

Insect Navigation: Brains as Symbol-Processing Organs

C. R. Gallistel

Department of Psychology

University of California, Los Angeles

1. Introduction

Digger wasps dig nest burrows in which they lay an egg. They cover the burrow entrance with pebbles, making it all but invisible, then fly off in search of insect prey. They sting their prey to paralyze it, carry it back to the burrow, open the burrow, drag the prey inside, exit the burrow and close it up again. The larva that hatches from the egg eats the paralyzed, still living prey. This aspect of the natural history of digger wasps cautions us against an overly sentimental view of the workings of mother nature, at least at the insect level. Quite aside from pointing up this moral, it raises two questions of profound interest to contemporary cognitive scientists and neurobiologists: (1) How does the wasp find her way back to her invisible nest? (2) What does the answer to that question tell us about the nature of the brain?

1.1 Behavior suggestive of a cognitive map

Tinbergen and Kruyt (1938) began the process of answering these questions when they arranged a circle of sticks and pine cones around the nest while a wasp was digging it, then displaced the circle while she was gone. The returning wasp landed in the center of the displaced circle. She could not find her nest, even though it was only a meter away. Their experiment demonstrated that the wasp located the entrance to it by reference to nearby landmarks, not by the sight or smell, etc. of the nest itself. This is but one example of the quite general principle that animals locate their goals by reference to the surrounding terrain, and not, contrary to what people commonly assume, by homing on a sensory beacon emanating from the goal itself (see Gallistel, 1990, Chap. 5, for examples in animals ranging from bees to humans).

The discovery that the wasp locates her nest by reference to the landmarks or terrain features that surround it implies that she

notes the position of her nest relative to these landmarks before her departure, and, in fact, it has since been shown that the departing wasp makes a special survey flight for that purpose (Zeil, 1993a). Moreover, it implies that when she returns she uses the remembered position of her nest relative to those landmarks to guide her flight path to a landing site. A record of the position of one or more points of interest relative to the positions of landmarks or salient terrain features is a **map*. A map that resides in the brain of an animal rather than on a sheet of paper or in the memory banks of a computer is called a **cognitive map*. Thus, the results of Tinbergen and Kruyt's experiment suggest that the wasp makes a cognitive map of the immediate vicinity of her nest and uses that map to choose a landing site.

The Tinbergen and Kruyt (1938) experiments did not, however, shed light on how the wasp navigated over the large distances covered by her foraging. They did not show how she found her way back from the place where she captured her prey to the vicinity of her nest, where the small-scale terrain features surrounding it would first become clearly visible. In 1950, Thorpe did some experiments that suggested that digger wasps relied on a cognitive map of their large-scale environment to find their way to the vicinity of their nest from wherever they happened to capture prey. He worked with a species of digger wasp that preys on caterpillars too large to be carried home in flight. The female of this species drags her prey home across the gravely ground. In the experiment shown in , Thorpe found a wasp dragging her prey homeward. He first placed obstacles in her path that forced her to deviate from her course (numbered marks on figure). As soon as she had cleared each obstacle, she resumed her course. This proved that the wasp was not marching blindly along like a wind-up toy tank, without regard to where her marching was actually getting her. When forced to deviate from her course, she corrected for the deviation she had made --a result that has been confirmed recently by similar experiments with ants (Schmidt, Collett, Dillier & Wehner, 1992).

Next, Thorpe picked up the wasp, placed her in a box, and carried her about 30 meters to the east northeast of the capture point (broken line on Figure 1), where he released her. At the release point, a clump of grass obscured the view of her nest site. It is, in any event, unlikely that the view of the nest site plays any role in controlling most of the wasp's homeward journey because she

views the world from a height of a few millimeters. This means that her view of level ground more than a fraction of meter distant is generally obscured by small bumps in the terrain. Although at the release site she could not see or otherwise detect the area immediately around her nest, the newly released wasp nonetheless headed straight through the clump of grass on a course that lay within 15° of the true course from the release site to her nest. In other words, she set her course for a site whose position she could not observe. She did so in all likelihood by reference to those parts of the terrain that she could observe—such as the small birch trees (tree symbols on Figure 1) and the horizon contour formed by the rim of the depression (indicated by the hatching on the periphery of the map of the experimental terrain in Figure 1).




Figure 1. *Digger-wasp capture and release experiment. Heavy line indicates the course of the insect. It was first observed at A. Numbers 1-4 indicated where obstacles were placed in its path. It went round the obstacles and resumed its course. At B it was captured and carried to C (broken line). Numbers 5-7 indicate places where obstacles were again placed in its path. N indicates its nest. The hatching around the perimeter of the map indicates a slight depression in an area of gravelly heath land. Tree symbols indicate small birch trees. Patches of scattered plants of Erica and Calluna indicated by the 'E' symbol. Tufts symbolize tussocks of juncus grass. (From Thorpe, 1950, by permission of E.J. Brill, publisher.)*

1.2 Maps and navigational computation

The ability to use currently perceived terrain features to set a course for an area you cannot presently perceive generally requires the use of a map (also called a chart), which specifies the position of your destination relative to the positions of the landmarks you currently perceive. It also implies the ability to do navigational computations utilizing the data on the map, because navigation (course-setting) requires computation. Thus, these results suggest the brain of an insect has computational capabilities analogous to those found in, for example, a GPS (Geographic Position Sensor).

A GPS is a miniature marvel that has recently become available to the general public in a package that weighs a few ounces and fits in the palm of the hand. It consists of a computer specialized for navigational computations, a precision clock, and a radio receiver tuned to the signals broadcast by an array of earth-orbiting navigation-satellites. The satellites provide signals from which it is possible to compute the position of a receiver--the latitude, the longitude, and the elevation of the GPS--to within about 100 feet.

