HUMAN FACTORS, 1995, 37(1), 137–148

Situation Awareness Is Adaptive, Externally Directed Consciousness

KIP SMITH[1] *and* P. A. HANCOCK, *University of Minnesota, Minneapolis, Minnesota*

We define situation awareness (SA) as adaptive, externally directed consciousness. This definition dispels the artificial and contentious division evident in the literature, according to which SA is either exclusively knowledge or exclusively process. This misdirected rivalry has more to do with general perspectives on the study of human behavior than with SA itself. Through defining SA as an aspect of consciousness, we hope to clarify two key issues. (1) The source of goals with respect to SA is a normative arbiter in the task environment; that is, the behavior that SA generates must be directed at an external goal. (2) SA is the invariant at the core of the agent's perception-action cycle that supports skilled performance; that is, relationships among factors or dimensions in the environment determine what the agent must know and do to achieve the goals specified by the external arbiter. We introduce a construct we call the *risk space* to represent the invariant relations in the environment that enable the agent to adapt to novel situations and to attain prespecified goals. We articulate this concept of a risk space through use of a specific example in commercial aircraft operations. The risk space structures information about the physical airspace in a manner that captures the momentary knowledge that drives action and that satisfies the goals and performance criteria for safe and efficient flight. We note that the risk space may be generalized to many different means of navigation.

INTRODUCTION

In the early part of this century, *consciousness* became an unacceptable term in the psychologist's lexicon. In reaction to the unquantifiable subjectivity of introspectionism, Watson (1913) rewrote the course of psychology by excising the "mental" from what is literally the "science of mental life." In the latter part of this century, the concept of consciousness has enjoyed a gradual rehabilitation, albeit cloaked in many guises. Constructs such as attention, mental workload, and now situation awareness (SA) have arisen for consideration. This renaissance is attributable in part to attempts to expand understanding of behavioral abilities beyond simple stimulus-response relationships to include humans performing complex tasks in dynamic environments.

Although human factors has found particular value in each of these constructs, consciousness has yet to be fully reinstated in explanations of an individual's behavior in human-machine systems. Therefore, although the main topic here is situation awareness, our definition of SA and its relation to consciousness has broader implications for the field of human factors and the study of behavior in general.

[1] Requests for reprints should be sent to Kip Smith, Human Factors Research Laboratory, University of Minnesota, 141 Mariucci Arena, 1901 Fourth St., S.E., Minneapolis, MN 55455.

Our first step is to propose a definition of SA. Second, we characterize SA in terms of the perception-action cycle introduced by Neisser (1976) and as applied by Tenney, Adams, Pew, Huggins, and Rogers (1992). Third, we discuss SA in the contexts of adaptation (Simon, 1982) and of the distinction between competence and performance (Anderson, 1990; Chomsky, 1965; Marr, 1982). Theories of competence specify what must be known in order to solve an information processing problem. Our analysis places SA firmly within the abstract level of explanation that characterizes theories of competence.

We argue that SA specifies what must be known to solve a class of problems posed when interacting with a dynamic environment. Further, we suggest that the existing literature frequently confounds competence with performance. This confound may be responsible for the current failure of SA to gain wide acceptance beyond a cadre of aviation specialists.

We then provide a concrete example of our characterization of SA that uses a novel representation of the en route environment for commercial air traffic operations. This "risk space" representation allows us to illustrate our conception of SA in a practical realm. We argue that the risk space captures the invariant in the perception-action cycle that guides the adaptive behavior of both pilots and air traffic controllers. We indicate how the risk space embodies an operational definition of SA that is amenable to empirical testing and note that it can be used in navigational realms beyond aviation operations.

The final section contrasts our definitions, both theoretical and operational, with the notion of SA as a mental model, which has been propounded by others (e.g., Endsley, 1988). We use this contrast to illustrate both the generality of the position we take and the power of the representation we offer.

SITUATION AWARENESS DEFINED

We propose that SA be defined as adaptive, externally directed consciousness. We take consciousness to be that part of an agent's knowledge-generating behavior that is within the scope of intentional manipulation. As shown in Figure 1, we view SA as generating purposeful behavior (behavior directed toward achieving a goal) in a specific task environment. The products of SA are knowledge about and directed action within that environment. We argue that SA is more than performance. More fundamentally, it is the capacity to direct consciousness to generate competent performance given a particular situation as it unfolds. Given our definition, it is not surprising that we regard SA as directly related to stress, mental workload, and other energetic constructs that are facets of consciousness (Freeman, 1948).

Adaptation to Externally Defined Goals

We see a sequential link from consciousness to SA to adaptation. Adaptation is the process by

