

UNCLASSIFIED

Defense Technical Information Center
Compilation Part Notice

ADP021346

TITLE: Knowledge, Understanding, and Behavior

DISTRIBUTION: Approved for public release, distribution unlimited

This paper is part of the following report:

TITLE: International Conference on Integration of Knowledge Intensive Multi-Agent Systems. KIMAS '03: Modeling, Exploration, and Engineering Held in Cambridge, MA on 30 September-October 4, 2003

To order the complete compilation report, use: ADA441198

The component part is provided here to allow users access to individually authored sections of proceedings, annals, symposia, etc. However, the component should be considered within the context of the overall compilation report and not as a stand-alone technical report.

The following component part numbers comprise the compilation report:
ADP021346 thru ADP021468

UNCLASSIFIED

Knowledge, Understanding, and Behavior

James Albus

Intelligent Systems Division , National Institute of Standards and Technology,
Gaithersburg, MD 20899
301-975-3418
james.albus@nist.gov

Abstract – *What does it mean to know something? What is understanding? How does understanding influence behavior? These are philosophical questions that have occupied the best human minds since time of the ancient Greeks. Only in the last 50 years have scientists been able to address these issues with empirical methods. Only in the last decade has the computational power become widely available to pursue these questions beyond the domain of toy problems. Recent results in the field of autonomous ground vehicles are yielding insights into how to acquire and structure knowledge so as to produce intelligent behavior in the domain of autonomous ground vehicles.*

Introduction

The dictionary defines *to know* as:

- 1) to apprehend with the conscious mind (e.g., to know what is going on), or
- 2) to be acquainted with by experience (e.g., to know that the stove is hot), or
- 3) to have acquired a skill (e.g., to know how to water ski), or
- 4) to be informed about (e.g., to know what time it is), or
- 5) to have committed to memory (e.g., to know the Gettysburg Address), or
- 6) to be convinced of (e.g., to know that Jesus loves me), or
- 7) to be aware of (e.g., to know that you are lost)

Knowledge is all of these things, and much more. Knowledge includes all kinds of information about the world and ourselves. Knowledge is information about things, places, events, states, situations, classes, relationships, causes, effects, tasks, skills, goals, motives, values, plans, behaviors, perceptions, opinions, prejudices, feelings, experiences, tastes, rules-of-thumb, laws of behavior, and rules of etiquette. Knowledge includes memories of past experience, and generalization from application of logical rules to the combination of many past experiences with

apriori knowledge. Knowledge includes models of how the world works, and knowledge of procedures for using models to generate expectations about what is likely to result from alternative choices of future behavior. Knowledge includes knowing how to do things such as how to tie a shoe, how to brush our teeth, how to read and write, how to recognize our friends. Knowledge includes knowing how to cope with everyday unexpected events and compensate for failures. Skills are knowledge of how to do tasks.

Most of us know how to ride a bike, drive a car, and use a telephone. Some of us know how to scuba-dive, to do calculus, to breed orchids, or read ancient Egyptian hieroglyphics. All of us know how to recognize objects and places, family and friends, events and relationships. We know what familiar things are called, and how to classify things that we have never seen before. We know how to speak and write and understand what is spoken and written. We know what we think and feel. We know what others think and feel about us. We know about time and space, places and things, events and situations. We know where we are and where things are around us. We know where things are in our homes. We know how to find our way home from distant places. We know what we have seen and heard. We know how things act and respond to our actions. We know what is probable and what is impossible. We even know what we don't know. We know if we have never before been somewhere, or done something, or seen someone.

We know what is good and bad. We know what is right and wrong. We know how to behave under a wide variety of circumstances. We know what we like and don't like. We know what hurts and what feels good. We know how things smell, taste, and feel. We know when something smells sweet or rotten. We know what things are worth. We know how much cost and risk we are willing to incur in order to achieve our goals, or to avoid failure. We know what it means to win or lose. We know the difference between reward and punishment, praise and criticism. We know when we are hungry, thirsty, or suffocating. We know when we feel happy or sad. We know when we feel lonely, and when we are sexually aroused. We know when we need a hug.