The brain of the GPS is its computer. You key into the long-term memory banks of the computer the latitudes and longitudes (and, if you like, also the elevations) of ten to several hundred points of interest to you--harbor entrances, buoys, headlands, the ends of runways, trailheads, turn-off points, campgrounds, etc--identified by a keyed-in name or number. These points of interest are called waypoints. Their latitudes, longitudes and elevations constitute a map. Thereafter, when you key in the name or number of a waypoint, the GPS uses the waypoint's latitude and longitude (retrieved from its memory banks) and the latitude and longitude of the GPS's current position (computed second by second from the radio signals it is receiving from the satellites) to compute the bearing and range of the waypoint. The *bearing* is the compass direction of the waypoint, for example, of the entrance of the harbor you are making for. The bearing tells you the course to steer (assuming that the waypoint in question is not more than a few hundred miles distant). The *range* is your distance from the waypoint. If you ask it, the GPS will also compute for you the range and bearing of any waypoint from any other, as well as things like your speed, and your estimated time of arrival.

1.3 Representational (symbol-processing) and non-representational (sub-symbolic) theories of mind

The computer in a GPS computes courses and distances to be traveled from its map, that is, from the positions of points of interest recorded in its memory banks. The fact that the digger wasp can set off on an appropriate course for a destination she cannot currently perceive after she has been displaced to a known point in her environment suggests that she, too, can compute courses from her cognitive map. What makes this of interest to cognitive scientists is that it suggests that the brain of an insect is a symbol processing organ. It suggests that the brain uses computation to construct a cognitive map and to compute courses from the positional values stored in that map. The positional data stored on a map together with the computations that utilize those data to set courses constitute a *representation*, a model of the world used to determine courses of action in that world. Thus, the homing abilities of insect might be taken as strong evidence in favor of a computational-representational theory of mind.

Computational-representational theories of mind are at the heart of a controversy that has raged through psychology, computer science, linguistics and neurobiology throughout much of this century. The community of cognitive, behavioral and neural scientists has long been divided over the issue whether the brain really represents the world within which the animal acts or merely appears to do so. Setting out toward destinations you cannot see suggests that you know where those destinations are relative to where you are, just as speaking a language fluently suggests that the speaker knows its grammar, but many participants in this debate believe that these conclusions are naive, that models positing symbolic knowledge inside the head are doomed to fail. They believe that behavior that suggests knowledge of the world is produced, not by a brain that computes symbolic representations of the world, but by a vast and plastic neural network that does not contain within it any symbolic knowledge of the world.

The debate between those who believe that much behavior is mediated by symbolic knowledge of the world in the brain and those who believe that it is not has been going on for a long time, but recently, this question has been formulated in the following terms: Is the brain a symbol-processing system? The rise of cognitive psychology in the sixties reflected the emergence of the (then novel)

view that the brain was a symbol-processing system analogous to a computer, and that it used its symbol-processing capacity to construct representations of the environment in which it functioned. This cognitive conceptual framework contrasts with the earlier framework of the behaviorist school, in which the brains of animals were not symbol-processing systems and did not construct representations of their environment.

The best known proponent of the behaviorist school of psychology were Clark Hull and B. F. Skinner. Although they were all contemporaries, psychologists like Hull and Skinner did not in those days pay much attention to the work of ethologists like Tinbergen and Thorpe, so I do not believe that either Skinner or Hull ever offered a non-representational account of performances like that shown in Figure 1. Hull (1930), in an immensely influential paper entitled "Knowledge and Purpose as Habit Mechanisms," offered an account of what he supposed to be the process by which a rat traversed a simple maze to reach food in a location where it had previously found food. This account was intended to illustrate Hull's contention that the rat's behavior did not depend on a representation of the position of the food within the spatial framework defined by the maze and the surrounding room and did not involve anything like the navigational computations that go on in a GPS. The essence of Hull's account was that the rat developed a chain of simple "habits" or stimulus-response units, which were in essence conditioned reflexes. Each individual habit formed because that particular response made to whatever stimuli impinged on the animal at that point in the maze, had repeatedly led to reinforcement, whereas other responses made to those stimuli had not. The reinforcement selectively strengthened the conducting pathway in the nervous system leading from the receptors excited by a particular stimulus configuration to the muscles whose contraction and relaxation generated the movement of the animal (the response). The animal's movements at different points in the maze were not controlled and integrated by a representation of a course to be followed. In fact, each response was a little independent atom of behavior. The only thing that integrated responses in different parts of the maze were their stimulus consequences. The stimuli generated by a preceding response were often an important part of the stimulus configuration for the next reflex. Also, responses the animal originally made only at sight or smell or taste of the food (e.g., salivation) could, Hull argued, come in time to be made early on in the maze (anticipatory goal

responses) and their stimulus consequences (the feel of salivation in the mouth) could become part of the stimulus configuration that elicited a movement in the maze, such as turning right at a choice point. The whole point of Hull's argument was that the rat did not have knowledge of the spatial layout of the maze and the location of the food (a cognitive map), it merely appeared to have such knowledge. And this was precisely the conclusion that Tolman (1948) questioned in his famous paper entitled, "Cognitive Maps in Rats and Men." (Tolman also paid no attention to the ethologists, even though their work provided more powerful evidence for his thesis than did his own experiments.)

The computational-representational or symbol-processing conception of the brain processes that generate behavior also contrasts with the conceptual framework urged by those proponents of neural net modeling who argue that the brain is not a symbol-processing organ, and that models of psychological function based on the assumption that it is have and must fail. Both the behaviorists and the champions of neural net models have argued that behaviors such as navigating through an environment or learning to speak a language make it appear that we have formed a representation of the spatial form of the environment (a cognitive map) or of the grammar of the language (a mental grammar), when in fact we have not. In this chapter, I will review experimental data on the navigational abilities of insects and argue that these data strongly support the view that brains are in fact symbol-processing organs.

