding based on past experience and common knowledge; i.e., that they are able to extrapolate beyond what is actually given in the text. Finally, evaluative inferences arise when the significance of an event depends on the reader's knowledge of its consequences, given the context in which it is

presented. For example, the fact that Jo has a whole hour to kill takes on different meanings depending on whether she is at the airport waiting for her plane to arrive or going for a walk in the park on a sunny day because she has the afternoon off from work. As these examples illustrate, understanding how readers understand written discourse must involve an examination of how readers' knowledge structures influence their comprehension.

Thus, for Sanford and Garrod, the primary question is how written discourse makes contact with knowledge in order to bring about an understanding of its meaning. They use the term knowledge-base to refer to all of the information stored in memory that is brought to bear in understanding a piece of text. "On the page before the reader is a linguistic object, be it a single sentence or a larger piece of discourse; and in the mind of the reader reside knowledge structures of various kinds....The problems are: how the words relate to knowledge structures, which knowledge structures seem to be essential, and how the knowledge structures work to produce a final representation" (Sanford & Garrod, 1981, p. 38). In addressing these issues, Sanford and Garrod have focused on the manner in which the reader's memory structures might be organized to assist in knowledge access during written discourse comprehension. In terms of memory, they view the comprehension process as one of retrieving the appropriate information and constructing a rational interpretation of the text. Retrieval processes and construction processes can each be specified in terms of three variables.

Retrieval processes can be specified in terms of (1) the domain to be searched, (2) a given partial description of the information to be found, and (3) the type of information to be returned. Similarly, construction processes may be specified in terms of (1) the domain of memory where the construction is recorded, (2) a description of the information to be incorporated, and (3) the type of structure to result. To clarify upon the retrieval and construction processes, Sanford and Garrod (1981) have proposed a number of distinct partitions of memory, which serve as distinct search domains. Specifically, they have proposed that four partitions are necessary to capture memory access during comprehension. These partitions result from the combination of two dimensions. First, search domains may be in current focus or not in current focus. Information that is in current focus can be accessed rapidly and easily; it is held in dynamic partitions of memory since the contents of the partitions change with the text. Information that is not in current focus is more difficult to retrieve; it consists of both general knowledge and long-term representations of the text held in static partitions of memory. Second,

memory representations may derive solely from interpretation of the text, or they may be comprised of knowledge from other sources. This second dimension corresponds roughly to the distinction between episodic and semantic memory that is commonly made in cognitive psychology. Episodic memory refers to knowledge of particular episodes; both the episode and the information regarding how it was acquired are retained in memory (e.g., knowing how to swim and recalling how and when one learned). Semantic memory represents general knowledge that is dissociated from the specific situations in which it was acquired (e.g., knowing that the capital of Ohio is Columbus).

The application of these two dimensions results in four partitions of memory. Partition 1 is referred to as explicit focus. It is a limited capacity store that contains representations of entities and events explicitly mentioned in the text. Partition 2 is the implicit focus, a subset of general knowledge corresponding to the current scenario in the text. Partition 3 is long-term memory for the discourse, a subset of episodic memory. Partition 4 is long-term semantic memory, or the knowledge-base. According to Sanford and Garrod (1981), written discourse comprehension can be framed in terms of retrieval and construction processes operating within the constraints of these four partitions of memory. In further clarifying these processes, they have chosen to focus on problems of reference; i.e., understanding what words refer to in order to make sense of them.

The two partitions of focus, explicit and implicit, provide a retrieval domain that incorporates the information most crucial to understanding the text at any given time. Thus, they represent the current "topic" of the text, and they change along with the topic of the text. The focus is useful because it provides a narrower domain in which a search can begin. For example, some words provide very limited partial descriptions of what they might refer to and could easily generate a lengthy time-consuming search. Examples include words such as "he," "it," and "the man." Unconstrained searches for referents would return a representation for every single entity matching this partial description. The search becomes much more manageable if it is limited to a "likely" search domain based on the current focus of the text. This search may take place in either explicit or implicit focus. Explicit focus contains representations of things mentioned in the discourse, referred to as tokens, whereas implicit focus contains current scenario information. Explicit focus is a short-term store whose capacity is limited. Hence, as new tokens are added, old ones gradually diminish until they are no longer in focus at all. Implicit focus, on the other

hand, is a partition of long-term memory that is not constrained by capacity limitations. It consists of information from long-term memory that has the advantage of ease of access given its current relevance to the text.

Explicit and implicit focus are beneficial because they provide ready access to information of current concern in the text. They help the reader easily keep track of the meaning of the words in the piece of discourse. Invariably, however, words will appear whose meaning cannot be resolved in focus. In these cases, secondary processing outside of focus will be necessary. Thus, readers will have to rely on the slower and less accessible long-term episodic and semantic stores; i.e., they will need to search the entire memory space. Once a search must occur outside of current focus, the number of potential returns becomes very large, particularly if the partial description is not very informative. Secondary processing will occur in any situation where primary-level descriptions are inadequate to select a unique referent. For example, if a text had previously mentioned two characters by the name of John, one a banker and the other a mechanic, the reader may need to rely on long-term memory to recall which John is referred to at a given time. With longer texts especially, there may no longer be suitable tokens in focus, even though the individual or item had been mentioned before. In these cases, the appropriate search domain would be the static, long-term memory partitions. Secondary processing is also used to provide the initial scenario for implicit focus. When the reader first begins a new piece of discourse, a scenario is not yet available. Hence, one must be selected from those available in long-term memory.

In summary, Sanford and Garrod's (1981) model of written discourse comprehension is an attempt to develop a single framework encompassing various aspects of language processing. Their primary purpose was to develop a model that would clarify the processes by which readers understand the printed word in a piece of discourse. "If the processor is to find a referent for anything mentioned in a text, then this can be expressed as a procedure for searching memory to find another procedure which will accommodate the thing being mentioned" (Sanford & Garrod, 1981, p. 210). Accordingly, their primary explanatory mechanisms include four partitions of memory, which provide four distinct domains for searching for the meaning of a word. Explicit and implicit focus provide rapid and easily accessible search domains, whereas episodic and semantic long-term memory stores provide slower and less accessible search domains. Explicit focus contains tokens for items mentioned explicitly in the text, and implicit focus contains a

representation of the current scenario in the text (situations, events, objects, and characters). Episodic long-term memory is a memory store for the discourse itself, whereas semantic long-term memory represents the individual's general knowledge base. These latter two partitions are searched only when the relevant information cannot be retrieved from focus.

Models of Situation Awareness

Adams, Tenney, and Pew's Model of Situation Awareness

Models of situation awareness (SA) represent attempts to explicate the processes necessary for sustaining the minute-by-minute state of cognizance required to successfully operate and maintain a system. The term is most widely used in the commercial and military aviation communities to refer to a pilot's or air traffic controller's mental model of the system. It has become recognized as a crucial construct that lies at the heart of decision making and performance in complex, dynamic systems such as aircraft (Endsley, 1995). In fact, as Endsley (1995) points out, the critical importance of SA for crews of military aircraft was acknowledged as far back as World War I. "The safe operation of the aircraft in a manner consistent with the pilot's goals is highly dependent on a current assessment of the changing situation, including details of the aircraft's operational parameters, external conditions, navigational information, other aircraft, and hostile factors" (Endsley, 1995, p. 33).

Two representative definitions of SA are those offered by Endsley (1988) and Regal, Rogers, and Boucek (1988). According to Endsley, situation awareness is "the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future" (p. 97). This definition of SA emphasizes the role of determining the relevance and implications of events in a timely and appropriate manner. Thus, according to Endsley's definition, SA is more than just the ability to notice and attend to signals and sources of information in the environment; the individual must also be able to interpret those events and ascertain what they imply about future states. According to Regal, Rogers, and Boucek (1988), SA "means that the pilot has an integrated understanding of factors that will contribute to the safe flying of the aircraft under normal or non-normal conditions. The broader this knowledge is, the greater the degree of situational awareness" (p. 65). This view of SA emphasizes the role of prior knowledge in enabling the individual to comprehend incoming information. The greater the depth and breadth of the

individual's prior knowledge, the more likely he/she will be to understand that information in terms of the vast range of situations, implications, and response options that might accompany it.

As can be noted from both definitions, the processes necessary for the development and maintenance of SA involve significant cognitive effort. SA does not just happen. The individual must work to achieve and maintain it, and this work requires a good deal of mental effort. Further, acquiring and maintaining SA becomes increasingly difficult as the complexity and dynamics of the environment increase. In dynamic environments, the operator must make numerous decisions rapidly on the basis of an ongoing, up-to-date analysis of the environment. Adams, Tenney, and Pew (1995) have attempted to clarify the processes that are involved in the acquisition and maintenance of SA. At the core of their model is Neisser's (1976) view of the perceptual cycle, which Adams et al. modified by drawing upon Sanford and Garrod's (1981) theory of written discourse comprehension. Originally, Sanford and Garrod's theory was intended to explain the comprehension of events in written discourse; however, Adams et al. extended it to the comprehension of events in flight. Thus, their model of SA represents a merging of Neisser's perceptual cycle with Sanford and Garrod's theory of event comprehension.