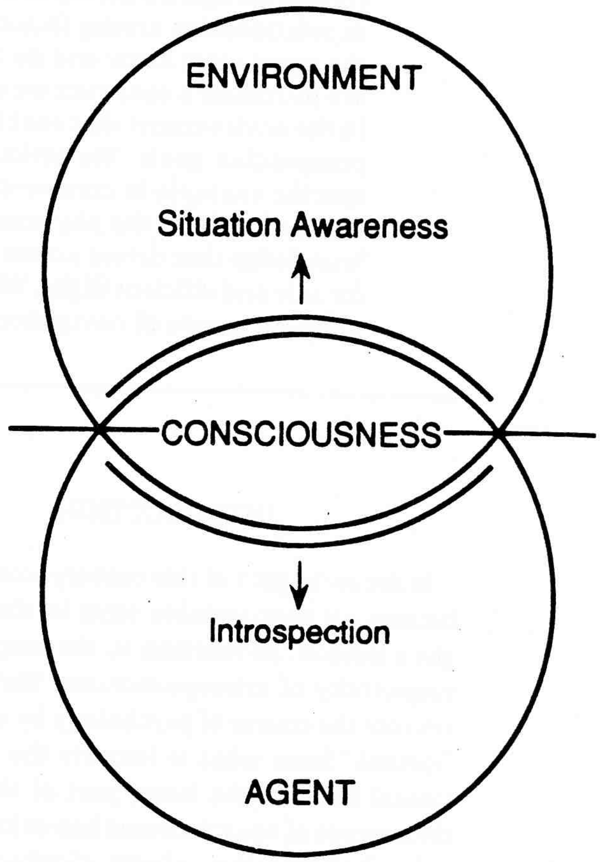

Figure 1. *An approach to defining situation awareness (SA) through explicit recognition of the centrality of externally oriented consciousness. The central (horizontal) line provides an arbitrary distinction between exogenous and endogenous orientations of consciousness and represents a distinction between SA and introspection.*

which an agent channels its knowledge and behavior to attain goals as tempered by the conditions and constraints imposed by the task environment (Holland, 1975/1992). Adaptation ensures that the agent's behavior and goals match the information made available and activity made possible by its environment (Simon, 1982).

SA, like adaptation, is a dynamic concept that exists at the interface between the agent and its environment. The study of SA requires assessment of the environment-agent relationship. As with adaptation, SA presumes experience in an environment and the development of an armory of appropriate alternative courses of action (Holland, 1975/1992).

We submit that SA presumes adaptation of a particular kind. As expressed in aviation, and as all skilled practitioners know, SA is about having the "right stuff." This notion implies complete and natural adherence to task goals and to criteria for performance. This, in turn, implies the existence of a specification of the task the agent is to perform and of measures for evaluating that performance. To possess SA—to have the right stuff—the agent must have developed a level of adaptive capability sufficient to match the specification of task goals and the criteria for assessing performance variables. Thus SA is adaptation to a singular source of constraint: a normative arbiter that defines the stuff that is right.

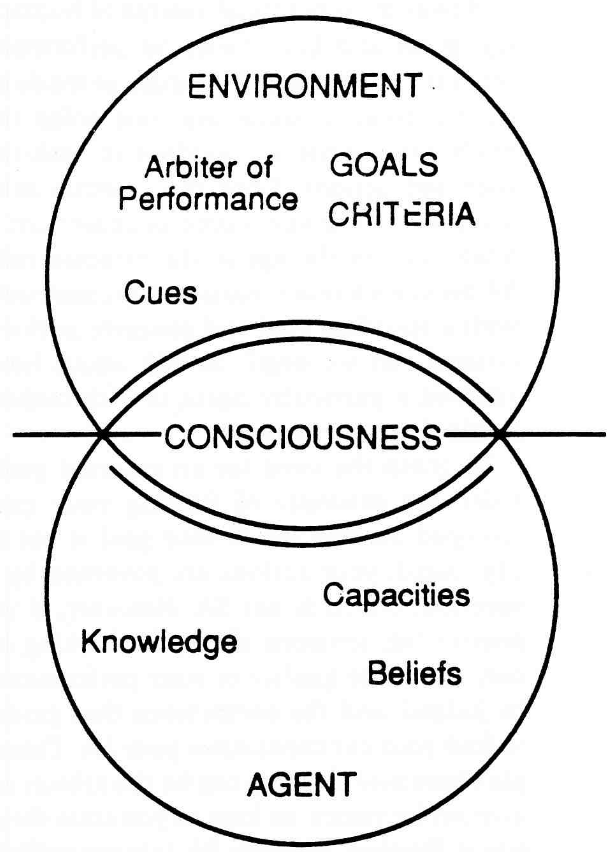

Figure 2. *Constraints on SA. The singular constraint is the presence of a normative arbiter of performance in the agent's task environment. The arbiter specifies for the agent task-relevant constraints and criteria for performance. Adaptation to the environment requires the agent to adopt the arbiter's specification of constraints and performance variables. Cues and demands are stimuli that unfold in the environment. The agent's internal constraints are those that shape its intentionality.*

As shown in Figure 2, we see the arbiter and its dicta residing in the task environment. A real problem in the current formulations of SA is the failure to articulate the presence in the environment of normative specifications and criteria for the performance of the agent's task. Although individuals may exhibit situated, outwardly directed consciousness, it is not until the externally defined task is made explicit that their behavior achieves the status we wish to reserve for SA. To qualify for SA, the agent first must intend its goals, beliefs, and knowledge to match the task and performance specified by dicta from its environment and, then, must succeed to some degree in meeting those expectations.

In our definition, we intend the phrase *externally directed* to indicate that the goal of the behavior that SA directs must reside in the task environment rather than in the agent's head. Failure to recognize the role of the normative arbiter of performance has clearly been a source of confusion in the literature on SA. Until an external goal and criteria for achieving it are specified, examination of greater or lesser degrees of SA or even of loss of SA remains impossible, being the same cul-de-sac in which introspectionism found itself. If the agent were to dictate private, incontestable (but dynamic) goals, SA would always be perfect because whatever was perceived would be that goal.