1.4 Symbol-processing systems

A symbol-processing system, such as a conventional computer, has two fundamental components: First, it has a set of symbols, for example, symbols for distances, symbols for positions, symbols for temporal intervals, symbols for numerosity, symbols for verbs, and so on. The symbols take on different values at different times, depending on the values of the variables that they denote. For example, the symbol for the distance between oneself and home takes on different values depending on how far one is from home. In thinking about symbols, it is essential to distinguish between the symbol itself--a mark on paper, a bit pattern in a digital computer, or a voltage in an analogue computer, representing a distance or an interval--and the variable that it denotes--the actual distance or the actual temporal interval. The essence of a symbol is that it denotes or refers to something other than itself. Second, a symbol-

processing system has a set of operations that may be performed on the values of the symbols, for example, the operation of adding value of one symbol to the value of another to generate a third value. If the symbols being added represent temporal intervals, then the value that is their sum represents the combined interval or sum of the intervals. Thus, symbols may be combined to generate either symbol strings or new symbols, whose values are determined by the values of the symbols that were combined and by the combinatorial operation (e.g., addition or subtraction or multiplication or exponentiation, etc.).

The use of symbols in conventional computation is sometimes said to be **compositional*, because the computation may be composed out of any of the values that the symbols happen to have regardless of how they came to have those values. The addition process (or operation) simply combines two values to generate a third value. In Figure 2, Calvin shows a proper appreciation of the deep mystery of this operation, which is too often taken for granted. Often, the values used in a particular addition operation are generated by earlier operations, and this leads to the fundamental idea of the composition of operations or functions, which the student will have encountered both in calculus and in computer programming. The notion of composition of functions is expressed algebraically as follows: If $f(x) = x + a$ and $g(x) = x^2$, then $f(f(x)) = (x + a) + a$ and $g(f(x)) = (x + a)^2$ and $f(g(x)) = x^2 + a$, and so on. In other words, because the operations (in this case the functions f and g) are not affected by the origins of the values they operate on, you can take the values generated by one of the operations and make it the subject of the other operation or of a repetition of the first operation.

Calvin and Hobbes

by Bill Watterson




Figure 2 *Calvin appreciates the essential mystery of the fundamental arithmetic operations. Neurobiologists are math agnostics: they do not understand how the nervous system combines value-specifying signals to generate new value-specifying signals in a manner that mimics the fundamental operations of arithmetic--addition, subtraction, multiplication, division, and ordination--and some doubt that it in fact does so. (Copyright by Bill Watterson and reproduced with his permission.)*

The values of variables need not come from previous computations, however. Sometimes the values of symbols come from sensing and measuring devices sensitive to the value of the real-world variable the symbol denotes or refers to. Thus, the values for symbols of temporal intervals may come from timers, the values for time symbols from clocks, the values for distance symbols from a radar ranging device, the values for a direction symbol from a magnetic compass, and so on. A symbol-processing device requires a mechanism or a procedure that establishes which symbols denote which real world variables, and it generally requires calibration or checking procedures to insure that the values of those symbols do in fact accurately represent the values of the real-world variables to which they refer.

Supplying a value for a symbol is sometimes called **binding* a value to a variable. For example, in making a computation on a reverse-Polish calculator, you first key in one value and hit Enter, thereby binding the keyed-in value to the addend; then you key in another value (binding that value to the augend); and finally, you hit the '+' button to set the addition operation in motion. The electronic process that mediates the addition operation determines and displays a new value, which is the sum of the addend and the augend. (Calvin's problem is that he is not sure how much to credit what the addition operation does.)

The digital nature of contemporary calculators is irrelevant in the above illustration. The use of a slide rule to multiply serves as an equally good example¹. First, you move the initial mark (index

¹It is a painful reminder of my advancing age and the rapidity of technical progress that I must tell the contemporary student what a slide rule is or was. It was an analog calculator used by technical people for multiplying and dividing until quite recently, when it was superseded by the incomparably more powerful pocket calculator. It used logarithmic scales marked on both a frame and a slider. The slider, with its scale, moved within the frame, with its scale. It exploited the fact that $\log(x*y) = \log(x)+\log(y)$. A good slide rule was a lovely instrument, much treasured by its possessor.

point) of the slider so that it is opposite the value of the multiplicand on the frame (binding the value of the multiplicand.) Then you find the value of the multiplier on the slider (binding the value of the multiplier). Finally, you note the value on the frame opposite the value of the multiplier on the slider. This is equivalent to hitting the 'x' button; in doing this, you perform the multiplication operation.

In both examples of the compositional nature of conventional computation, a human operator supplies the values to be added or multiplied (operated on) and reads the results, but that is also irrelevant. When conventional computers are used to control processes, as they are in modern factories, automobiles, and appliances, the values for the symbols are supplied by sensors and timers, and the results of their computations are used to control actuators. Human participation is not necessary. If the brain is a computer, then it is a process-control computer; it does not require an **homunculus* --a little man in the head who operates the computer.

Finally, the use of numbers to be the values of the variables is also irrelevant. In electronic analog computers, voltages specify the values of variables such as wind velocity, roll-of-the ship, and so on. These voltages are fed into operational amplifiers configured either to add or to multiply their input voltages in producing their output voltage. What is crucial is that there is some machinery whereby the values can be bound to variables (symbols) within the symbol-processing device and whereby these values may then be combined with other values to generate still other values that may themselves enter into further computation. It is also crucial that there is some machinery that establishes what it is that a symbol denotes.