Neisser's view of the perceptual cycle, which is portrayed in Figure 13, was designed to depict the interdependence of memory, perception, and action. Neisser argued that knowledge (in the form of schemata or mental models) leads to the expectation of certain types of information. Thus, the schemata that are active at a given time will structure the subsequent flow of events. That is, they will serve to increase the individual's receptivity to certain aspects of the environment and to particular interpretations of the available information. These concepts are depicted by the inner circle in Figure 13. An individual's present frame of mind (i.e., schemata) will direct where he/she looks, which in turn affects what information is selected for further processing. Only that information that is selected has the capacity to affect the individual's schemata, whereupon the cycle repeats itself. The outer circle in the figure represents a more general exploratory cycle, which Neisser added in order to handle cases where perceptual exploration uncovers information that the schema did not expect or fails to obtain the data that were anticipated. Thus, the general exploratory cycle can include actions taken to secure information that is not present in the immediate environment. It has access to the individual's larger knowledge of the relevant world and its possibilities.

Adams et al. (1995) chose Neisser's perceptual cycle as the core of their model in part because the components central to SA are inherent to the framework. For example, it can be used to explain SA as both product and process, a distinction commonly made in the SA literature. The *product* of SA refers to the state of awareness with respect to information and knowledge. The *process* of SA refers to the perceptual and cognitive activities involved in forming and revising the state of awareness. As Adams et al. point out, the processes of SA not only determine the products but also are affected by them as well. Thus, the processes of information acquisition and revision determine the ultimate state of awareness. However, the processes that are employed are themselves determined by expectations, hypotheses, and familiarity with the situation. Products and the processes are interdependent, a relationship that is reflected nicely by Neisser's framework. As a product, SA is the state of the active schema, the conceptual frame or context that determines the selection and interpretation of events. As a process, SA is the state of the perceptual cycle at any given moment. Product and process influence each other cyclically.

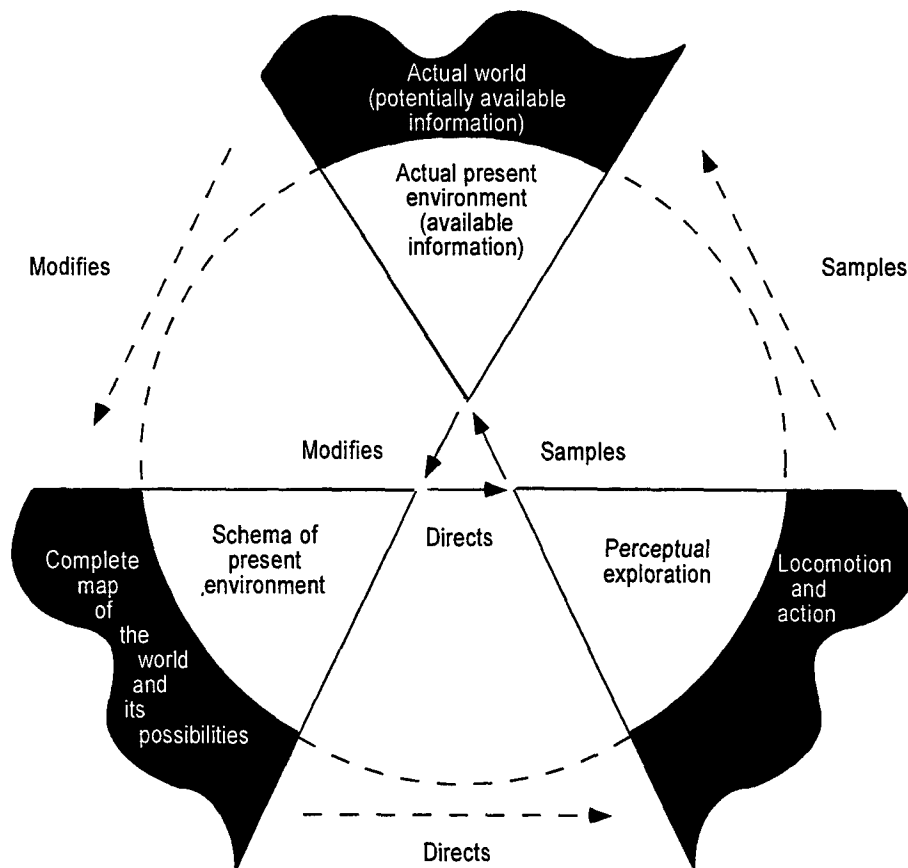


Figure 13. Neisser's perceptual cycle. The inner circle depicts the perceptual cycle and the outer circle depicts the general exploratory cycle.

Given that a critical component of SA is the ability to manage multiple tasks effectively, Adams et al. modified Neisser's perceptual cycle in order to produce a framework encompassing multiple task management. In the dynamic and multidimensional environments of flight management and air traffic control, the operator must know which tasks to perform and when to perform them. For example, while attempting a landing in stormy weather, the pilot needs to monitor the descent, perform the prelanding checklist, set the flaps/slats, monitor the copilot's performance, look out the window for traffic, receive and respond to radio messages from air traffic control, and monitor airspeed, among other activities (Adams, Tenney, & Pew, 1995). As Adams et al. point out, "the tasks in the queue must be prioritized and interleaved with deference to both the temporal requirements on their execution and their overall importance to the management of the situation as a whole" (p. 91). One benefit of SA is that the operator is better prepared to cope with upcoming events. At the same time, however, the mental effort needed to generate and maintain this level of SA may detract from task completion at times. Considerable cognitive effort is required not only to remember what tasks need to be completed and when, but also to actually carry out the tasks when the time comes.

In order to deal with such considerations, Adams et al. (1995) incorporated Sanford and Garrod's (1981) work on event comprehension. As described earlier, Sanford and Garrod maintain that event comprehension can be understood in terms of the functioning of active memory and long-term memory, each of which can be subdivided into two separate components. Active memory consists of an explicit focus and an implicit focus. Explicit focus, which corresponds to what is commonly referred to as working memory, contains a limited number of "tokens" that serve as pointers to larger knowledge structures in long-term memory. The maintenance of a token in explicit focus depends both on the recency of its direct activation by the situation and on its relevance to the current state of the situation. The other component of active memory, implicit focus, contains the complete representation of the schema that is partially represented in explicit focus. Information relevant to the knowledge in implicit focus cannot be elicited as quickly as in explicit focus, but it can be interpreted much more rapidly and at a lower cost to workload than information unrelated to the contents of explicit focus.

As with active memory, inactive or long-term memory can also be subdivided into two bins: episodic and semantic memory. Episodic memory is said to contain all of the knowledge structures that have been constructed over the course and context of the current situation (e.g.,

reading text or flying a plane). Semantic memory contains general knowledge that an individual has accumulated throughout a lifetime. Knowledge from either type of long-term memory can be brought to consciousness only as a result of considerable effort or strong environmental cueing and thus at a great cost to workload.

The result of merging Sanford and Garrod's (1981) theory with Neisser's (1976) perceptual cycle is depicted in Figure 14. As can be seen in the figure, explicit focus and implicit focus replace Neisser's "schema of the present environment." Episodic and semantic long-term memory replace Neisser's "cognitive map of the world and its possibilities." The new model is better equipped to handle the multi-task environment associated with SA because it enables one to make predictions as to how the operator will cope with tasks in the queue. For example, events that are relevant to those aspects of the task on which the individual is currently working should be readily assimilated because they will relate to knowledge currently in explicit focus. Events that relate to the task but not to the particular aspect of current interest can also be interpreted fairly quickly since they map onto knowledge in implicit focus. On the other hand, if the interpretation of an event requires consideration of inactive knowledge in long-term memory, the probability that it will be processed will depend on its significance and the time available for working on it.

The Adams et al. model can also be used to predict what will happen when task completion is interrupted by another event. If interpretation of the interrupting event requires task-incompatible use of knowledge already in focal memory, then the mental records for the original and interrupting events may become confused. If the interpretation requires knowledge distinct from that which currently occupies focal memory, then the interrupting event can be dealt with only at the expense of the original task (i.e., by displacing the current contents of explicit focus). Thus, the difficulty of reinstating an interrupted task would depend on its similarity to the main task in terms of the potential confusability of the two sets of information. Furthermore, the model can be used to predict which of many tasks the operator will choose to perform next. According to the model, the availability of tasks will be modulated by the operator's larger knowledge of their status and structure. Thus, tasks will be selected in proportion to their determined urgency or criticality on the basis of information contained in episodic and semantic memory.

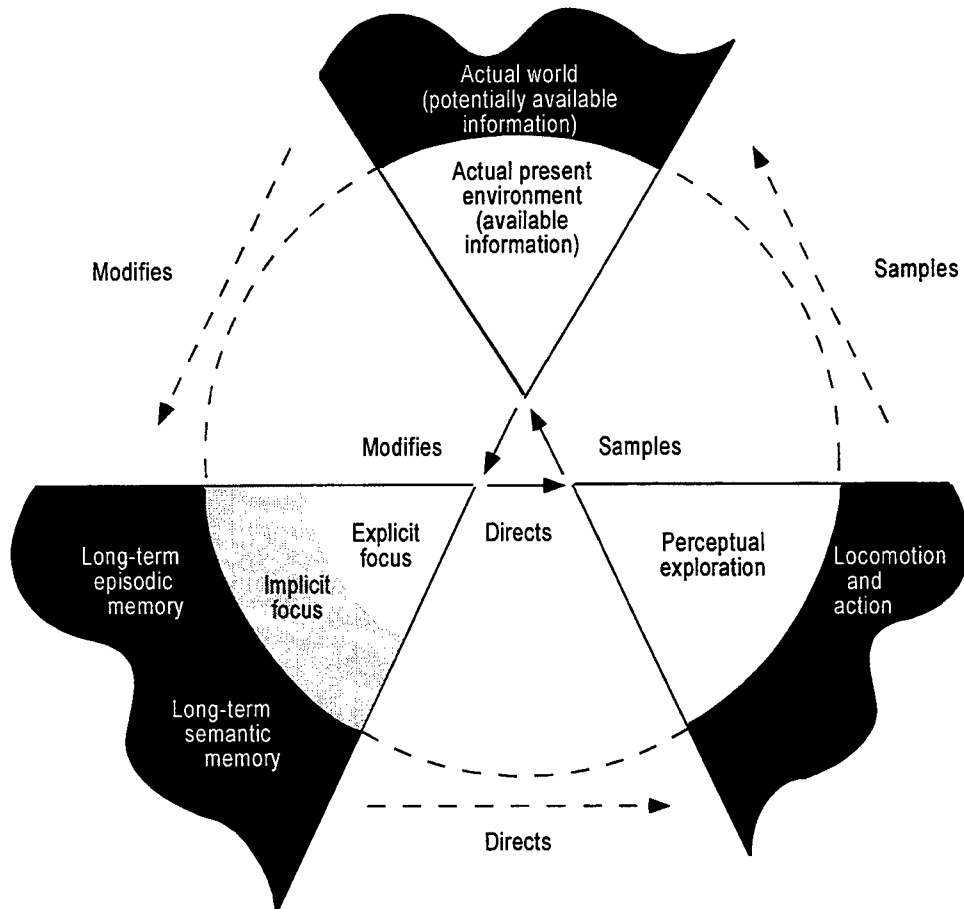


Figure 14. Adams, Tenney, and Pew's (1995) modification of the perceptual cycle (p. 91).