However, in practical realms of human activity, goals and boundaries on performance are often set by others or by nature or made explicit by ourselves at some previous point in time. Such constraints are evident in task-relevant cues and actions. Once the external arbiter is accepted as the key source of constraint on the adaptation of the agent, the directed nature of SA becomes a more manageable construct. Only with a specified goal and concrete performance criteria can we begin to talk about how well adapted a particular agent is with respect to a particular environment.

To grasp the need for an external goal, consider the example of finding your car in a crowded parking lot. If your goal is not explicitly stated, your actions are governed by introspection, which is not SA. However, if you explicitly tell someone else, "I'm looking for my car," then the quality of your performance can be judged, and the competence that guides you to find your car constitutes your SA. This example illustrates that you can be the arbiter of your own performance, as long as you state the goal a priori. Further, to assess SA, it is not sufficient to think you understand the actions of another. Those actions must be assessed in the light of explicitly stated goals.

As it is the relationship between organism and environment that is crucial, SA is appropriately discussed within the framework of the ecological movement (e.g., Gibson, 1979; see also Flach, Hancock, Caird, and Vicente, 1995; Hancock, Flach, Caird, and Vicente, 1995). The ecological approach affirms the importance of the interaction of an agent with its environment in the mutual shaping of the agent's actions. The approach elevates this interaction to make it the "unit of analysis" in any assessment of adapted behavior. Explanation is found in the interaction of agent and environment rather than in the agent (human or machine), separately considered. Given this perspective, to comprehend SA without a viable understanding of the interaction between agents and their task environment would be virtually impossible. An ecological view emphasizes that if SA is continued to be studied with the techniques that focus overwhelmingly on the agent as the individual unit of concern, the critical invariants in the agent-environment interaction are about as likely to emerge as is the *Encyclopaedia Britannica* from the combined efforts of Sir Arthur Eddington's anthropoid typists (but see Koestler, 1978).

To return to our definition, an agent's SA generates knowledge and action, given the structure of its environment and the goals and performance criteria specified in that environment. The explicit use of the term *given* reveals our conviction that the root of SA is the adaptive capacity to guide behavior in response to dynamic situations. It is important to note that this definition denies a claim of SA to any agent who is merely conscious and attending to its environment. Further, it denies a claim of SA to an agent that may be fully capable of SA but that is not actively pursuing the goal specified by the arbiter of performance. Rather, to stake a claim to SA, an agent must be seeking information and taking action in pursuit of an externally specified goal. Without the normative focus of an externally specified goal, SA degenerates into introspection.

Competence, Performance, and Situation Awareness

Adapted, externally directed consciousness (SA) endows the agent with the competence to generate appropriate behavior in response to dynamic situations. The distinction between competence and performance is a persistent theme in the study of cognition (Anderson, 1990; Chomsky, 1965; Marr, 1982). Competence directs behavior but is independent of the situation. Competence melds knowledge, abilities, and cognition to generate appropriate behavior given conditions in the task environment. Competence resides in the agent and enables the agent to generate skilled performance. Performance is action situated in the world, a momentary phenomenon that is guided by competence but must be distinguished from it. Performance is contingent on information available in the environment; competence is independent of the particulars of a situation. The utility of this distinction is the leverage that an analysis of competence provides to developing an understanding of the normatively focused behavior that characterizes SA.

An analysis of competence asks a simple question (Marr, 1982): What is the problem for which this agent's behavior is the solution? Specification of the problem focuses on the agent's knowledge and goals, the information available in the environment, and the actions the agent may take to meet such goals. Specification of competence details what the agent must know and do to attain its goal (to solve the problem). An analysis of competence is unconcerned with the actual processes (e.g., representations, mental models) that produce the agent's performance. As Newell (1981) pointed out, the resulting description of competence places constraints on behavior but does not prescribe performance. Full prescription or emulation of the agent's behavior requires a complete accounting of the agent's representation and processing.

The problem to be solved by SA is coming to know what must be known in order to behave in accord with the mandates of the arbiter of performance. SA is the competence that directs the agent's sampling of factors in the environment that determine how the agent can come to know what it must do. In contrast, the behavior that competence generates is the solution to a specific problem. The behavior that SA generates is the agent's solution to the problem of knowing those cues and demands in the environment that enable it to take action that aligns with the dicta of the arbiter of performance.

Tenney et al. (1992) proposed that Neisser's (1976) perception-action cycle provides a framework for understanding how SA works. In essence, an agent and its environment interact in a manner to satisfy the arbiter by generating skilled performance. We agree. The perception-action cycle is shown in Figure 3. Information and action flow continuously around the cycle. Starting arbitrarily at the top, the environment informs the agent, modifying its knowledge. Knowledge directs the agent's activity in the environment. That activity samples and perhaps anticipates or alters the environment, which in turn informs the agent. The informed, directed sampling and/or anticipation capture the essence of behavior characteristic of SA.