1.5 Neural nets

The compositional property is so obvious an attribute of conventional computation that the student may wonder why it is singled out for comment. I single it out because this property is missing in neural net "computation," or at least in many of the best known models of this style of computation (see Rumelhart & McClelland, 1986) . The rise of neural net theories in the eighties was in part a reaction to the view that the brain performed computations in the conventional (symbol-processing) sense. Some of the originators of the neural net approach to psychological

modeling argued that the brain is not a symbol-processing system (Rumelhart & McClelland, 1986), that it is a sub-symbolic system (Smolensky, 1986). The development of neural net theorizing represents a resurgence of the behaviorist view that the brain does not really represent the external world in any interesting sense of the term representation, it merely appears to represent it².

Artificial neural net models explain behavioral phenomena as the result of the flow of signals through a network of interconnected nodes, with a learning process of some kind that adjusts the strengths of the connections. Each node takes a sum of the weighted signals it receives. The strength of an input connection determines the weight that the input signal it carries is given. The node applies a sigmoidal (thresholding) transformation to the weighted sum of its inputs to determine the strength of its output signal. (See the chapter by Anderson in this volume for a fuller description of these ideas.)

In many nets, there is nothing that functions as a symbol. The strength of the signal sent out by an internal (hidden) node does not, generally speaking, correspond to or specify the value of a real-world variable (e.g., an actual distance or actual a temporal interval). Nor does the pattern of signals sent out by an array of internal hidden nodes denote or refer to any real world variable. Only the patterns at the input and output nodes denote or refer to anything, and, because these denotations are generally supplied by the human operators of the network there has been little attention in neural network models to problems of establishing reference or denotation. If you do not believe that the internal variables in the net denote or refer to real-world variables, then, of course, you are not concerned with mechanisms that establish and maintain the mapping from real-world variables to internal variables.

The patterns of connection weights that result from the learning process are sometimes said to constitute the net's knowledge of the world. However, there is no provision in the architecture of the network for using these patterns of weights as the inputs to computational operations. There is, for example, no

²An example of this approach are models in which a network learns to produce past tenses of verbs without having any notion of a verb (a word that symbolizes a particular act or relation) let alone any representation of the inflection rules of English grammar (Rumelhart & McClelland, 1987). This model stands in deliberate contrast to models in which the rules for inflecting regular verbs are programmed into the computer (a conventional symbol-processing approach).

provision for adding a pattern of weights that specifies one distance to the pattern of weights that specifies another distance in order to obtain a pattern of weights that specifies a third distance. A fortiori, there are no processes for binding values to the non-existent symbols; nor are there any operations on these banished elements of what some neural net theorists regard as a misguided approach to the understanding of mental processes (the conventional artificial intelligence approach, which assumes symbol-processing capacity).

1.6 Difficulty of reconciling symbol-processing with our current understanding of neurobiology

Neural net theorists take their inspiration from our current understanding of the nervous system, which is not easily reconciled with the view that the brain is a symbol-processing system, because we do not know how the elements of symbol processing might be implemented in the nervous system. The values a symbol-processing device uses in its computations are often not specified by a current input. The value of the cumulative displacement in the dead-reckoning process to be described later has this property, as do many other values in the navigational computations to be described. Thus, a symbol-processing device must have a mechanism for preserving the values of variables until they are used in a computation. In short, it must have a memory. The storage process is the process whereby the value to be preserved is entered into memory. It is called the **write* or (**store*) operation in a conventional computer. The **read* (or **fetch*) operation in a conventional computer is the process by which the values preserved in memory are made available to a symbol-processing operation, such as adding, multiplying or comparing. The mechanisms by which the nervous system stores and retrieves the value of a variable, if they exist, have yet to be identified.

Other fundamental elements of symbol processing are the mechanisms for combining the values of variables in accord with the laws of arithmetic and logic--the source of Calvin's incredulity (see Figure 2). If the nervous system contains processes equivalent to the fundamental operations in arithmetic and logic, the processes built into any conventional computing device, then those processes remain to be identified. Thus, a conceptual chasm separates our current understanding of the nervous system from those cognitive theories of brain function that posit symbol-processing capacity.

Many neural net theorists seek to eliminate this conceptual chasm between neurobiology and cognitive psychology by denying the assumption that the brain has a symbol-processing capacity, which is equivalent to the assumption that the brain computes in the conventional sense of that term. Our current inability to explain basic symbol-processing operations neurobiologically is also a major reason why neurobiologists tend to favor the view that the brain is not a symbol-processing organ.

The other possibility is that those mechanisms in the nervous system that mediate its symbol-processing capacity remain to be discovered or identified. If this is the case, then our current understanding of the nervous system is seriously deficient. The brain is the organ of mind. If the essence of mind is the processing of symbols, then we have not begun to acquire a serious understanding of that organ until we can account for its symbol-processing capacity.

In short, the stakes in the debate about whether the brain is a symbol-processing organ are large. Either much of current psychological theorizing rests on fundamentally erroneous assumptions about the nature of mental processes. Or, our understanding of the behaviorally important properties of the nervous system does not extend to those processes that are at the heart of all higher function. Small wonder that the debate sometimes grows heated.

1.7 Mental representations

Symbol-processing theories of mind and representational theories of mind are intimately related. Representational theories of mind maintain that the brain constructs mental representations and uses these representations in the decision processes that generate action. The mental representations exist in the form of computed values for symbols. Thus, for example, the mental representation of an animal's position consists of the value of its position vector, which is computed by its dead-reckoning mechanism (see below).