Endsley's Model of Situation Awareness

Yet another model of SA, in addition to Adams, Tenney, and Pew's (1995), is the one developed by Endsley (1987, 1988, 1995). Her model serves to expand upon and clarify the concepts comprising her definition of SA: "the perception of the elements in the environment within a volume of time and space, the comprehension of their meaning, and the projection of their status in the near future" (Endsley, 1988, p. 97). As stated earlier, SA has come to be recognized as crucial for pilot decision-making and performance. Consequently, as depicted in Figure 15, Endsley's model of SA places SA as a precursor to both decision-making and performance. According to her model, SA can be described in terms of three different phases or levels. Level 1 SA involves the perception of elements in the environment. It forms the most basic foundation for SA. Only those elements that the operator perceives can receive further consideration in later stages. Thus, the misperception of items or the failure to notice critical information at this stage

can lead to serious distortions in operator SA that will subsequently affect decision-making or performance. Level 2 SA involves the comprehension of the meaning of the elements that have been perceived in Level 1. In this stage, the operator attempts to synthesize the elements that have been perceived to form a coherent picture of the current situation. Thus, the operator goes beyond simple awareness of environmental elements to an understanding of their significance. Finally, Level 3 SA involves the projection of future status. At this level, the operator uses what has been perceived and what is known about the meaning and significance of those perceptions to determine the status of the system in the near future.

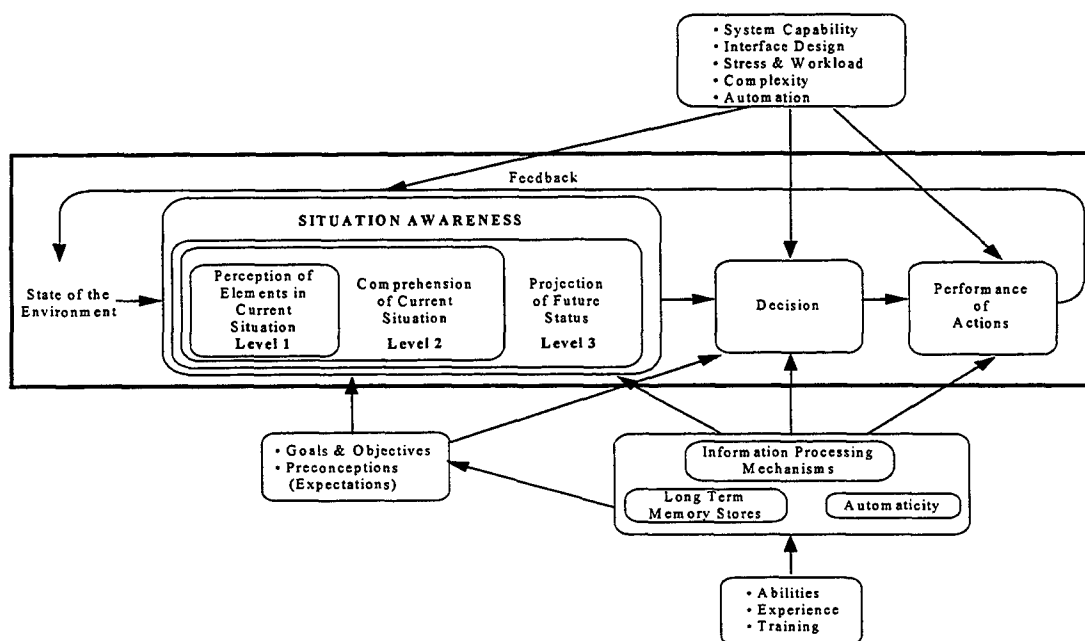


Figure 15. Model of situation awareness in dynamic decision making (Endsley, 1995, p. 35).

Individual factors affecting SA. As can be seen in Figure 15, both individual and task/system factors can influence operator SA. Endsley has expanded on the role of individual characteristics in the acquisition and maintenance of SA, as portrayed in Figure 16. The mechanisms of short-term sensory memory, perception, working memory, and long-term memory form the basic structures on which SA is based. First, the features of environmental elements are initially processed by means of preattentive sensory stores, which detect properties such as color, proximity, shape, and motion. Detection of these features provides cues for further focused attention. Because the features that are most salient are most likely to receive further processing, cue salience will be a primary determinant of which areas of the environment the operator will attend to at the first level of SA.

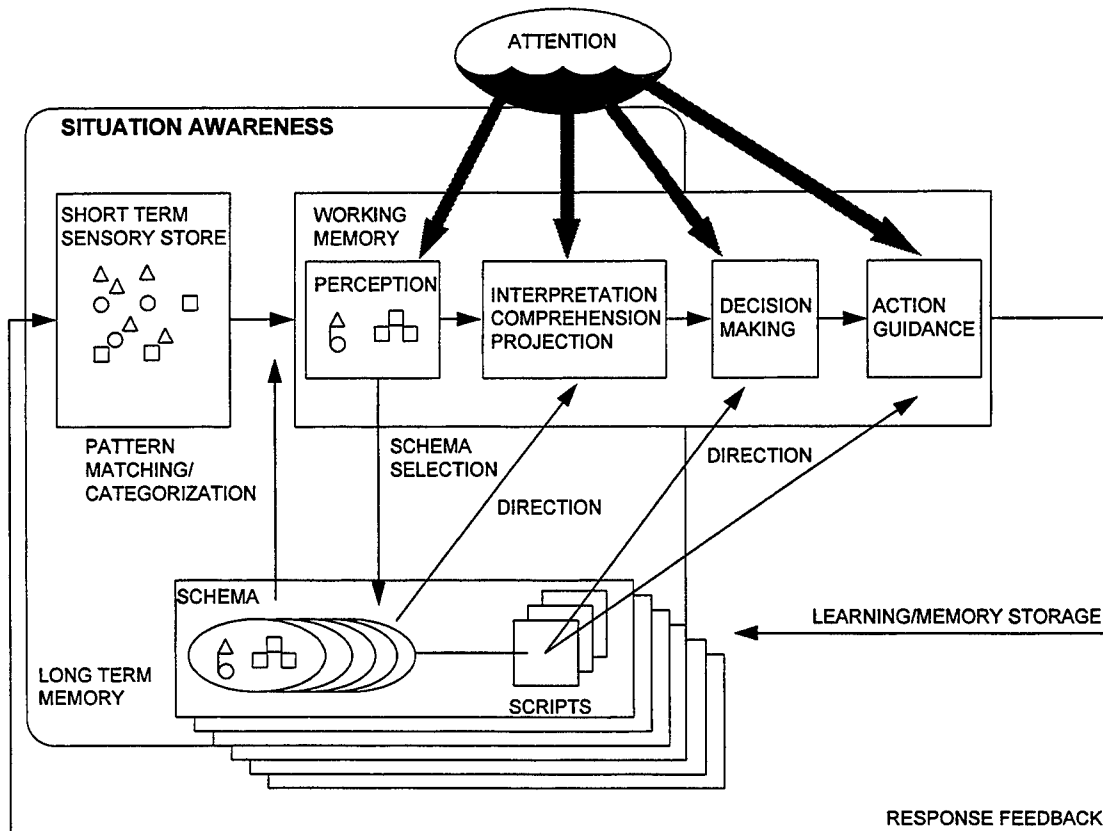


Figure 16. Mechanisms of situation awareness (Endsley, 1995, p. 41).

The manner in which elements of the environment are perceived is controlled by both working memory and long-term memory. Once perceived, information is stored in working memory. As can be seen in Figure 16, the bulk of the activity crucial to SA occurs in working memory. It is here that new information is combined with existing knowledge from long-term memory to form a coherent picture of the current situation (Level 2 SA). It is here also where projections about the future status of the system are made (Level 3 SA). Decision-making regarding the actions to be taken occur in working memory as well. Thus, a heavy load is placed on working memory during the development and maintenance of SA. However, existing knowledge in the form of schemata and scripts from long-term memory stores as well as repeated experience in a particular environment can facilitate the perception of information. Schemata provide frameworks for organizing and comprehending information efficiently. They contain general or representative information regarding a concept that can be easily retrieved and subsequently fleshed out with the details of the particular situation at hand. For example, a

“classroom” schema might contain items such as a blackboard, desks, chalk, students, a teacher, maps, and so on. These are items that one generally expects to see in a classroom, based on extensive prior experience. Retrieval of this schema would enhance the perception of items in a new classroom because there are many items one would expect to see. Even though the classroom may be different, it is still a variation on the “classroom” schema. A script is a type of schema that represents a series of appropriate actions to be taken in a given situation (e.g., a “restaurant” script). When scripts and schemata are available, the load on working memory is lessened because the relevant information can be elicited automatically. In terms of Endsley’s model of SA, schemata are primarily involved in the second and third levels of SA, whereas scripts come into play during decision-making and response selection/execution.