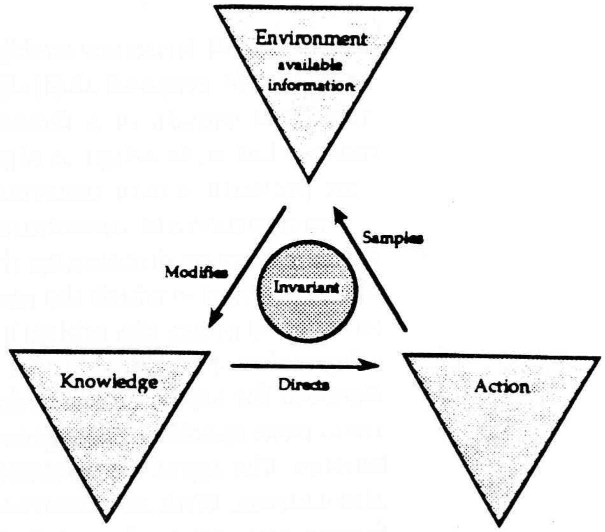

Figure 3. *Neisser's (1976) perceptual cycle. The invariant at the core of the cycle specifies the agent's adaptation to its environment. It structures the information made available by the environment, the agent's knowledge, and the actions the agent takes to meet the constraints specified by the arbiter of performance.*

To go beyond performance, to capture the competence that is SA, we include in Figure 3 the invariant that links the elements of the perception-action cycle. The invariant is the structure of the agent's adaptation to the environment: It forms the linkage among information, knowledge, and action that produces competent behavior. Specifically, the invariant codifies the information that the environment may make available, the knowledge the agent requires to assess that information, and the action the knowledge will direct the agent to take to attain its goals. We propose therefore that SA is specified by the invariant at the core of an adapted agent's perception-action cycle.

Some readers may be puzzled by the fact that we conjoin adaptability and invariance. We contend that there is no incompatibility. Rather, it is the invariant structure in the agent's interaction with its environment that generates up-to-the-minute knowledge and appropriate action.

The invariant structure enables the agent to anticipate and respond skillfully to novel conditions and signals in a dynamic task environment—that is, to adapt to high rates of change, time pressure, and/or uncertainty.

The environment contributes to the invariant those factors or dimensions that define sources of information to which the agent must attend to in order to attain the arbiter's goal. The agent's adaptation contributes the knowledge and methods for assessing that information. The environment specifies appropriate alternative behaviors. The agent's adaptation generates those alternatives. With its consciousness directed at factors relevant to the arbiter's goal, the agent modifies its knowledge in a manner that affords appropriate action (e.g., Gibson, 1979). Although each situation that the agent faces is likely to present a novel array of information, the invariant relations that exist in each array guide successful problem solving and skilled performance of the task.

By defining SA as a generative process of knowledge creation and informed action taking, we expressly deny that SA is merely a snapshot of the agent's current mental model. Rather, SA guides the process of modifying knowledge—that is, of constructing a representation of current and likely events. The experience of air traffic controllers "losing the picture" illustrates this point. As Hopkin (1988) recounted, the job of the controllers is to construct knowledge of their sectors and to take action on the basis of this knowledge. They call the knowledge they generate "the (big) picture." At times, they "lose the picture"—their knowledge becomes insufficient to support their task. Experienced controllers on the job are often sufficiently self-aware to recognize they are losing the picture as it happens.

This metaknowledge argues our point: SA is not the controllers' picture. Rather, it is the controllers' SA that builds the picture and that enables them to know that what they know is insufficient for the increasing demands. SA not only supports the construction of the picture but also guides the assessment of its integrity.

THE RISK SPACE REPRESENTATION OF THE FLIGHT DECK/ATC INVARIANT

Commercial aviation may epitomize a task that qualifies for SA. The task environment contains clear and unequivocal normative arbiters of performance: airlines and the Federal Aviation Administration (FAA, and its counterparts around the world). The agents are pilots and controllers. The arbiter establishes the agents' goals and performance criteria for safe and efficient flight. In what follows we present a hypothesis about the invariant structure of the aircraft-airspace system that guides skilled performance on the flight deck and in the air traffic control (ATC) suite.

Our candidate representation for the invariant in commercial aviation is a multidimensional *risk space* (Smith and Hancock, 1992). The risk space is a generalization of a mathematical formulation of human performance in process control tasks (Moray, 1986; Phatak and Bekey, 1969). The axes of any risk space are those factors in the environment that compromise safety. These safety factors define sources of information to which the agent attends in order to satisfy the arbiter's norms for performance. Thresholds of safety parse the risk space into decision regions. The thresholds are performance criteria defined by either the arbiter of performance or the agent who differentiates alternative control decisions. Each decision region is associated with a unique set of appropriate alternatives (actions). Information generated by the task environment (the airspace) is continually posted and updated within the risk space. Figure 4 schematically presents a portion of the flight deck/ATC risk space.