Many kinds of knowledge can be expressed by rules. There are rules of the road. There are rules of physics and chemistry. There are rules of cause-and-effect. There are rules that can be expressed in equations as laws of physics or theorems in mathematics or logic. There are rules for medical diagnosis and spectrum analysis. There are rules of causality and pragmatics. There are rules of syntax and semantics.

Knowledge Representation

Knowledge can be represented as iconic images or maps,¹ or as symbolic variables or abstract data structures such production rules, equations, objects, and classes. Iconic representations are well suited to support geometry and navigation, but they have limitations on range and resolution. Symbolic representations are well suited to support language, logic, and mathematics, but they have limits on vocabulary and ontology. Relational pointers support syntax, grammar, and semantics. They have direction and type. Symbolic knowledge can be encoded into language and communicated by vocalization, writing, or sign language.

Iconic representations are typically grounded in the world, (e.g., visual images on the retina, tactile images on the skin, patterns of sound in the ears, inertial signals from vestibular sensors, or kinematic sensors from tendons and joints.) Symbolic representations are abstracted from sensory data, and grounded through links back to iconic representations. For example, a physical object in the world can be represented in an intelligent system both as an iconic group of pixels in an image, and as a symbolic data structure with a name, attributes, class, state, and worth. An event can be represented both as a temporal pattern of signals (e.g., amplitude, frequency, phase, duration), or as a symbolic data structure with name, attributes, class, state, and worth. Relationships can be represented as labeled pointers that link iconic and symbolic representations to each other (e.g., *is_a*, *belongs_to*.)

Representation in the brain of knowledge about the world lies at the core of the phenomenon of mind. Arrays of signals from the retina produce immediate visual experience. Arrays of signals from tactile sensors produce the sensation of tactile feeling. Signals from the ears produce awareness of the acoustic environment. Signals from chemical detectors in the nose and tongue produce the sensations of smell and taste. Signals from thermal sensors tell us what is hot or cold. Signals from structural damage sensors produce the sensation of physical pain. All of these signals are interpreted by the mind in the context of what it knows or believes about how the world works and how things behave. Knowledge enables the mind to analyze the past, to

evaluate the present, and plan for the future. Knowledge of what is going on in the world enables computational processes of the mind to generate behavior that is most likely to succeed in accomplishing system goals.

Understanding

Understanding can be said to occur when the system's internal representation of external reality provides a model of the world that is adequate for generating successful behavior. Understanding enables a system to predict what is likely to happen in the future, both in the near term (for filtering and interpreting immediate sensory experience), and in the long term (for planning and evaluating alternative future behaviors.) Understanding enables the system can translate what it knows into actions that achieve its goals.

We say a person understands if s/he can answer questions with information that is not explicitly stored in memory, i.e., if s/he can show evidence of correct reasoning about the subject. We say a person understands if s/he can make reasonable predictions about what will happen under a variety of circumstances. We say a person understands if s/he can take actions that preempt problems or failures, or can optimize results in a variety of situations.

Understanding requires that knowledge have internal semantic structure that is able to support causal reasoning. Understanding also requires that internal data structures have semantic links to external objects and events in the world. In semiotic terms, understanding implies "symbol grounding."

Symbol grounding occurs through sensing the world and comparing sensory observations with predictions generated by the internal world model. Differences between what is observed and what is predicted by the internal representation are used to correct and refine the model. This is known as recursive estimation, or Kalman filtering. It is what establishes the links between internal data structures and external reality.

Figure 1 illustrates an iconic representation of the world as seen by a camera through the front windshield of an automobile on a four-lane road. Observed pixel brightness, color, and range images are processed to detect edges, lines, and surface features. Detected features are grouped into objects that can be given labels and assigned to classes. Relationships between objects represent patterns and situations. At each level, sensory processing algorithms focus attention on regions of importance in the image, group pixels and features into entities, compute attributes and state of those entities, perform recursive estimation on attributes and state

variables, compare attributes of observed entities with attributes of stored class prototypes, classify observed entities on the basis of that comparison, and establish

links between regions in the image and symbolic data structures.