In addition to information processing capabilities, two other individual characteristics that can influence SA include automaticity and operator goals. Automaticity has both advantages and disadvantages for SA. Because automatic processing is fast, autonomous, and effortless, automaticity can relieve some of the burden on working memory and free up attentional mechanisms for other critical processes. When tasks become automatic, they can be achieved with minimal attention allocation. Because they are completed automatically, however, the individual subsequently has little awareness of exactly how they were completed. Consequently, one disadvantage of automaticity is a reduced responsiveness to new stimuli. Since automatic processes have taken over, the individual is not attentive to the occurrence of new or unexpected stimuli. Thus, in an atypical situation, SA can be reduced by automaticity.

Goals are also critical to the development and maintenance of SA because they serve as directors to guide the individual’s search for the information needed to meet those goals. In the context of top-down decision processing, goals and plans direct which aspects of the environment are attended to in the development of SA. With these goals in mind, the attended information is then integrated and interpreted to form Level 2 SA. At the same time, however, bottom-up processing occurs. Patterns in the environment are recognized, and these may indicate that new plans are needed to meet existing goals or that different goals are needed altogether. Thus, goals and plans may direct information-gathering processes, but they can also be modified by what is perceived.

Task and system factors affecting SA. In addition to individual factors, characteristics of the task and the system can influence operator SA. These factors include system design, interface design, stress, and workload. System and interface design are particularly critical since they are used in the perception process. Perceptions of elements in the environment, which form the basis of Level 1 SA, may come directly from the operator's senses or from system displays that first alter the information into a format more suitable for human use. Thus, deficiencies in system or interface design can affect the quality of SA. First, the system may not capture all of the information that the operator would like to have available. Second, the interface may present all of the information within its capabilities, but it may do so ineffectively in a manner ill-suited for human perception and comprehension. Hence, the way in which the information is presented can have a great impact on SA.

Physical and psychological stress and mental workload are further task determinants of SA. High levels of stress or workload can affect SA by narrowing the operator's field of attention, which may cause the operator to make decisions without fully considering all available information. They can also interfere with the functioning of working memory, where most of the activity crucial to SA occurs. High stress or workload may reduce the capacity of working memory, so that the individual cannot retain as much information as would be possible under less stressful conditions. They may also interfere with the retrieval of information from working memory.

Errors in SA. As Endsley (1995) points out, the model of SA that she has developed can be used to understand the origins of errors in SA. That is, errors can be classified as Level 1, Level 2, or Level 3 SA errors in an attempt to determine how and why they occurred. For example, Level 1 SA errors occur when the operator simply fails to perceive some relevant item in the environment that is critical for SA in the given situation. Failure to perceive critical information may result from physical characteristics of the stimuli (e.g., the intensity of a signal light may be weak and not readily discernible). On the other hand, errors of this type may occur when signals are apparent but are simply not noticed. For example, under conditions of overload, the operator may attend to those tasks that are momentarily most urgent and neglect less critical tasks (e.g., by not attending to them as often as necessary, thereby missing critical information). Finally, Level 1 SA errors may result from true misperception of information, such as misreading a *C* for an *O*.

Level 2 SA errors occur when the operator fails to comprehend the significance of the information that has been perceived. Such errors can occur if the operator is not yet experienced enough to have a richly developed mental model of the situation that is equipped to relate incoming information to knowledge based on similar past situations. In other cases, the operator may apply the wrong mental model to the situation and thus misinterpret all subsequent information. Level 2 SA errors may also occur when no mental model is available. In this instance, the operator will have to rely solely on working memory to perceive elements in the environment and attempt to make sense of them while simultaneously attending to a multitude of other tasks. Thus, errors may occur as a result of working memory overload.

Finally, Level 3 SA, which involves the projection of the status of information into the near future, may be lacking or incorrect. The operator may have perceived the information and comprehended its significance but lack the ability to understand its future implications. According to Endsley (1995), it often takes a highly developed mental model to be able to project system status accurately. In this manner, her tripartite model of SA can be used not only to understand the various components of SA but also to develop a deeper understanding of what lies at the root of many types of errors.

CONTRIBUTIONS TO INFORMATION WARFARE: UNDERSTANDING HUMAN INFORMATION PROCESSING IN THE THIRD WAVE BATTLESPACE VIA THE MODELS

Having described each model in detail, one can now begin to evaluate them in terms of their ability to enhance our understanding of human information processing in the Third Wave Battlespace. What guidelines and insights do they offer? Can they prescribe what we should and should not do in the context of IW? According to Kantowitz (1985), the ultimate criterion for the utility of any theoretical concept is its ability to predict behavior. Theories are neither right nor wrong; they are merely explanations, and some explanations may be better than others. Some may be well-suited for specific situations, whereas others may be more global in nature. Ultimately, the usefulness of any theory depends on its explanatory power. A useful model is one that can provide accurate behavioral guidelines.

Rasmussen's Skills-Rules-Knowledge Framework

Along these lines, one method for evaluating the utility of each model is an approach similar to Rasmussen's skills-rules-knowledge (SRK) framework (Harwood & Sanderson, 1986; Rasmussen, 1983, 1986). The SRK framework was developed in response to the burgeoning use of automation in modern systems, which increasingly requires humans to monitor, process, and manipulate information. As Rasmussen points out, it is crucial to be able to predict human performance and the various modes of errors that frequently occur within such systems. He further specifies that what is needed is "a set of models which is reliable for defined categories of work conditions together with a qualitative framework describing and defining their coverage and relationships" (Rasmussen, 1983, p. 257). Toward that end, his SRK framework provides several basic distinctions that are useful in defining the categories of human performance for which separate development of models is possible. The SRK framework was intended to distinguish categories of human behavior according to different ways of representing the constraints in the relationships among events in the environment and between human actions and their effects. Given this approach, three typical levels of performance emerge: skill-, rule-, and knowledge-based performance. The three levels of performance correspond to decreasing levels of familiarity with the environment or task. Further, each level can be differentiated in terms of the information that is observed from the environment.

Skill-based behavior represents sensory-motor performance during activities that occur without conscious control in highly familiar environments. Such actions appear as smooth, automated, and highly integrated patterns of behavior. Performance depends upon a very flexible and efficient dynamic internal world model. Sensory input is generally not selected or observed; rather, the senses are directed towards aspects of the environment needed subconsciously to update and orient the internal map. Thus, the constraints in the behavior of the environment at the skill level are represented by prototypical temporal-spatial patterns. The flexibility of skilled performance is due to the ability to draw upon a large number of automated subroutines to compose the sets suitable for the specific purposes. In some instances, performance is one continuous integrated dynamic whole (e.g., bicycle riding, piano playing). At the skill-based level, information from the environment is perceived as time-space *signals*, continuous quantitative indicators of the temporal and spatial behavior of the environment. The signals have no meaning except as direct physical time-space data.

Rule-based behavior occurs in a familiar work situation and is controlled by stored rules or learned procedures. Such rules may have been derived empirically on previous occasions or learned from others through a series of step-by-step instructions. Thus, performance is goal-oriented but is controlled through stored rules that are selected from previous successful experiences. The rule reflects the functional properties that constrain the behavior of the environment, but usually in properties found empirically in the past. As Rasmussen (1983) points out, the boundary between skill-based and rule-based behavior is not distinct and depends considerably on the individual's level of training and attention. In general, however, skill-based performance occurs without conscious attention, and the individual often cannot verbalize how the behavior is controlled or what information is used. With rule-based behavior, the individual can normally report the rules that are used to control performance. At the rule-based level, information from the environment is generally interpreted as *signs*. Information is defined as a sign when it serves as a cue to activate stored patterns of behavior. Signs refer to situations or to appropriate behavior based on convention or prior experience; they do not refer to concepts. Further, signs can be used only to select or modify the rules controlling the sequencing of behavior; they cannot be used for functional reasoning to generate new rules or to predict the response of a system to unfamiliar situations.

Finally, in unfamiliar situations in which expertise or stored rules from previous encounters are unavailable, the control of performance progresses to a higher conceptual level where performance is goal-controlled and *knowledge-based*. Knowledge-based behavior is characterized by planning, reasoning, and problem-solving. Here, the goal is explicitly formulated on the basis of an analysis of the environment and the individual's overall aims. Alternative plans are considered and tested (either physically or conceptually) until a useful plan is adopted. At this level of reasoning, the structure of the system is represented explicitly by some type of mental model. At the knowledge-based level, information from the environment is perceived as *symbols*. Symbols refer to concepts relating to functional properties. Unlike signals and signs, they can be used for reasoning. They are defined by and refer to the internal conceptual representation that is the basis for reasoning and planning.