Two safety factors—horizontal separation between aircraft and relative velocity at altitude—define the axes of the portion of the ATC risk space sketched in Figure 4. Maintaining separation is one of the FAA's prime directives. Avoiding high-velocity convergence is a paramount concern. Separation and relative velocity determine the information that both pilots and

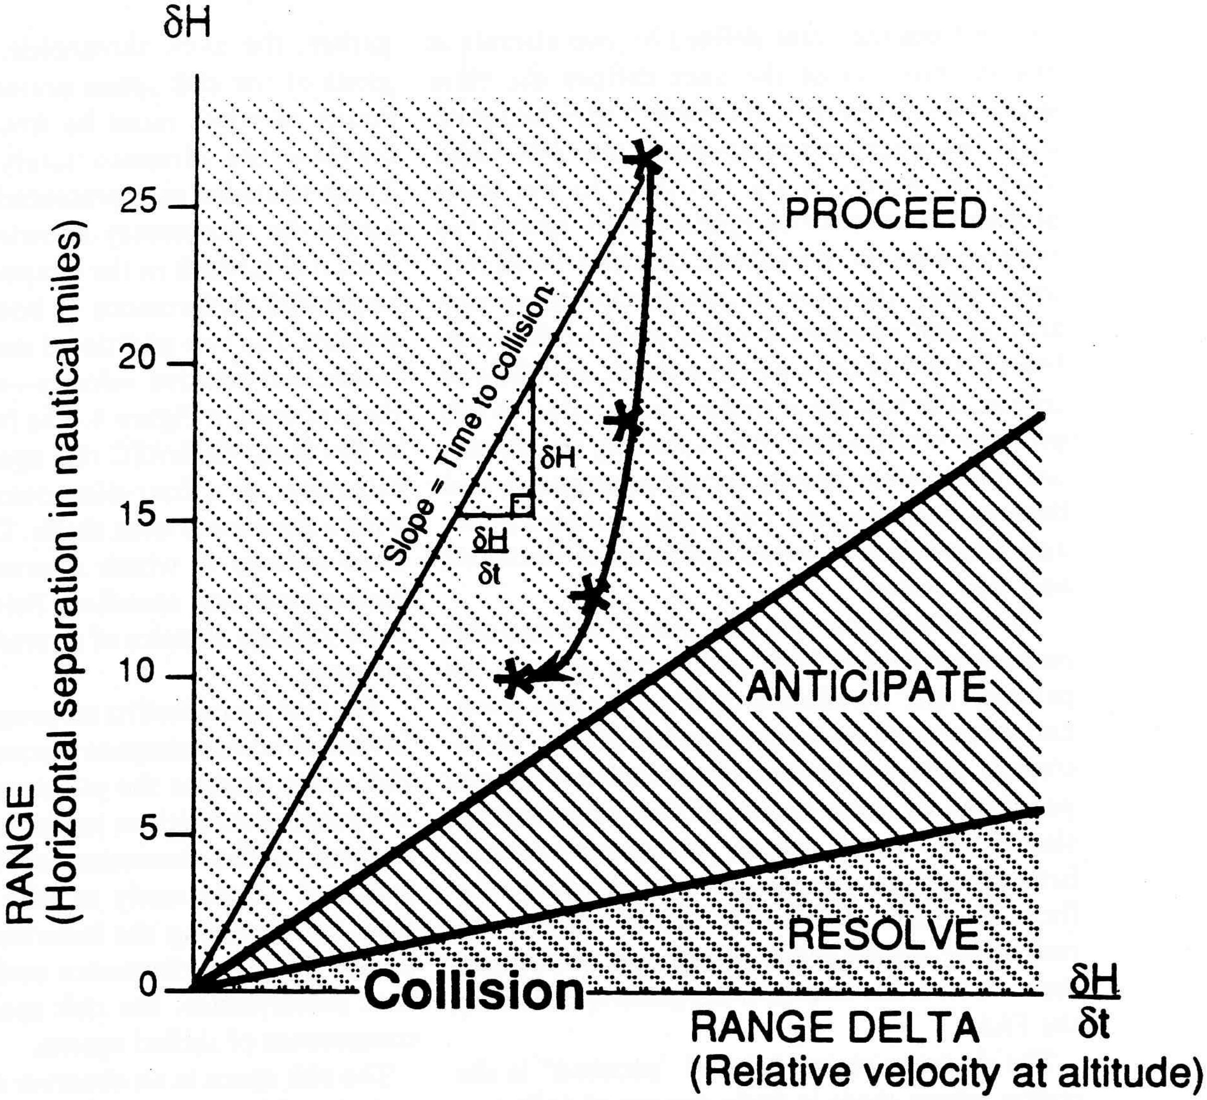

Figure 4. *The horizontal plane of the flight deck/ATC risk space for the en route environment. The y axis is the horizontal distance (separation, range) between two aircraft. The x axis is the relative velocity at which the aircraft are approaching or diverging. At any time, positions and airspeeds define a point in the risk space. Flight through the airspace traces a path through the risk space. The slopes of lines in the risk space have units of time (e.g., 30 and 300 s). The line connecting the point defined by two aircraft to the origin of the axes represents the time remaining before the aircraft collide. We propose that two such lines form thresholds for anticipating and resolving a conflict. These thresholds of safety, shown by heavy lines, parse the risk space into decision regions. Each region is associated with a set of safe and effective decision alternatives. The set of alternatives is defined by the invariant in the flight deck/ATC perception-action cycle.*

controllers must attend to in order to know what they must know to attain the goals specified by the FAA. Note that time is not an axis of the risk space. The risk space represents time both implicitly and explicitly. At every instant, the locations and airspeeds of two aircraft define a unique point in the risk space. As the aircraft fly through the airspace, the point charts a path through the risk space. Figure 4 shows one such path: two aircraft are converging at a rate that is decreasing; one (or both) may be making a course adjustment.