Figure 1. A scene looking through the windshield of an automobile on a four-lane road.

Estimates of range at each pixel enable the viewpoint of Figure 1 to be transformed into the overhead viewpoint of Figure 2, which is more suitable for planning and executing driving behavior. In Figure 2, abstract data structures with attributes and state of objects and groups are shown. The knowledge represented in Figure 2 enables understanding, in the sense that an autonomous driving algorithm can use it to make plans and react to situations that achieve safe and efficient driving behavior on a typical four-lane road. Similar data structures and

mechanisms of perception can be extended to enable autonomous driving behavior in other situations, including intersections, freeways, city streets, dirt roads, and cross-country driving through fields and woods. For off-road driving terrain features such as slope, roughness, ground cover, stumps, fallen trees, bushes, and soil mechanics need to be represented. For city streets, knowledge about one-way streets, pedestrians, school zones, and emergency vehicles may be required.




Figure 2. A map representation of the scene from Figure 1. For each object of interest, there is a pointer to a symbolic data structure that describes the attributes and class of that object. This level of understanding enables behavior that can achieve system goals of safely driving to a desired destination.

Behavior

Behavior is the sequence of actions that a control system produces. Action results from actuators that exert influence on the world. Output from computational processes in the brain causes actuators to generate forces that move arms, legs, hands, and eyes. Actuators point sensors, excite transducers, manipulate objects, use tools, and steer and propel

locomotion. A biological intelligent system may have tens, hundreds, thousands, even millions of actuators, all of which must be coordinated in order to perform tasks and accomplish goals. Natural actuators are muscles and glands. Machine actuators are motors, pistons, valves, solenoids, displays, and transducers for acoustic and electromagnetic radiation.

Intelligent behavior is designed to achieve or maintain a goal. A goal is a desired result, or desired state of the world. Goal seeking is the *raison d'être* of intelligent behavior. The ultimate goal is the intelligent system's reason for being. A goal may come from outside the intelligent system in the form of a command, or be generated internal to the system in response to a need, drive, or urge. An example of an external goal might be an order from a superior in a military chain of command to be at a specified place at a specific time. An example of an internal goal might to find food or shelter in response to hunger or feeling cold. A measure of intelligence is the degree to which the system is successful in achieving or maintaining its goals. Natural selection favors behaviors that are most likely to result in the survival and reproduction of the individual's genes. The brain is thus a mechanism for realizing the gene's goal of maximizing the likelihood of its own reproduction.

The brain is a well-organized society of computational processes designed by natural selection to generate and control behavior. The brain is first and foremost a control system. All brains that ever existed, even those of the tiniest insects, generate and control behavior. Of course, the ultimate goal of survival and propagation cannot be achieved by any single action, or any fixed series of actions. To survive and propagate in the natural world a creature must decompose its ultimate goal into subgoals such as find food, acquire territory, build shelter, escape predators, attract a mate, rear a family, and adapt to the changing seasons. For example, survival may require hibernation or migration toward the equator during the winter. Each of these subgoals must be further decomposed into lower-level shorter-term sub-subgoals, involving more localized and specific behavior. At each level of decomposition, a goal from a higher level causes behavior consisting of actions designed to achieve strings of lower level goals that work together to accomplish the higher level goal. At each level, feedback from the environment is used to compensate for errors and perturbations so that goals are accomplished despite mistakes and unexpected events. Eventually at the lowest level, strings of neural impulses are delivered to muscles and glands that act on the world.