Rasmussen (1983) has identified the manner in which his SRK framework can be used to facilitate the modeling of human performance in various systems. As he notes, "qualitative

models identifying categories of behavior and the limiting properties of the related human resources will serve designers a long way in the design of systems which allow humans to optimize their behavior within a proper category” (p. 264). Further, the cognitive level of the behavior under consideration will have certain implications for the modeling of behavior. At the skill-based level, individuals are highly trained and have adapted to their environment; hence, models of optimal human performance will be primarily models of the behavior of the environment. At this level of performance, quantitative models of human behavior in well-structured tasks are possible. At the knowledge-based level, on the other hand, individuals are reacting to unfamiliar situations. Thus, models that attempt to match categories of system requirements with human resources will be important. According to Rasmussen (1983), in order to be useful, both quantitative and qualitative models must reflect individuals’ mental models, the type of data being handled, and the rules or strategies used to control the processes. Further,

we do not need a single integrated quantitative model of human performance but rather an overall qualitative model which allows us to match categories of performance to types of situations. In addition, we need a number of more detailed and preferably quantitative models which represent selected human functions and limiting properties within the categories. The role of the qualitative model will generally be to guide overall design of the structure of the system including, for example, a set of display formats, while selective quantitative models can be used to optimize the detailed designs. (Rasmussen, 1983, p. 264)

Thus, in Rasmussen’s view, we need models of the present type that tend to be specific in application and scope, but we also need guidelines as to how their individual application should be coordinated. His SRK classification provides one means of capturing categories of human performance to facilitate the integration of various sub-models.

Model Evaluation

In order to evaluate the utility of the 17 human information processing models described here, one could use Rasmussen’s SRK framework to clarify what the models assert about particular categories of behavior. However, classifying the models as primarily skill-, rule-, or knowledge-based and describing their contributions solely in this regard would be somewhat

difficult. Some of the models defy classification according to Rasmussen's scheme. For example, Baddeley's model of working memory seems to have a role in all three categories of behavior. On the other hand, the AI models are primarily rule-based and little more. Further, this system would ultimately prove not to be very useful. For the purposes of IW, we have a greater need for models that specify the nature of information processing within the various stages of the OODA Loop. Consequently, the approach taken here will be similar to what Rasmussen has proposed with his SRK framework, but different categories will be used to classify the models. Specifically, the OODA decision-making framework will be used to classify the models and determine whether they apply primarily to the Observe, Orient, Decide, or Act phase, or to all four phases. Subsequently, the contributions of each model will be described with regard to this classification scheme. Where appropriate, Rasmussen's SRK terminology will be used to further describe the contributions of the models. A summary of the key features of each model will be presented first, followed by an evaluation of its contributions.

Models of Memory and Attention

The six models of memory and attention that were described included Atkinson and Shiffrin's modal model of memory, Baddeley's model of working memory, controlled and automatic information processing, attention to action, global workspace theory, and multiple resource theory. The primary characteristics of each of these models are summarized in Table 1.

With respect to the four stages of the OODA Loop, the models of memory and attention portrayed in Table 1 can be categorized under all four stages. That is, they have something to contribute about the nature of human information processing during all four stages. During the Observe phase, the individual is engaged in monitoring the situation, which involves gathering and detecting data and storing and recalling information. Nearly all of the models of memory covered here maintain the importance of working memory during this phase of the OODA cycle. Working memory will serve as the storehouse for incoming data. Hence, it will be critical to retain important information in working memory through control processes such as rehearsal so that it may be transferred to long-term memory. Because of the capacity and duration limitations of working memory, some data will be lost from memory and may need to be observed repeatedly.

Table 1

Characteristics of Six Models of Memory and Attention

Models of Memory and Attention
Atkinson and Shiffrin's Modal Model of Memory
<ul style="list-style-type: none">• Critical importance of short-term memory<ul style="list-style-type: none">• Transfers information to long-term memory• Receives information from long-term memory• Center for most of our cognitive activity• Short-term memory is limited in duration and capacity<ul style="list-style-type: none">• Information lost quickly unless actively maintained• Capacity limited to about 5 to 9 items
Baddeley's Model of Working Memory
<ul style="list-style-type: none">• Multi-component model of working memory• Articulatory loop<ul style="list-style-type: none">• Phonological store• Holds speech-based information• Duration about 2 seconds• Visuo-spatial sketchpad<ul style="list-style-type: none">• Stores visuo-spatial images• Important for geographical orientation and for planning spatial tasks• Central executive<ul style="list-style-type: none">• Controls selection, initiation, and termination of encoding, storage, and retrieval processes
Controlled and Automatic Processing
<ul style="list-style-type: none">• Memory consists of inter-related nodes<ul style="list-style-type: none">• Long-term store consists of inactive nodes• Short-term store is part of long-term store<ul style="list-style-type: none">• Portion that is currently activated• Workspace for decision-making, thinking, and control processes• Automatic processing<ul style="list-style-type: none">• Sequence of nodes in long-term store activated without active control or attention• Requires considerable training to develop• Difficult to suppress or alter once learned• Not hindered by capacity limitations of short-term store• Does not require attention for completion• Virtually unaffected by cognitive load• Familiar situations• Controlled processing<ul style="list-style-type: none">• Temporary sequence of nodes activated under individual's control and attention• Easy to establish and alter without training• Highly limited by capacity of short-term store• Only one process can be completed at a time• Requires attention• Heavily dependent on cognitive load• Underlies development of automatic processing• Novel situations

(table continues)

(Table 1 cont'd.)

Attention to Action

- Role of attention in both automatic and deliberate actions
 - Two types of control structures
 - Horizontal threads
 - Control simple or well-learned automatic actions
 - Autonomous, self-sufficient set of processing structures
 - Complete actions without conscious or attentional control
 - Vertical threads
 - Control novel or complex tasks
 - Permit conscious control of performance
 - SAS control structure
-

Global Workspace Theory

- Parallel distributed nervous system
 - No central executive
 - Global database
 - Permits information exchange among individual processors
 - A form of short-term or limited capacity working memory
 - Consciousness occurs when global workspace distributes information to all processors
-

Multiple Resource Theory

- Multi-task performance situations
 - Processing resources enable performance of a task
 - Limited availability
 - Under voluntary control
 - Performance suffers when demand exceeds supply
 - Multiple independent resource pools based on 3 dimensions
 - Stage of processing
 - Modality of input
 - Codes of processing and output
 - Timesharing is more efficient when tasks demand separate rather than common resources on any of the 3 dimensions
-

During the Orient phase, data that has been observed in the previous phase is integrated with information from the individual's existing knowledge in an attempt to form a coherent picture of the situation. Thus, the individual begins creating, evaluating, and selecting possible hypotheses to explain the state of the environment. More than one hypothesis may be applicable if there is some uncertainty or ambiguity about the data or if there is more than one cause of a problem. Hence, multiple hypotheses may be valid. Both working memory and long-term memory will have critical roles in the Orient phase. Information that is retrieved from long-term memory will be used to interpret information in working memory relating to the current situation. As hypotheses are created, they will be maintained in working memory for evaluation and selection and will therefore be subject to the capacity limitations of working memory. Thus, the individual may be expected to maintain approximately five to nine hypotheses in working memory at one time. The hypotheses will also be compared and evaluated in working memory,

which will further increase the cognitive load on the individual and make it more difficult to complete concurrent tasks.

Similarly, working memory and long-term memory will have comparable roles during the Decide and Act phases. During the Decide phase, the individual begins to create, evaluate, and select response alternatives in line with plausible hypotheses. During the Act phase, the individual plans, organizes, and executes a response. Within each of these phases, the majority of the activity will occur within working memory, with additional information being provided by long-term memory.

Thus, the modal model of memory, Baddeley's model of working memory, and the controlled and automatic processing model maintain that working memory will be highly critical during all phases of the OODA Loop. Shiffrin and Schneider's model of controlled and automatic processing further implies that most of the processing will be effortful controlled processing heavily dependent on cognitive load. Many of the situations that will be encountered in the Third Wave Battlespace will be novel situations. For example, data may have been manipulated in ways not witnessed before; information may have been distorted. Thus, the chances are low that the situation will be similar to one encountered before to which the individual can respond automatically. Consequently, because the individual will be engaged in controlled processing much of the time, performance may suffer if the task load is too great. The individual will have to devote considerable attention to the task at hand since it cannot be completed automatically. According to Rasmussen's framework, such behavior occurring in an unfamiliar situation for which stored rules from previous encounters are unavailable can be classified as knowledge-based. Automatic processing, on the other hand, would represent skill-based behaviors that can be completed automatically without conscious control in highly familiar environments. Similarly, Norman and Shallice's attention to action model holds that the vertical threads of the Supervisory Attentional System would be needed to control tasks in the novel and complex environment of the Third Wave Battlespace.

Finally, multiple resource theory provides some guidelines to help ease the burden of a heavy cognitive load imposed by multiple tasks. If possible, tasks should be presented in such a way that they draw upon separate rather than common resources. For example, a task that requires manual output can be performed more effectively with a task that requires vocal output

as opposed to one that also requires manual output. Similarly, since timesharing is more effective between modalities than within, it would be beneficial to present some information aurally and some visually.

AI Models

Five AI models of information processing were described in the present document, including the General Problem Solver, ACT-R, Soar, GOMS, and EPIC. The key features of these models are summarized in Table 2.