As shown in Figure 4, the slope of any straight line in the risk space has units of time. A line drawn from the point defined by two aircraft to the intersection of the axes defines the time when those aircraft will collide if they maintain their headings and airspeeds. Lines like those shown in Figure 4 are solutions to the basic problem of en route navigation: They specify the time to collision. We submit that pilots and controllers pay attention to information about separation and airspeed in order to solve the problem of whether (and when) aircraft are likely to collide. The risk space makes the solution to this problem transparent. By capturing the invariant relationship among location, airspeed, and time to collision, the risk space enables pilots and controllers to make appropriate decisions and take action.

Thresholds parse the risk space into decision regions associated with alternative actions. We propose that thresholds of time prompt pilots and controllers to invoke specific classes of decisions. Figure 4 depicts three decision regions populated by three classes of decisions (and actions): proceed on course, anticipate potential future conflicts, and resolve impending conflicts. The thresholds and decision regions in the risk space define what pilots and controllers must do in order to attain the goals specified by the FAA.

The decision region labeled "proceed" is the region where there is little danger of collision. The goal of all control decisions is to keep all aircraft in the decision region labeled "proceed." Conversely, the small region labeled "resolve" corresponds to situations where collision is imminent and pilots must take evasive action. The region in the middle, labeled "anticipate" is an oversimplification, but in general it is the region where ATC and pilots go on alert and come to decisions that are intended to prevent a conflict from materializing.

The risk space framework integrates information critical to safety considerations in a manner that specifies the action appropriate for a given sector at a particular time. The decision regions specify (and can be used to communicate) the actions that pilots and controllers need to take given the unfolding of events in the airspace. Together, the axes, thresholds, and decision regions of the risk space present a complete account of what must be known and done to traverse the airspace safely and efficiently. These relations, as represented by the risk space, afford the momentary knowledge of the distribution of aircraft in the airspace that guides the successful performance of both pilots and controllers. The two additional axes—vertical separation and relative velocity—are orthogonal to those shown in Figure 4. The full representation of the flight deck/ATC risk space requires visualization of a four-dimensional hyperspace where thresholds form shells. The thresholds enclose regions in which alternative actions are appropriate and specified. Points pass through these regions as pairs of aircraft fly through the airspace.

The risk space facilitates prognosis but is itself invariant across airspace (sectors) and time. The risk space enables the pilot/controller to adapt to changing conditions in the task environment, to focus on the dimensions of risk that are critical, and consequently to take appropriate action. By capturing the invariant relations that define skilled performance and by prescribing that performance, the risk space specifies the competence of skilled agents.

The risk space is an observer construct: It is a paramorphic representation (Dawes, 1979) of the adaptation of skilled pilots and controllers. We do not claim that any individual pilot or controller necessarily visualizes a four-dimensional risk space. Our claim is that the risk space framework (when fully developed) specifies competence at decision making and, therefore, can be used to explain and predict skilled performance. The risk space explains and predicts skilled performance in commercial aviation. Its axes, thresholds, and decision regions capture competence at knowing what must be known and doing what must be done when traversing the airspace.

The risk space stands as an operational definition of SA in the en route environment. Specifically, the risk space codifies the relevant information—separation and relative velocity—

that the environment makes available. Further, the risk space codifies as thresholds of risk the knowledge the agent requires to assess that information. Finally, the risk space embeds in decision regions the actions that an adapted agent's knowledge will direct it to take to meet its goals.

To our knowledge, this specificity is unparalleled in SA. Our current research is directed at defining (1) the locations of the thresholds under a wide range of airspace configurations and (2) the set of appropriate actions given those configurations. The risk space hypothesis of the SA of pilots and controllers in the en route environment is open to empirical challenge. It is comforting, however, to see that this same conceptualization has been independently generated by others for use in the realm of road vehicle collision avoidance (Fancher et al., 1994).

THE CURRENT ALTERNATIVE: A SYMBOL WITH MULTIPLE REFERENTS

Published accounts of SA appear rooted in the assumption that SA is behavior. Such a definition allows SA infinite variability. SA would become whatever an agent has and/or does that makes an observer believe the agent is skilled. At best, such an approach is descriptive. For instance, Endsley's (1988; 1995 [this issue, article on theory]) pioneering account equates SA directly with "an important component of . . . performance" (Endsley, 1988, p. 97). She stated that SA "can be conceived of as the pilot's internal model of the world around him at any point in time" (1988, p. 97). The problem here is the attempt to define a new term (SA) with reference to a second—mental model—which is itself ill defined and the subject of much contention and confusion. Like SA, mental models are intuitively appealing but difficult to corral.

We assume that the term *mental model* is used in the literature to refer to an organized structure of task-relevant information that allows mental simulation of processes, causes, or events in the environment (e.g., Gentner and Stevens, 1983; Johnson-Laird, 1983). With respect to Endsley's assertion, in order to define SA, we would have to develop a complete specification of an agent's mental model. Such a specification would obviously vary from agent to agent, even in the same situation. To equate SA with momentary knowledge and mental models is to run the risk of allowing SA to degenerate rapidly into "whatever is inside your (skilled) head."