The definition of a mental representation can be given in a few well chosen words, but, like most definitions of complex concepts, the definition then requires extensive explanation. A **mental representation* is a functioning isomorphism between a set of processes in the brain and a behaviorally important aspect of the world. This way of defining a representation is taken directly from

the mathematical definition of a representation. To establish a representation in mathematics is to establish an isomorphism (formal correspondence) between two systems of mathematical investigation (for example, between geometry and algebra) that permits one to use one system to establish truths about the other (as in analytic geometry, where algebraic methods are used to prove geometric theorems).

An *isomorphism* is a relation between two systems. Each system is composed of entities and relations or processes involving those entities. Take for purposes of illustration a navigational computer as one system and the landmarks and goals represented on the computer's map as the other system. The position of the Los Angeles airport is represented in the computer by a number pair--the latitude and longitude values that, by convention, are used to represent positions on the surface of the earth in navigational computations. These number pairs are called *position vectors*. They are physically realized as bit patterns in the computer's RAM or by patterns of magnetization on its hard disk. The bit patterns in a RAM chip are not themselves positions in the world, they are bit patterns in a RAM chip. These vectors, or, if you like, the bit patterns that physically realize them, are the entities in one system, the *representing* system. The entities in the representing system are called *symbols*. The set of these entities, that is, the set of position vectors in memory constitute the computer's map. (To be useful, there must also be a list or table associating other information with these position values, e.g., the names of cities, but it is the position values that enter into navigational computations.) The corresponding entities in the other system--the *represented* system--are the actual positions themselves. These positions are, of course, not bit patterns, they are places in the world.

A navigational computer performs operations on position vectors. These operations yield answers to navigational questions, such as, Does Lake Superior lie on the great circle route between Los Angeles and Nairobi? These operations constitute the other essential aspect of a representing system--the operational or relational aspect. The corresponding aspect of the represented system is the set of geometric relations between places in the world. The three places just mentioned instantiate a collinearity relation because Lake Superior does lie on the great circle defined by Los Angeles and Nairobi.

The computer with its map (its set of position symbols) is the representing system. The actual places out there in the world, the places whose positions are symbolized inside the computer, constitute the other system. The two systems are isomorphic systems because the surveying procedures for determining numerical latitudes and longitudes and the program for performing navigational computations are together such that the computer says that three places lie on a great circle if and only if they really do. In this case, the computer says that Lake Superior lies on the great circle defined by Los Angeles and Nairobi if and only if the route that minimizes the distance flown on a trip from Los Angeles to Nairobi takes the airplane over Lake Superior. Finding the distance-minimizing route is the operation in the real world that is isomorphic to computing the great circle route. It is not easy to find the distance-minimizing route between Los Angeles and Nairobi by trial and error or any other readily performed non-computational real-world procedure, which is what makes computation so valuable in navigation.

The statement that Lake Superior lies on the great circle route if and only if the computer says it does not elevate the computer to the status of a reality-determining god. The if-and-only-if relation between the outcome of processes in the computer and the true state of the world arises only because of the care with which the surveying procedures and the navigation program are chosen and implemented. The if-and-only-if relation between the two systems is therefore fragile. It can fail with tragic consequences, as when the Korean Airlines 747 was shot down for flying over militarily sensitive portions of the Soviet Union. The plane, which was flying under the control of its navigational computer, was hundreds of miles off course as a consequence of a mistake in the after-take-off initializing input to the computer.

The navigational computations of animals also often fail, with fatal consequences. Thus, the processes in the nervous system that implement navigation are subject to intense natural selection. They are subject to natural selection because, like the computations of a navigational computer, they are *functioning* isomorphisms. The isomorphism between the representing system and the represented system is used by the representing system in coping with--surviving and reproducing in--the represented system.

1.8 Summary

A cognitive map is a mental encoding (symbolic representation) of the relative spatial positions of goals and landmarks. Its place symbols (position vectors in memory) are used in navigational computations, the processes that mediate the setting and holding of courses. If there is a cognitive map, then somewhere in the nervous system there must be neurophysiological states or activities that physically realize the position vectors that constitute the map (the neurobiological equivalent of the bit patterns in the RAM chip). These neurobiological position symbols are used in the neurophysiological processes that mediate the setting and holding of courses. If the cognitive map hypothesis is correct, then these neurophysiological processes will prove to be isomorphic to the computations that mediate conventional navigation in a navigational computer. That is, where the computation requires the addition of two variables (as does the dead reckoning computation to be described later), there will someday be found in the nervous system a process equivalent to addition. This process will take a neural signal that specifies the animal's cumulative displacement as of a moment ago and add to that a signal that specifies its most recent displacement to generate a signal that specifies its cumulative displacement as of right now. Thus, the use of a cognitive map, if it can be convincingly demonstrated, provides a paradigmatic example of the brain as a symbol-processing organ.

The advantage of studying animal navigation rather than other behavioral phenomena that have inspired intense debate between proponents of symbolic and non-symbolic analyses of mental function is that navigation is better understood from a computational standpoint than are most of the other areas in which this debate rages. Surveying and navigation have been intensively studied for thousands of years. While a proper humility in the face of possibly undreamed of limits to our understanding is always appropriate, it is nonetheless reasonable to believe that we may understand all there is to understand about navigation from a formal/computational standpoint. By contrast, our understanding of pattern recognition algorithms, language parsing algorithms, and language learning algorithms is rudimentary and incomplete. If we want to test whether the brain does or does not solve problems symbolically, it is best to work in an area where we believe that we may understand the "solution space," that is, the set of conceivable solutions.

2. Breaking the question down

The question whether insects have cognitive maps is too big a question to be addressed all at once. Like all major scientific questions, it needs to be broken down into a series of smaller, more readily answerable questions. Our extensive and long-standing understanding of geometry and navigation makes it possible to frame a series of specific questions.