Table 2

Characteristics of Five AI Models

AI Models
<p>General Problem Solver</p> <ul style="list-style-type: none"> • AI model that simulates human problem-solving processes • Means-end analysis and heuristic search strategy used to achieve the Goal of transforming an initial State into a desired State via Operators <ul style="list-style-type: none"> • States are conditions of the environment or task • Operators are actions that change problem from one state to another • Goals are desired ends • Has been used to solve symbolic logic problems, Missionaries and Cannibals task, cryptarithmic problems, grammatical analyses of sentences, logic proofs, and trigonometry problems • Very limited in scope
<p>ACT-R</p> <ul style="list-style-type: none"> • AI model based on the production system • Cognitive skills are composed of production rules <ul style="list-style-type: none"> • IF-THEN clauses • Condition-Action pairs • Three types of memory <ul style="list-style-type: none"> • Working--active, current • Declarative--"what is" knowledge • Procedural--"how to" knowledge • Repetitive three-stage process for completing complex cognitive tasks <ul style="list-style-type: none"> • Match conditions of various production rules to information in working memory representing current problem • Select rule providing best match based on conflict resolution <ul style="list-style-type: none"> • Minimize computational cost • Retrieve rule that leads to goal • Fire selected rule

(table continues)

(Table 2 cont'd.)

Soar

- AI model that attempts to provide a unified theory
 - Three cognitive levels
 - Memory
 - Long-term production memory storing procedural, declarative, and episodic knowledge
 - Working memory
 - Decision
 - Two-phase elaborate-decide cycle for decision-making
 - Elaborate phase involves retrieval of all productions relevant to current problem from production memory
 - Decide phase involves selection of most acceptable and desirable production
 - Goal
 - Established whenever impasse in decision procedure is reached
 - All tasks are formulated in a problem space wherein an initial state can be transformed into a desired state via the application of operators
 - Has been applied in search-based tasks, knowledge-based tasks, robotic tasks
 - Possesses many features of general human intelligence, but still limited
-

GOMS

- AI model depicting procedural knowledge needed to complete a task
 - Goals define a state of affairs to be achieved
 - Operators are perceptual, cognitive, and motor acts that change user's mental state or environment
 - Methods are sequences of operators
 - Selection rules are used to choose among potential Methods
 - Based on rationality principle of task analysis
 - Has been used to depict computer manuscript editing
 - Currently limited to modeling error-free behavior
-

EPIC

- AI model emphasizing multiple-task performance
 - Collection of processors and memories
 - Cognitive processor
 - Working memory
 - Long-term memory for declarative information
 - Production memory for procedural knowledge
 - Production rule interpreter
 - Perceptual processors
 - Visual
 - Auditory
 - Tactile
 - Motor processors
 - Ocular
 - Vocal
 - Manual
 - Task completion occurs through execution of production rules
 - Three-phase cognitive processing cycle
 - Update contents of working memory based on previous cycle
 - Test conditions of production rules
 - Execute actions of all rules whose conditions are satisfied
 - Executive control processes
 - Ensure that all tasks are completed without conflict
 - Are handled by their own production rules
 - Cognitive processor is not limited by capacity
 - Has been used to predict reaction times and task completion times
 - Does not provide a complete theory of cognition
-

The AI models of information processing speak primarily to the Decision phase of the OODA Loop, which occurs after the individual has monitored the situation (Observe) and formed a coherent picture (Orient). The Decision phase deals with the subsequent decision-making process as to how to respond to the situation. In general, the AI models propose that an IF...THEN decision-making process is used to select a rule that best matches the conditions in working memory. Thus, the AI models would be primarily applicable to highly familiar situations where rules for behavior exist and can be retrieved to find the best match. In Rasmussen's terminology, they would be classified as rule-based models. By and large, given the unfamiliar nature of the situations that will be representative of IW, the AI models will not be applicable. They often invoke the IF...THEN processing to transform an initial state into a desired state, and these states themselves may not be clearly defined in IW. Further, many of the AI models function well only in very tightly controlled, limited domains and not in the ambiguous and complex situations of the real world.

Models of Visual Attention

The two models of visual attention described in this document included the theory of uniform connectedness and the CODE theory of visual attention. The major features of these models are depicted in Table 3.

Models of visual attention would be applicable primarily to the Observe phase of the OODA Loop when the individual is monitoring the situation and gathering and detecting data. The models of visual attention would specify what visual attention selects in the environment. According to the theory of uniform connectedness, visual attention selects UC regions--regions of the visual field with relatively homogeneous surface characteristics. According to the CODE theory of visual attention, visual attention depends on both sensory data and personal bias and priority. Thus, the distribution of visual attention might be modified if the individual has been instructed to focus on particular areas of a display or has been told to search various regions in some particular order. Unlike many of the other information processing models, the CODE theory of visual attention is quantitative rather than qualitative. Thus, it represents the sort of detailed quantitative model of a specific behavior that Rasmussen called for in his approach.

Table 3

Characteristics of Two Models of Visual Attention

Models of Visual Attention
<p>Uniform Connectedness</p> <ul style="list-style-type: none"> • Distribution of attention in space explained according to the principle of uniform connectedness (UC) • When visual field has homogeneous surface characteristics, the regions of the visual field will be perceived as single units or percepts, leading to an object-based visual selection advantage <ul style="list-style-type: none"> • Object-based performance will be enhanced when task requires processing of multiple properties of a UC region • Shape judgments do not necessarily imply object-based visual selection • Object-based visual selection does not entail mandatory processing
<p>CODE Theory of Visual Attention</p> <ul style="list-style-type: none"> • Computational model that specifies what visual attention selects by integrating space-based and object-based theories • Merges the CODE theory of perceptual grouping by proximity with Bundesen's theory of visual attention • CODE provides feature catch <ul style="list-style-type: none"> • Probability of sampling features of items in above-threshold region of CODE surface • Feature catch depends on proximities of items in display, variability of feature distributions, and threshold applied to CODE surface • Bundesen's theory of visual attention specifies how choices among perceptual inputs are made <ul style="list-style-type: none"> • Selection is made by choosing the categorization corresponding to greatest amount of sensory evidence • Can be modified by personal bias and priority for attending to particular items • Final merged model <ul style="list-style-type: none"> • CODE provides input in the form of a feature catch representing sensory data • Bundesen's theory provides values for bias and priority to permit selection of an appropriate response • Can predict reaction time and accuracy in a variety of tasks • Abstract, deals only with grouping by proximity

Language Comprehension Models

The two language comprehension models included in this review were the construction-integration model and Sanford and Garrod's model of written discourse comprehension. The basic features of these models are summarized in Table 4.

Models of language comprehension apply to the Observe and Orient phases of the OODA Loop. During the Observe phase, the gathering and detection of data may entail comprehension of written material. Further, during the Orient phase, the integration of information to form a coherent picture of the situation may require referring back to the written material or recalling it. The construction-integration model emphasizes the data driven processes involved in comprehension. Thus, word recognition begins with the word. According to this model, a number of potentially applicable meanings are generated initially. Subsequently, irrelevant material is deleted. According to Sanford and Garrod's model, the meaning of a word can be determined by searching in one of four partitions of memory. Material in either explicit or implicit focus relating to the current topic of the text will be easier and faster to access than

material outside of focus. If word meaning cannot be resolved in focus, then a slower search of the less accessible episodic or semantic memory structures will be necessary. Thus, according to this model, the attempt to interpret ambiguous terms will first occur in focus with information relating to recent material. If this search fails, a more laborious search of long-term memory will commence. This type of search process can have important implications for the presence of ambiguous terms in the context of IW. Specifically, if an individual encounters an ambiguous term, it will most likely be interpreted in terms of recent material, which may no longer be applicable and may produce misinterpretations. Further, if time is limited, long-term memory structures may not be searched at all.

Table 4

Characteristics of Two Models of Language Comprehension

Models of Language Comprehension
<p>Construction-Integration Model</p> <ul style="list-style-type: none"> • Cognitive architecture that combines production systems and connectionist approaches • Construction process <ul style="list-style-type: none"> • Weak rule-based production system • Text base is constructed from text input and comprehender's knowledge base <ul style="list-style-type: none"> • Knowledge base is an associative network • Relevant and irrelevant knowledge will be activated • Integration process <ul style="list-style-type: none"> • Integrates text base into coherent whole • Eliminates any irrelevant material • Has been used to describe word recognition in discourse and solving word arithmetic problems • Emphasizes bottom-up, perception-like aspects of comprehension
<p>Sanford and Garrod's Model of Written Discourse Comprehension</p> <ul style="list-style-type: none"> • Addresses the manner in which written discourse makes contact with knowledge in order to bring about an understanding of its meaning • Focuses on the organization of memory structures and how they assist in knowledge access during written discourse comprehension • Four partitions of memory provide four distinct domains for searching for word meaning <ul style="list-style-type: none"> • Explicit focus <ul style="list-style-type: none"> • Limited capacity store • Contains representations of items explicitly mentioned in text • Implicit focus <ul style="list-style-type: none"> • Subset of general knowledge • Corresponds to current text scenario • Episodic memory <ul style="list-style-type: none"> • Long-term memory for the discourse • Semantic memory <ul style="list-style-type: none"> • Long-term knowledge base • Explicit focus and implicit focus provide a retrieval domain representing current topic of text <ul style="list-style-type: none"> • Rapidly and easily accessed • Processing outside of focus occurs when meaning cannot be resolved in focus <ul style="list-style-type: none"> • Occurs in episodic and semantic memory partitions • Slower and less accessible

Models of Situation Awareness.

Two models of situation awareness included in this document were Adams, Tenney, and Pew's model and Endsley's model. The principal characteristics of these models are summarized in Table 5.