Equating SA with mental models also appears to confound knowledge with process. For instance, Endsley (1988) defined three levels of SA: perception of the elements in the environment within a volume of space and time, comprehension of their meaning, and projection of this status in the near future. These levels sound very similar to three stages of information processing. We suggest that the three levels are, in fact, isomorphic with the processes posited in Newell and Simon's (1972) generic model of the human problem solver. In Newell and Simon's account, the agent samples the environment, generates a representation of that environment, and invokes methods that transform the representation and, on occasion, effect change in the environment. The outputs of the processes of sampling, representation, and transformation (projection) are knowledge and behavior. The processes themselves are neither knowledge nor behavior. Accordingly, it is inconsistent to define SA in terms of three levels of processing while equating it with momentary knowledge.

The evident confusion in the assertion that SA is performance and/or mental model spurs our desire to offer an alternative account. Lest we be viewed as mere critics, we are the first to acknowledge that our definition of SA draws heavily on previous work, especially Endsley's. We take, however, a more ecological approach to define SA as the invariant in the agent-environment system that generates the momentary knowledge and behavior required to attain the goals specified by an arbiter of performance in the environment.

DISCUSSION

We have argued that SA is a facet of consciousness. It is one in a long line of energetic constructs (e.g., attention, stress, and workload) that in a progressive fashion have reintroduced consciousness into scientific discussion of human behavior. As consciousness has re-emerged, age-old questions as to its nature—product or process, knowledge or performance—have percolated to the surface. We submit that a philosophical resolution of these questions is unlikely to be achieved in any facile manner. However, for practical purposes, those in the field of human factors can impose some operational bounds that allow the definition of terms.

We have articulated such bounds. We posit first that SA is externally directed toward a task environment. This means SA is a facet of consciousness, not necessarily all of consciousness. Second, for SA to exist, constraints on performance must be made explicit in the environment. SA is referenced to those goals and boundaries on performance. Finally, we recognize SA as an invariant component in an adaptive cycle of knowledge, action, and information. In this cycle, knowledge directs adaptive behavior, which modifies the environment, which in turn then informs knowledge; and so the cycle continues. Adaptive behavior that satisfies the arbiter of performance is the desired by-product of a cycle that is driven by competence.

Our discussion of agents and environments supports our proposal that SA is restricted to tasks that present clear and defined objectives that can be externally verified. Consequently, tasks with goals that are accessible only to introspection do not engender SA in the sense that we define it. SA is the competence that generates skilled performance in task environments that define the agent's goals. As a competence-level construct, SA guides our evaluation of an agent's performance. We therefore advocate developing representations of an agent's SA, like the risk space of Figure 4, for their potential to guide analyses of the interaction of skilled agents and dynamic task environments. However, we caution that much remains to be understood about how SA relates or coincides with other facets of consciousness that have enjoyed a similar vogue in the past (see Wickens, 1993).

We do not apologize if our approach proves disturbing to those who desire static definitions. Any useful notion of SA must take into account the central problems of generativity and nonstationarity posed by environments that are continually changing and by agents who are continually adapting and learning. The risk space is an observer construct that, we contend, captures the structure of a skilled agent's adaptation to a dynamic environment filled with moving targets. The risk space is defined by sources of information, by criteria for action, and by the actions themselves; together, these factors address the goals and standards for performance specified by a normative arbiter. The sources of information, criteria, and actions are invariant across situations; the risk space is the structure that makes them operative. Like a grammar for language, the risk space generates all possible situations and prescribes adaptive, appropriate behavior. The risk space is an operational definition of the SA of pilots and controllers in the midair (en route) environment. The invariant structure it captures predicts the cycle of adaptive behavior that all recognize as the product of SA.

At a very simple level, SA is an appropriate awareness of a situation. This implies that there are some things in a situation that have central importance and other things that are not relevant at all. Because the situations of interest to those studying human factors are ubiquitously dynamic, the things that are important and the things that are irrelevant can change with little warning. Others have seen fit to seize on this point to describe SA as a characteristic of the individual agent, who may, therefore, be said to have good or poor SA, depending on his or her propensity. What we have tried to affirm here is that SA is not resident in the agent but exists in the invariant interaction of the agent and his or her environment.

We suggest that SA may be good or poor depending on factors that constrain that interaction. It is correct to note that the behavior that SA generates is amenable to manipulation through change in the way information is displayed. With others in the ecological movement, we suggest that context-specific emergent properties of a task may be displayed so that the critical aspects of a situation are brought to awareness (Bennett and Flach, 1992; Hansen, 1995; Vicente and Rasmussen, 1992).

We have proposed here that for the context we have specified—en route navigation—the emergent property of the environment is described in a multidimensional hyperspace with axes representing separation and the rate of change of separation. To achieve the arbiter's goal of collision avoidance, pilots and controllers need be aware only of interobject location in this space. Through this specification we bound the nature of the navigation problem and overcome the infinite regress of context change and awareness change with every fluctuation of the dynamics.

The specifications of the risk space also address the question of how future interfaces may be constructed. It is clear that the evolution of the human-machine interface has been driven by the mediation of computer control. Equally clearly, fostering facile interactions has mandated progressively more use of human spatial abilities. Alphanumeric strings have become vestigial descriptions, whereas icons and three-dimensional representations have flourished. As a result, we expect that the emergent spaces that describe the constraints on system performance will be used to generate visual boundaries in virtual reality.

Interaction in virtual reality has yet to match prognostications (Kozak, Hancock, Arthur, and Chrysler, 1993), but the concept of graphic boundaries to operational status should, in time, enhance performance by adaptively displaying the invariants in dynamic operational phase spaces. In sum, we see value in the phenomenon called situation awareness. We expect to see it and other, even clearer, manifestations of consciousness as we design and evaluate purposive human-machine-environment systems.