The structure of the brain has evolved to facilitate this principle of task and goal decomposition. In humans, plans for the long-term future are developed in the forebrain, the most recently evolved brain structure. The frontal cortex uses reason and logic to manipulate abstract models of the world including maps, symbols, rules of mathematics, and laws of physics for synthesizing high level concepts and plans. The frontal limbic system also uses high level value judgment to

make decisions based on social customs, religious beliefs, and legal obligations. In the frontal lobes, plans may be developed for organizing a hunt, building a house, writing a book, planning a military campaign, pursuing a career, or performing scientific research.²

Plans developed by the forebrain are expressed in terms of sequences of goals and actions for which there exist behavioral skills residing in the evolutionarily older pre-motor cortex and associated limbic regions. Behavioral skills may enable activities such as capturing prey, escaping a predator, winning a fight, consummating a sexual encounter, going to the store, getting a haircut, buying a car, or going out for dinner.

Plans developed by the pre-motor regions may be expressed as sequences of simple tasks that can be performed by the still older primary motor cortex and the basal ganglia. Simple tasks may consist of activities such as climbing over a log, following a path, putting on a coat, or opening a door.

Plans developed by the primary motor cortex and basal ganglia are expressed in terms of elementary movements for various parts of the motor system such as the arms, legs, and torso. Elementary movements may consist of activities as such as reaching, grasping, lifting, throwing, stepping, and jumping. Elementary movements are further decomposed and coordinated with balance, posture, and orienting information in the midbrain and cerebellum.³

At the lowest level in the spinal cord, there are computational modules containing the final motor neurons that control stretch in the muscles, tension in the tendons, and position of the joints. These produce the output drive signals to muscles that cause limbs and digits to move in coordinated ways so as to accomplish the behavioral goals generated at higher levels.

There is a sensory processing hierarchy that runs parallel to this behavioral hierarchy. Modules in the sensory processing hierarchy operate on signals from sensors to extract features, detect patterns, recognize events, classify entities, analyze situations, and recognize concepts. At each level, feedback from sensory processing is integrated with commands from higher levels to modify behavior so as to accomplish high level goals despite unforeseen events in the world. The sensory-motor system contains many parallel and cross-coupled hierarchies of computational modules. In the spinal cord, sensory signals interact with motor commands at several levels – at the final motor neurons, in the spinal motor coordination centers, in the cerebellum, pons, red nucleus, reticular formation, and substantia nigra. The primary sensory cortex lies

immediately adjacent to the primary motor cortex, so that the sensory areas for the hands, fingers, arms, legs, torso, and head lie adjacent to the motor areas for each of these body parts.

A Cognitive Architecture

4D/RCS is a cognitive architecture that provides a theoretical foundation for designing, engineering, integrating, and testing intelligent systems software for unmanned vehicle systems. 4D/RCS is the latest in a long line of RCS (Real-time Control System) designs that began in the 1970's with robot manipulators, evolved during the 1980's into automated manufacturing facilities, space station telerobotic systems, autonomous undersea vehicles, and during the 1990's into unmanned ground vehicles.³ 4D/RCS was developed specifically for the Army Research Laboratory Demo III program.^{4,5}

4D/RCS consists of a multi-layered multi-resolutional hierarchy of computational nodes each containing elements of sensory processing (SP), world modeling (WM), value judgment (VJ), and behavior generation (BG). At the lower levels, these elements generate goal-seeking reactive behavior. At higher levels, they enable goal-defining deliberative behavior. Throughout the hierarchy, interaction between SP, WM, VJ, and BG give rise to perception, cognition, imagination, and reasoning. At low levels, range in space and time is short and resolution is high. At high levels, distance and time are long and resolution is low. This enables high-precision fast-action response over short intervals of time and space at low levels, while long-range plans and abstract concepts are being formulated over broad regions of time and space at high levels.

4D/RCS closes feedback loops that enable reactive behavior at every level. SP processes focus attention (i.e., window regions of space or time), group (i.e., segment regions into entities), compute entity attributes, estimate entity state, and assign entities to classes at every level. WM processes maintain a rich and dynamic database of knowledge about the world in the form of images, maps, entities, events, and relationships at every level. Other WM processes use that knowledge to generate estimates and predictions that support perception, reasoning, and planning at every level.