Table 5

Characteristics of Two Models of Situation Awareness

Models of Situation Awareness
Adams, Tenney, and Pew's Model of SA
<ul style="list-style-type: none">• Combines Neisser's view of the perceptual cycle with Sanford and Garrod's theory of written discourse comprehension• Neisser's perceptual cycle depicts interdependence of memory, perception, and action<ul style="list-style-type: none">• Knowledge leads to expectation of certain types of information• Active schemas structure subsequent flow of events by increasing receptivity to some aspects of the environment and to particular interpretations of the environment• General exploratory cycle used when information does not correspond to expectations• Sanford and Garrod's work on event comprehension incorporated to deal with multi-task aspects of SA<ul style="list-style-type: none">• Explicit and implicit focus replace Neisser's concept of the schema of the present environment• Episodic and semantic memory replace Neisser's cognitive map of the world and its possibilities• New model better equipped to handle SA in multi-task environments<ul style="list-style-type: none">• Can predict how operator will cope with tasks in the queue• Events relevant to current task can be readily assimilated because they relate to knowledge in explicit focus• If event interpretation requires inactive knowledge in long-term memory, its probability of being processed will depend largely on the available time
Endsley's Model of SA
<ul style="list-style-type: none">• SA can be described in terms of 3 phases or levels<ul style="list-style-type: none">• Level 1 involves perception of events in the environment• Level 2 involves comprehension of their meaning• Level 3 involves projection of system status in near future• SA affected by individual and task factors<ul style="list-style-type: none">• Individual factors<ul style="list-style-type: none">• Sensory memory• Perception• Working memory• Long-term memory• Automaticity• Operator goals• Task factors<ul style="list-style-type: none">• System design• Interface design• Stress• Workload• Most of the activity crucial to SA occurs in working memory• Errors can be understood in terms of which level of SA is implicated

With respect to the OODA Loop, models of situation awareness fall primarily into the Observe and Orient phases. In fact, in her model of SA, Endsley specifically points out that SA occurs prior to decision-making and action. Further, in a recent document, she has gone on to demonstrate explicitly how her model of SA fits into the OODA Loop (Endsley & Jones, in press). Endsley proposes that SA be viewed as a more detailed description of the Observe and Orient phases. Thus, Level 1 SA, which involves the perception of elements in the environment, replaces the Observe phase. Level 2 SA, which involves comprehending the meaning of the elements that are observed, and Level 3 SA, which involves the projection of the status of the system in the near future, replace the Orient phase. Thus, integrating information with what is already known to form a coherent picture of the situation involves not only understanding what has been observed but also projecting what it implies about the near future. The expansion of the Observe and Orient phases of the OODA Loop with Endsley's model of SA further implies that any information processing errors that occur may be traced to faulty Level 1, 2, or 3 SA.

Conclusions and Recommendations

The preceding evaluation of the utility of each information processing model reveals first that current cognitive models apply primarily to the Observe and Orient phases. This outcome is illustrated in Table 6, which shows the major contributions of each class of models to the four phases of the OODA Loop. Of the five classes of models, only the AI models could potentially apply to the Decide phase, and they are essentially inadequate for capturing the complexity and ambiguity of the Third Wave Battlespace. Hence, the present evaluation indicates that models applicable to the Decide and Act phases are lacking. In addition, the models that are relevant for the Observe and Orient phases do not greatly enhance our understanding of the processes that are taking place during those stages. For example, given the present evaluation of 17 models of information processing, we can now ask whether we are able to portray the Observe and Orient phases better than before. Such an attempt reveals that although the models do enlighten us somewhat, their contributions are disappointingly minor.

Table 6

Contributions of Five Classes of Cognitive Models to the Phases of the OODA Loop

COGNITIVE MODEL	PHASE OF THE OODA LOOP			
	O	O	D	A
Memory and Attention	<ul style="list-style-type: none"> Working memory will be the critical storehouse for incoming data Capacity and duration limitations may be a factor in performance 	<ul style="list-style-type: none"> Working memory and long-term memory will be important in the interpretation of data and formation of hypotheses 	<ul style="list-style-type: none"> Working memory and long-term memory will be critical in the creation, evaluation, and selection of possible responses to the situation 	<ul style="list-style-type: none"> Working memory and long-term memory will be important in planning, organizing, and executing a response
AI Models	--	--	<ul style="list-style-type: none"> IF...THEN processing for response selection will be inadequate in the complex and unfamiliar environment of IW 	--
Visual Attention	<ul style="list-style-type: none"> Visual attention will select relatively homogenous areas of surface features in visual field Sensory data as well as priority and personal bias will be important in determining what visual attention selects 	--	--	--
Language Comprehension	<ul style="list-style-type: none"> Many different meanings of ambiguous written material may be generated 	<ul style="list-style-type: none"> Ambiguities may be resolved by referring to the context of recent material, which may lead to misinterpretations 	--	--
Situational Awareness	<ul style="list-style-type: none"> Perception of events in the environment 	<ul style="list-style-type: none"> Comprehension of the meaning of events in the environment Projection of their status in the near future 	--	--

During the Observe phase when the individual is engaged in monitoring the situation and gathering and detecting data, a huge amount of data will be coming in and will need to be processed rapidly and efficiently. The particular source of information the individual chooses to attend to at any given time will be determined not only by sensory input but also by priority and personal bias. Much of the activity needed to cope with this input will take place in working memory and will therefore be prey to this structure's duration and capacity limitations. If cognitive load is heavy due to the need to perform multiple concurrent tasks, performance may suffer. The information processing itself will most likely be effortful controlled processing that will be highly demanding of attention and conscious control since many of the tasks will likely be novel and complex. The multi-task burden may be eased somewhat by attempting to utilize multiple resource pools to distribute the cognitive load.

During the Orient phase, the individual will attempt to form a coherent picture of the situation by linking current information in working memory with prior knowledge in long-term memory to make sense of it. At this stage, the individual begins to create, evaluate, and select different hypotheses to account for the overall picture. As in the Observe phase, much of the activity will take place in working memory and be subject to its capacity and duration limitations. The individual will most likely attempt to resolve ambiguities in the data by relying on recent information, which may lead to misinterpretations if recent material is no longer relevant. If time is limited, the individual may forego a more laborious search of long-term memory altogether.

In general, the information processing models reviewed here point to the critical importance of working memory and the need to be wary of overburdening it, given its duration and capacity limitations. One problem with applying many of the models is their breadth. In attempting to capture all of human information processing, they have become too general. Along the lines suggested by Rasmussen, perhaps it would be more profitable to develop models of specific processes that can then be placed into the overall framework provided by the OODA Loop.

REFERENCES

- Adams, M. J., Tenney, Y. J., & Pew, R. W. (1995). Situation awareness and the cognitive management of complex systems. *Human Factors*, 37, 85-104.
- Anderson, J. R. (1983). *The architecture of cognition*. Cambridge: Harvard.
- Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Lawrence Erlbaum.
- Anderson, J. R. (1993). *Rules of the mind*. Hillsdale, NJ: Lawrence Erlbaum.
- Anderson, J. R., Kushmerick, N., & Lebiere, C. (1993). Navigation and conflict resolution. In J. R. Anderson (Ed.), *Rules of the mind* (pp. 93-119). Hillsdale, NJ: Lawrence Erlbaum.
- Atkinson, R. C., & Shiffrin, R. M. (1968). Human memory: A proposed system and its control processes. In K. W. Spence & J. T. Spence (Eds.), *The psychology of learning and motivation: Advances in research and theory* (Vol. 2, pp. 89-195). New York: Academic Press.
- Baars, B. J. (1983). Conscious contents provide the nervous system with coherent, global information. In R. J. Davidson, G. E. Schwartz, & D. Shapiro (Eds.), *Conscious and self-regulation: Advances in research and theory* (pp. 41-79). New York: Plenum.
- Baddeley, A. (1986). *Working memory*. Oxford: Clarendon Press.
- Baddeley, A. (1990). *Human memory: Theory and practice*. Boston: Allyn and Bacon.
- Baddeley, A. D., & Hitch, G. (1974). Working memory. In G. H. Bower (Ed.), *The psychology of learning and motivation* (pp. 47-89). New York: Academic Press.
- Barsalou, L. W. (1995, March/April). ACT-R: Cognitive theory's simple, elegant tool. *American Scientist*, 83, 185-186.
- Boyd, J. R. (1987, August). *A discourse on winning and losing*. Unpublished briefing (Report no. MU-43947), Air University Library, Maxwell AFB.
- Bundesen, C. (1990). A theory of visual attention. *Psychological Review*, 97, 523-547.
- Card, S. K., Moran, T. P., & Newell, A. (1983). *The psychology of human-computer interaction*. Hillsdale, NJ: Lawrence Erlbaum.
- Chong, R. S. (1995, February). *EPIC-Soar: Perception, cognition, and action: One small step towards unification* [On-line]. Available <http://ai.eecs.umich.edu/people/rchong/epic-soar.html>

Cohen, A., & Ivry, R. (1989). Illusory conjunctions inside and outside the focus of attention. *Journal of Experimental Psychology: Human Perception and Performance*, 15, 650-663.

Compton, B. J., & Logan, G. D. (1993). Evaluating a computational model of perceptual grouping by proximity. *Perception and Psychophysics*, 53, 403-421.

DiNardo, R. L., & Hughes, D. J. (1995). Some cautionary thoughts on information warfare. *Airpower Journal, Winter 1995*, 1-10.

Endsley, M. (1987). *SAGAT: A methodology for the measurement of situation awareness* (NOR DOC 87-83). Hawthorne, CA: Northrop Corp.

Endsley, M. (1988). Design and evaluation for situation awareness enhancement. In *Proceedings of the Human Factors Society 32nd Annual Meeting* (pp. 97-101). Santa Monica, CA: Human Factors and Ergonomics Society.

Endsley, M. (1995). Toward a theory of situation awareness in dynamic systems. *Human Factors*, 37, 32-64.