The first two questions are can insects determine and remember angles and distances? Angles and distances are fundamental to the representation of position. If a device cannot determine angles and/or distances, then it cannot represent position; hence, it cannot have a map. On the other hand, if it can both determine and remember either angles or distances, then it has the basic ability to determine the values that define positions on a map. We can ask separately about angles and distances. However, it is important to know that from a formal standpoint, the two are intimately connected. Together, they constitute the metric properties of a geometry. (See Gallistel, 1990, Chap. 6 for a discussion of the grouping of geometric properties.)

The third question is whether insects do dead reckoning. Dead reckoning is the computation of net displacement (change in position) by either integrating velocity with respect to time, or by the discrete equivalent of this computation, which is adding up a series of small displacements to determine their sum. Dead reckoning is a fundamental mode of navigation because it enables one to move into unfamiliar terrain and still know the direction and distance back to home. By its very nature, it also generates a vantage-point-independent coding of positions. Thus, it may be used to relate one vantage point to another, which is crucial in computing a map of the terrain.

The fourth question is how insects hold a course? To hold a course, you must be able to continuously estimate which way is north. ("North" in this context just means some arbitrarily chosen but thereafter fixed reference direction.) One way of judging which way is north is by reference to a map of the terrain around you, which will identify for you the landmarks that lie to the north of your current position. Another way is by reference to the sun. Using the sun requires learning of the solar ephemeris function, which is the direction of the sun as a function of the time of day. Either of these methods implies the computation of a map of some kind.

The fifth question is how insects recognize a terrain from different angles and distances of view? The importance of this question is that in conventional navigation, such recognition is mediated by a map of the terrain. That is, knowledge of the three-dimensional lay of the land is used to interpret the two-dimensional projection of the terrain onto the retina. It is an open question whether it is possible to recognize the retinal projections of familiar terrains purely on the basis of the two-dimensional image, without computing aspects of the three-dimensional terrain of which the image is a projection. Thus, this question is intimately related to the question of whether insects have a map and, if so, what use they make of it.

The final question is whether insects have an integrated map, a map that enables them to compute the course between any two positions represented in their nervous system, whether or not they have ever traveled between those two positions in the past. If they do, then they have a map in the fullest sense of the term.

2.1 Can insects determine and remember angles?


Position on a plane may be defined by two distances (Cartesian coordinates), by an angle and a distance (polar coordinates), or by two angles (as in a three-point fix, see Figure 4). Thus, angles and distances are fundamental to the definition of position. If the nervous system cannot encode angles and distances, then it cannot encode position. If it can encode either angles or distances or both, then it can, in principle, encode positions. Thus, one fundamental question is whether the nervous system of an insect can encode an angle.

2.1.1. The dance of the honey bee symbolically expresses an allocentric angle. The discovery by von Frisch and his students of the dance language of the bees (von Frisch, 1967) provided exceptionally direct evidence that the nervous system of a bee could encode an allocentric angle, store the value of that angle in its nervous system, and use the stored value to control later behavior in which it communicated the value of the stored angle to other bees. When a foraging bee discovers a rich source of nectar and pollen, it records the angle formed by the position of the hive, the sun's azimuth, and the position of the food source. The sun's *azimuth* is the position on the horizon directly beneath the sun. An *allocentric* (other or elsewhere centered) angle is an angle that cannot be determined from sensory information available to an

observer at the point of observation. Because the bee often cannot see its hive from the nectar source, it cannot determine the angle between the hive, the source, and the sun from sensory information impinging on it while at the source. We shall see below that the bee computes this angle by dead reckoning using its sun-compass mechanism (Sections 2.3 and 2.4).

The bee takes some of the nectar into its crop, accumulates the pollen on its leg and body hairs, and then flies back to the hive. After entering the hive, it performs a waggle dance on the vertically hanging honeycombs out of sight of the sun. The waggle dance is in the form of a crude figure eight, with the wagging part of the dance constituting the segment common to the two loops. The angle of this waggle segment relative to the gravitational vertical equals the angle formed by the hive, the sun's azimuth, and the food source (Figure 3). This angle is called the *solar bearing* of the food source from the hive. In this way, the dance symbolizes the direction of the source from the hive. Other foragers follow the waggle portion of the dance, transferring to their own nervous systems the direction specifying value. They then leave the hive and fly in the indicated direction. These bees are the bees recruited to the source by its discoverer. When they return, they also dance, so that the number of recruits exploiting the source increases rapidly.

Figure 3. *The dance of the returned forager bee indicates the solar bearing (α) of the food source by the deviation of the waggle run from the vertical. The solar bearing is the angle between the sun's azimuth and the source. The solar bearing varies as a function of time because the sun's azimuth (or the compass bearing of the sun, that is, the angle θ) changes dramatically in the course of the day (sun's path). However, the dancing bee corrects for the passage of time, whether or not it can see the sun (see Figure 8), so that its dance continues to indicate the unvarying compass bearing (β) of the source.*



The direction of the source from the hive is called its compass direction (or compass bearing). The *compass direction* of something is the angle formed by two horizontal lines intersecting at the point of observation, that is, at the *vantage point* (the angle \square in Figure 3). The direction relative to the earth of one line is fixed and is referred to in all compass directions. Commonly, this fixed direction is the north-south direction, which is given by a line through the vantage point and parallel to the earth's axis of rotation. The other side of the angle is the line from the vantage point to the position in question, in this case, the position of the food source. At first glance, it would not seem that the dance of the returning forager specifies a compass direction, because it specifies the direction of the food source relative to the direction of the sun. The sun's direction is not fixed. The sun's direction relative to the earth's axis changes steadily (see sun's path in Figure 3). However, as we shall see in Section 2.4, both the returned forager and the bees that respond to its dance correct for the change in the sun's direction, whether they can see the sun or not. Thus, what the dance really conveys is the compass direction of the food source. The compass direction does not vary as a function of the time of day, whereas the solar bearing does.