4D/RCS has been implemented in part on the Demo III experimental unmanned vehicle (XUV) program and in related research.^{6,7,8,9} Four levels (Servo, Primitive, Autonomous Mobility Subsystem, and Vehicle level) of the 4D/RCS behavior generation, world modeling,

and value judgment processes have been more or less fully implemented on the Demo III XUVs and the NIST HMMWV, and the first two levels of sensory processing have been partially implemented. Even this partial implementation enabled performance of the Demo III vehicles that was extremely impressive. The vehicles were regularly able to traverse more than a kilometer of challenging terrain including dirt roads and trails, woods and fields, hills and valleys, filled with tall grass, weeds, stumps, fallen trees, and brush – without human intervention. Occasionally, the vehicles had to be stopped to prevent a dangerous situation, or because they became trapped by difficult terrain features (such as loss of traction on steep sandy slopes.) But more often, they were able to keep out of trouble, and accomplish their mission of reconnaissance, surveillance, and target acquisition. Particularly at night, the XUVs were often able to traverse difficult terrain as fast and reliably as skilled human drivers in HMMWVs. In a series of tests during the winter of 2002-03, the XUVs logged more than 400 km over a variety of terrains, including wooded hills, desert, and urban terrain. More than 96% of this distance was accomplished without any human assistance.¹⁰ As a result of these tests, it now is widely anticipated that autonomous vehicles will be capable of useful military missions in laying smoke, delivering supplies, serving as pack mules, and acting as forward observers for scouts by the year 2010.

References

- [1] Farah, M. J. "The Neural Bases of Mental Imagery", pp. 963-975, In *The Cognitive Neurosciences*, (Gazzaniga, M., Ed.) Bradford Books, Cambridge, MA, 1995.
- [2] Robin, N., and Holyoak, K.J. Relational complexity and the functions of prefrontal cortex, pp. 987-997, In *The Cognitive Neurosciences*, (Gazzaniga, M., Ed.) Bradford Books, Cambridge, MA, 1995.
- [3] Albus, J. S., and Meystel, A. M., *Engineering of Mind: An introduction to the science of intelligent systems*, Wiley, NY, 2001.
- [4] J. A. Bornstein, "U.S. Army ground robotics research program", *Proceedings of SPIE Aerosense Conference Vol. 4715*, Orlando, Florida, April, 2002
- [5] C. Shoemaker, J. Bornstein, S. Myers, and B. Brendle, "Demo III: Department of Defense testbed for unmanned ground mobility", *Proceedings of SPIE Vol. 3693, AeroSense Session on Unmanned Ground Vehicle Technology*, Orlando, Florida, April 7-8, 1999.

- [6] T. H. Hong, S. Balakirsky, E. Messina, and M. Shneier, "A Hierarchical World Model for an Autonomous Scout Vehicle", *Proceedings of SPIE Aerosense Conference Vol. 4715*, Orlando, Florida, 1-5 April, 2002.
- [7] T. H. Hong, T. Chang, C. Rasmussen, and M. Shneier, "Road Detection and Tracking for Autonomous Mobile Robots", *Proceedings of SPIE Aerosense Conference Vol. 4715*, Orlando, Florida, 1-5 April, 2002.
- [8] A. Lacaze, "Hierarchical Real-Time Path Planning for Obstacle Avoidance and Path Optimization for One or More Autonomous Vehicles", *Proceedings of SPIE Aerosense Conference Vol. 4715*, Orlando, Florida, 1-5 April, 2002.
- [9] S. Balakirsky and A. Lacaze, "Value-Driven Behavior Generation for an Autonomous Mobile Ground Robot", *Proceedings of SPIE Aerosense Conference Vol. 4715*, Orlando, Florida, 1-5 April, 2002.
- [10] Camden, R., Bodt, B., Schipani, S., Bornstein, J., Phelps, R., Runyon, T., French, F., Shoemaker, C., "Autonomous Mobility Technology Assessment Interim Report – February, 2003," *Army Research Laboratory ARL-MR-565*, June, 2003