Endsley, M. R., & Jones, W. M. (in press). Situation awareness, information dominance, and information warfare (Report No. AL/CF-TR-1997-). Wright-Patterson AFB, OH: Armstrong Laboratory.

Ernst, G. W., & Newell, A. (1969). *GPS: A case study in generality and problem solving*. New York: Academic Press.

Harwood, K., & Sanderson, P. (1986). Skills, rules and knowledge: A discussion of Rasmussen's classification. In *Proceedings of the Human Factors Society 30th Annual Meeting* (pp. 1002-1006). Santa Monica, CA: The Human Factors Society.

Isreal, J. (1980). *Structural interference in dual task performance: Behavioral and electrophysiological data*. Unpublished doctoral dissertation, University of Illinois, Champaign.

Jensen, O. W. (1994). Information warfare: Principles of third-wave war. *Airpower Journal, Winter 1994*, 35-43.

Joint Chiefs of Staff (1996). *Joint Vision 2010. America's Military: Preparing for Tomorrow*. Washington, D. C.: Office of the Chairman, JCS, 1996.

Kahneman, D. (1973). *Attention and effort*. Englewood Cliffs, NJ: Prentice-Hall.

Kantowitz, B. H. (1985). Channels and stages in human information processing: A limited analysis of theory and methodology. *Journal of Mathematical Psychology*, 29, 135-174.

- Kieras, D. E., & Meyer, D. E. (1994). *The EPIC architecture for modeling human information-processing and performance: A brief introduction* (Tech. Rep. No. 94/ONR-EPIC-1). Ann Arbor: University of Michigan, Division of Research Development and Administration.
- Kintsch, W. (1988). The role of knowledge in discourse comprehension: A construction-integration model. *Psychological Review*, 95, 163-182.
- Kintsch, W. (1992a). A cognitive architecture for comprehension. In H. L. Pick, Jr., P. Van Den Broek, & D. C. Knill (Eds.), *Cognition: Conceptual and methodological issues* (pp. 143-163). Washington, D. C.: American Psychological Association.
- Kintsch, W. (1992b). How readers construct situation models for stories: The role of syntactic cues and causal inferences. In A. F. Healy, S. M. Kosslyn, & R. M. Shiffrin (Eds.), *From learning processes to cognitive processes: Essays in honor of William K. Estes* (Vol. 2, pp. 261-278). Hillsdale, NJ: Lawrence Erlbaum.
- Kintsch, W. (1994a). Discourse processing. In G. d'Ydewalle, P. Eelen, & P. Bertelson (Eds.), *International perspectives on psychological science. Vol. 2: The state of the art* (pp. 135-155). Hove, UK: Lawrence Erlbaum.
- Kintsch, W. (1994b). The psychology of discourse processing. In M. A. Gernsbacher (Ed.), *Handbook of psycholinguistics* (pp. 721-739). San Diego: Academic Press.
- Kramer, A. F., & Watson, S. E. (1996). Object-based visual selection and the principle of uniform connectedness. In A. F. Kramer, M. G. H., Coles, & G. D. Logan (Eds.), *Converging operations in the study of visual selective attention* (pp. 395-414). Washington, D. C.: American Psychological Association.
- Laird, J. E., Newell, A., & Rosenbloom, P. S. (1987). SOAR: An architecture for general intelligence. *Artificial Intelligence*, 33, 1-64.
- Logan, G. D. (1995). Linguistic and conceptual control of visual spatial attention. *Cognitive Psychology*, 28, 103-174.
- Logan, G. D. (1996). The CODE theory of visual attention: An integration of space-based and object-based attention. *Psychological Review*, 103, 603-649.
- Meyer, D. E., & Kieras, D. E. (1997). A computational theory of executive cognitive processes and multiple-task performance: Part I. Basic mechanisms. *Psychological Review*, 104, 3-65.
- Miller, G. A., & Frick, F. C. (1949). Statistical behavioristics and sequences of responses. *The Psychological Review*, 56, 311-324.

- Moscovitch, M., & Klein, D. (1980). Material-specific perception for visual words and faces. *Journal of Experimental Psychology: Human Perception and Performance*, 6, 590-603.
- Navon, D. (1984). Resources--A theoretical soup stone? *Psychological Review*, 91, 216-234.
- Navon, D., & Gopher, D. (1979). On the economy of the human-processing system. *Psychological Review*, 86, 214-255.
- Neisser, U. (1976). *Cognition and reality: Principles and implications of cognitive psychology*. San Francisco: Freeman.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Newell, A., & Simon, H. A. (1963). GPS, a program that simulates human thought. In E. A. Feigenbaum & J. Feldman (Eds.), *Computers and thought* (pp. 279-293). New York: McGraw-Hill.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
- Norman, D. A., & Bobrow, D. G. (1975). On data-limited and resource-limited processes. *Cognitive Psychology*, 7, 44-64.
- Norman, D. A., & Shallice, T. (1986). Attention to action: Willed and automatic control of behavior. In R. J. Davidson, G. E. Schwartz, & D. Shapiro (Eds.), *Consciousness and self-regulation: Advances in research and theory* (Vol. 4, pp. 1-18). New York: Plenum Press.
- Pierce, J. R. (1980). *An introduction to information theory: Symbols, signals, and noise* (2nd ed.). New York: Dover.
- Rasmussen, J. (1983). Skills, rules, and knowledge; signals, signs, and symbols, and other distinctions in human performance models. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-13, 257-266.
- Rasmussen, J. (1986). *Information processing and human-machine interaction: An approach to cognitive engineering. Series Vol. 12*. New York: North-Holland.
- Regal, D. M., Rogers, W. H., & Boucek, G. P. (1988). Situational awareness in the commercial flight deck: Definition, measurement, and enhancement. In *Proceedings of the Seventh Aerospace Behavioral Technology Conference and Exposition* (pp. 65-69). Warrendale, PA: Society of Automotive Engineers.

- Rosenbloom, P. S., Laird, J. E., Newell, A., & McCarl, R. (1991). A preliminary analysis of the Soar architecture as a basis for general intelligence. *Artificial Intelligence*, 47, 289-325.
- Sanders, M. S., & McCormick, E. J. (1987). *Human factors in engineering and design* (6th ed.). New York: McGraw-Hill.
- Sanford, A. J., & Garrod, S. C. (1981). *Understanding written language*. New York: Wiley.
- Shannon, C. E., & Weaver, W. (1949). *A mathematical theory of communication*. Urbana: The University of Illinois Press.
- Shiffrin, R. M. (1993). Short-term memory: A brief commentary. *Memory & Cognition*, 21, 193-197.
- Shiffrin, R. M., & Atkinson, R. C. (1969). Storage and retrieval processes in long-term memory. *Psychological Review*, 76, 179-193.
- Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending, and a general theory. *Psychological Review*, 84, 127-190.
- Toffler, A., & Toffler, H. (1991). War, wealth, and a new era in history. *World Monitor*, 4 (5), 46-52.
- Toffler, A., & Toffler, H. (1993). *War and anti-war: Survival at the dawn of the 21st century*. New York: Little, Brown, and Co.
- van Oeffelen, M. P., & Vos, P. G. (1982). Configurational effects on the enumeration of dots: Counting by groups. *Memory and Cognition*, 10, 396-404.
- van Oeffelen, M. P., & Vos, P. G. (1983). An algorithm for pattern description on the level of relative proximity. *Pattern Recognition*, 16, 341-348.
- Whitaker, R. D., & Kuperman, G. G. (1996). *Cognitive engineering for information dominance: A human factors perspective* (Report No. AL/CF-TR-1996-0159). Wright-Patterson Air Force Base, OH: Armstrong Laboratory.
- Wickens, C. D. (1976). The effects of divided attention on information processing in tracking. *Journal of Experimental Psychology: Human Perception & Performance*, 2, 1-13.
- Wickens, C. D. (1984). Processing resources in attention. In R. Parasuraman & D. R. Davies (Eds.), *Varieties of attention* (pp. 63-102). Orlando, FL: Academic Press.
- Wickens, C. D. (1992). *Engineering psychology and human performance* (2nd ed.). New York: Harper Collins.

Wickens, C. D., & Kessel, C. (1980). The processing resource demands of failure detection in dynamic systems. *Journal of Experimental Psychology: Human Perception & Performance*, 6, 564-577.

Wickens, C. D., Sandry, D., & Vidulich, M. (1983). Compatibility and resource competition between modalities of input, output, and central processing. *Human Factors*, 25, 227-248.

Widnall, S. E., & Fogleman, R. R. (1995, October). *Cornerstones of Information Warfare*. Washington, D. C.: Dept. of the Air Force.

Wohl, J. G. E., Entin, E., & Eterno, J. S. (1983). *Modeling human decision processes in command and control* (Report no. TR-137). Burlington, MA: Alphatech, Inc.

GLOSSARY

ACT	Adaptive Control of Thought
ACT-R	Adaptive Control of Thought--Rational
AI	Artificial Intelligence
CM	Controlled Mapping
CODE	COntour DETector
CTVA	CODE Theory of Visual Attention
EPIC	Executive Process-Interactive Control
GOMS	Goals, Operators, Methods, Selection Rules
GPS	General Problem Solver
IDW	Information Dominance Warfare
ISW	Information Systems Warfare
IW	Information Warfare
ms	Millisecond
SA	Situation Awareness
SAS	Supervisory Activating System
SHOR	Stimulus-Hypothesis-Options-Response
Soar	State, Operator And Result
SRK	Skill-Rule-Knowledge framework
UC	Uniform Connectedness
VM	Varied Mapping
WME	Working Memory Element