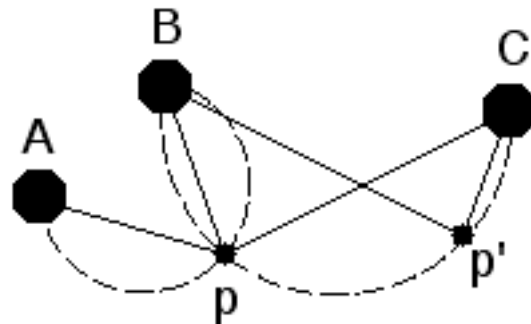
2.1.2 Bees record the angle subtended by the landmarks near a food source. More recently, Cartwright & Collett (1983) have shown that bees learn the angle formed by landmarks in the vicinity of a food source. An angle formed by contours or points visible from the point of observation is sometimes called an *egocentric* (self centered) angle, because it is the angle between the self, located at the point of observation (the food source), and two other points observable from that position. If the angle in question is formed by points at the extremities of an object (for example, the top and bottom of a tree), it is also called the *apparent size* of the object, because it is proportional to the size of the image of the landmark on the retina of the observer. Cartwright and Collett allowed bees to fly through a curtain into a featureless room and find food at a location defined by its relation to one or more tall black cylindrical landmarks. From trial to trial, the landmark configuration was translated within the room. (*Translated* means moved without rotating.) The food source went with the configuration (maintained its geometric relation to the landmarks). When, on a test trial, the food source was removed, the bees searched at the appropriate distance from the landmark(s). If a single landmark was used and it was made bigger on some test trials,

then the bees searched farther away as they would if they were locating the source by reference to the apparent size of the landmark. (Or as they would do if they were locating the source by its distance from the landmark and estimating distance from the apparent size of the landmark, a form of triangulation. There is no way of distinguishing these possibilities so long as there is only one landmark

When three cylinders were used, the bee was no longer influenced by their individual apparent sizes. It centered its search for the food source at a point defined by the angular separations of the landmarks. In other words, the bees behaved as if during training they had taken what in conventional navigation is called a three-point fix (Figure 4). A three-point fix specifies a position in the framework defined by three landmarks. It specifies that position by recording two angles: (1) the angle formed by the leftmost landmark, the point of observation, and the middle landmark ($\angle ApB$ in Figure 4); (2) the angle formed by the middle landmark, the point of observation, and the rightmost landmark ($\angle BpC$ in Figure 4).

Figure 4. *One way of defining a position is by a three-point fix, in which you identify three familiar landmarks (A,B & C) and measure the angles between them, with, for example, a sextant. The angular separation of only two landmarks does not define a position, because*

any one angle ($\angle ApB$, $\angle BpC$, or $\angle ApC$) defines an arc consisting of infinitely many points from all of which those two landmarks have the same angular separation. To illustrate, the angle $\angle BpC$ equals the angle $\angle Bp'C$, and this equality holds for any p' located on the dashed arc connecting B to C. Similarly, the angle $\angle ApB$ defines an arc of infinitely many possible positions for the point p. The intersection of these two arcs, however, uniquely defines the position of the point p. This is the mathematical principle behind the practice of taking three-point fixes. It also shows that the position of a point may be defined by two angles just as readily as it is defined by two distances.



2.1.3. Bees record the compass directions of the landmarks around a food source. If the only thing a bee recorded about the relation between a single landmark and a food source was the apparent size of the landmark (the size of its image on the retina), then given only a single cylindrical landmark the bee would search for the food source in an annulus (or ring) around the cylinder, because the arc of position defined by the apparent size of a cylinder is a complete circle centered on the cylinder. However, the bees in Cartwright and Collett's experiment, when trained with a single landmark, did not search in an annulus. Their search was centered on a single point, the point at which the landmark had the correct compass direction and apparent size. How the bees were able to determine the north-south direction while in the room out of sight of the sun is not clear. One possibility is that they are sensitive to the earth's magnetic field (Collett & Baron, 1994; Gould, 1980). Alternatively, as they entered the room, they may have noted the compass direction of the translucent curtain through which they entered. By remembering that the translucent wall was, say, the west wall, the bees could note and record the compass directions of landmarks in the room.

When there were three landmarks (cylinders), the bees were sensitive to the compass direction of the landmark configuration from the food source. If the configuration was rotated more than 45°, so that, from the point defined by the three-point fix, the array of landmarks all had compass directions seriously discrepant from the directions they had during the training trials, the bees no longer tried to find the food source by reference to those landmarks. Again, the analogy to conventional navigational practice is startling. It is standard practice to record the compass directions of landmarks as well as their angular separations. In using those landmarks to find the point again, you first check (with a bearing compass) that their compass directions are approximately correct before using a sextant to make more precise determinations of their angular separations (to get a precise three-point fix). If there is serious error in the compass directions of presumed landmarks when one is in the presumed vicinity of the point sought, then something is wrong, and there is no point in taking a three-point fix. Either one is not in the vicinity of the point sought or the landmarks are the wrong ones.