ACKNOWLEDGMENTS

Parts of this work were presented at the First Annual Center for Applied Human Factors in Aviation International Meeting on Situation Awareness, Orlando, Florida, February 1993, and at the 37th Annual Meeting of the Human Factors and Ergonomics Society in Seattle, Washington, October 1993. Support for portions of the present work came from the Federal Aviation Administration, Grant 93-G-048 (John Zalenchak, technical monitor) and from the U.S. Navy, Contract N/N62269-92-0211 (Jeff Morrison, technical monitor) to the second author. The views expressed are those of the authors and do not necessarily represent those of the named agencies.

REFERENCES

Anderson, J. R. (1990). *The adaptive character of thought.* Hillsdale, NJ: Erlbaum.

Bennett, K. B., and Flach, J. M. (1992). Graphical displays: Implications for divided attention, focused attention, and problem solving. *Human Factors, 34,* 513–533.

Chomsky, N. (1965). *Aspects of a theory of syntax.* Cambridge: MIT Press.

Dawes, R. M. (1979). The robust beauty of improper linear models in decision making. *American Psychologist, 34,* 571–582.

Endsley, M. R. (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.

Fancher, P. S., Ervin, R. D., Bareket, Z., Johnson, G. E., Trefalt, M., Tiedecke, J., and Hagleitner, W. (1994, April). *Intelligent cruise control: Performance studies based upon an operating prototype.* Paper presented at the 1994 Annual Meeting of IVHS America, Atlanta, GA.

Flach, J., Hancock, P. A., Caird, J. K., and Vicente, K. (Eds). (1995). *Global perspectives on an ecological approach to human-machine systems.* Hillsdale, NJ: Erlbaum.

Freeman, G. L. (1948). *The energetics of human behavior.* Ithaca, NY: Cornell University Press.

Gentner, D., and Stevens, A. L. (1983). *Mental models.* Hillsdale, NJ: Erlbaum.

Gibson, J. J. (1979). *The ecological approach to visual perception.* Boston: Houghton-Mifflin.

Hancock, P. A., Flach, J., Caird, J. K., and Vicente, K. (Eds.). (1995). *Local applications in an ecological approach to human-machine systems.* Hillsdale, NJ: Erlbaum.

Hansen, J. P. (1995). Representation of system invariants by optical invariants in configural displays for process control. In P. A. Hancock, J. M. Flach, J. K. Caird, and K. H. Vicente (Eds)., *Local applications in an ecological approach to human-machine systems.* Hillsdale, NJ: Erlbaum.

Holland, J. H. (1992). *Adaptation in natural and artificial systems.* Ann Arbor: University of Michigan Press. (Original work published 1975)

Hopkin, V. D. (1988). Air traffic control. In E. L. Wiener and D. C. Nagel (Eds.), *Human factors in aviation* (pp. 639–663). San Diego, CA: Academic.

Johnson-Laird, P. N. (1983). *Mental models.* Cambridge, MA: Harvard University Press.

Koestler, A. (1978). *Janus.* New York: Random House.

Kozak, J. J., Hancock, P. A., Arthur, E. J., and Chrysler, S. T. (1993). Transfer of training from virtual reality. *Ergonomics, 36,* 777–784.

Marr, D. (1982). *Vision.* San Francisco: Freeman.

Moray, N. (1986). Monitoring behavior and supervisory control. In K. R. Boff, L. Kaufman, and J. P. Thomas (Eds.), *Handbook of perception and human performance* (Vol. 2, pp. 40.1–40.51). New York: Wiley.

Neisser, U. (1976). *Cognition and reality: Principles and implications of cognitive psychology.* San Francisco: W. H. Freeman.

Newell, A. (1981). The knowledge level. *Artificial Intelligence, 18,* 87–127.

Newell, A., and Simon, H. A. (1972). *Human problem solving.* Englewood Cliffs, NJ: Prentice-Hall.

Phatak, A. V., and Bekey, G. A. (1969). Decision processes in the adaptive behavior of human controllers. *IEEE Transactions of Systems Science and Cybernetics, 5,* 339–352.

Simon, H. A. (1982). *Sciences of the artificial* (2nd ed.). Cambridge: MIT Press.

Smith, K., and Hancock, P. A. (1992). Managing risk under time stress. In *Proceedings of the Human Factors Society 36th Annual Meeting* (pp. 1019–1023). Santa Monica, CA: Human Factors and Ergonomics Society.

Tenney, Y. J., Adams, M. J., Pew, R. W., Huggins, A. W. F., and Rogers, W. H. (1992). *A principled approach to the measurement of situation awareness in commercial aviation* (NASA Contractor Report 4451). Langley, VA: NASA.

Vicente, K. J., and Rasmussen, J. (1992). Ecological interface design: Theoretical foundation. *IEEE Transactions on Systems, Man and Cybernetics, 22,* 589–606.

Watson, J. B. (1913). Psychology as the behaviorist views it. *Psychological Review, 20,* 158–167.

Wickens, C. D. (1993). Workload and situation awareness: An analogy of history and implications. *Insight: The Visual Performance Technical Group Newsletter, 14,* 4.

Date received: May 6, 1993
Date accepted: October 4, 1994