In summary, experiment establishes that bees record both allocentric and egocentric angles. Moreover, they can use recorded

angles to define a remembered position and to find it on subsequent visits. That is, they make three-point fixes. How they use those fixes to navigate to the point sought is an unresolved question (see Section 2.5)

2.2 Can insects determine and remember distances?


2.2.1. The dance of the honey bee symbolically expresses a distance. The dance of the honey bee also symbolizes the approximate distance to the food source; the greater the distance, the greater the number of waggles in the waggle segment of the dance. Thus, if the nervous system of a bee is itself devoid of symbols, it has the curious property of generating behavior that is itself symbolic. In fact the dance of the honey bee can be used to illustrate the concept of a functioning isomorphism. The two systems are: (1) the dancing bee and the recruits following it; (2) the direction and distance of the food source from the hive. The direction (an entity in the represented system) corresponds to the angle of the waggle segment with respect to the gravitational vertical (a symbol in the representing system). The distance (an entity in the represented system) corresponds to the number of waggles (another symbol in the representing system). The correspondence between elements of the dance and the geometry of the terrain outside the hive is used to control the food-seeking flight of the recruited bees, that is, in the navigational computations that control the distance and direction in which they fly prior to their beginning to search for the source. This use is what makes it a functioning isomorphism (as opposed to the many accidental isomorphisms that may be identified once you understand what an isomorphism is--an accidental isomorphism serves no purpose).

2.2.2. An internal distance code controls the length of the homing run of the displaced desert ant. The distance indication in the dance of the honey bee is crude. Evidence that the insect nervous system can encode a distance with some precision comes from experiments in which foraging desert ants are picked up just as they depart from a feeding source headed back to the nest and are displaced far outside their foraging territory (Wehner & Srinivasan, 1981). Unlike the displaced wasp in Figure 1, these displaced desert ants, which are on featureless and unfamiliar terrain, do not set a new compass course, one that will get them to the nest from their release site. Rather, they set off on the course they were on when captured. They run this course for a distance that is within 10% of their distance from their nest at the time of

their capture (Figure 5), then they break off this straight run to begin a tortuous, circling search for their nest. The Wehner and Srinivasan results establish that the nervous system of this species of ant can encode the distance of the ant from its nest out to about 50 meters with an accuracy of no worse than about $\pm 10\%$.

2.2.3. Locusts and bees compute distance by triangulation. In conventional surveying, before the advent of the laser interferometer, distances that could not readily be measured with a surveyor's chain measure were computed by triangulation, that is, from the change in the direction of the target produced by a measured displacement in the point of observation. The locust computes the distance it must jump to reach a target by the same principle (or by a closely related principle if, as may well be the case, the locust relies on image velocity rather than the change in image direction). Before jumping to a target, it engages in a curious behavior called "peering." While staring straight ahead, it shifts its upper body slowly back and forth, counterrotating its head as it

Figure 5. *Homing path in desert ant displaced to an unknown territory is traced on a 1 m gridwork. The ant was captured as it turned to head home from a feeding station 30 m south of its nest and displaced to an unfamiliar portion of the desert on which the gridwork had been painted. S = release point; N* = fictive nest (open circle marks place where it would have found its nest had it not been captured and displaced); O = sharp turn with which the ant initiates its search for the nest entrance (point of turn marked by filled circle). (Redrawn from Wehner and Srinivasan 1981, p. 318, by permission of author and publisher.)*




shifts so as to keep it oriented straight ahead. Ordinarily, the lateral shift in the locust's eye position shifts the angular position of the

target image on their retina. The distance to the target is the lateral distance of the locust's peering movement (the measured displacement = b in Figure 6) divided by the tangent of the change in the retinal angle of the target image (that is, the neurally measured change in the direction of the target = α in Figure 6). Experiments that cleverly alter the coupling between the target direction and the peering motion (Sobel, 1990) demonstrate that the locust's take-off velocity is determined by the change in target direction during these peering movements (Figure 6). This implies that the nervous system of the locust can perform the computation shown in Figure 6 (or its kinetic equivalent, see immediately below).

An insect that can estimate its speed can use the kinetic equivalent of the triangulation principle to determine its distance from a landmark: At any given visual direction, the distance of the landmark is inversely proportionate to the rate of change in its visual direction. This is called the motion parallax cue to distance. Recent experiments with bees flying through tunnels have shown that they maintain their distance from the tunnel walls by means of motion parallax. (Srinivasan, 1992) . In this case, the nervous system appears to be solving the equation $D = s[\sin(\alpha)/\alpha']$, where D is the distance from the landmark, s is the insect's measured speed, α is the angle between the insect's course and the landmark, and α' is the rate of change of the angle (the apparent angular velocity of the landmark).

Figure 6. *The locust measures the distance, D , from its point of observation to a target by some process equivalent to measuring the magnitude, b , of its lateral displacement during peering and dividing this variable by the tangent of the change in the direction of the target (α). The resulting distance estimate determines how hard it jumps. (Redrawn from Sobel (1990) by permission of the author and publisher.)*



In summary, experiment establishes that the insect nervous system can encode distance and use encoded distances to control a

variety of behaviors. The encoded distances may be derived either from stepping them off during dead reckoning (Figure 5) or by triangulating landmarks from an observation point (Figure 6).

2.3 Can insects do dead reckoning?

The experiment by Wehner and Srinivasan (1981) (see Figure 5) showed much more than that the nervous system of the ant could encode a distance. It showed that the nervous system of the ant was capable of dead reckoning, which is one of two general ways of estimating one's position in conventional navigation. (The other is called *piloting*; it unlike dead reckoning requires a map.) *Dead reckoning* is the process of estimating your position by integrating your velocity with respect to time. (Remember that velocity is a *vector* variable, not a *scalar* variable like speed, that is, it has two or three dimensions, depending on whether we mean velocity in a plane or velocity in a space. Put another way, it has a direction.) Dead reckoning may also be thought of as keeping track of the change in your position (your *displacement*) by summing successive small displacements. (Displacement is also a vector variable; it also has two or three dimensions.) Summing small displacements is the discrete equivalent of the integration operation. Summing successive small increments is the method that computers use when they compute integrals numerically rather than analytically.

In reaching the conclusion that insects can navigate by dead reckoning, behavioral scientists rely strongly on the assumption that, when it comes to navigation, we know all possible solutions. The ants in the Wehner and Srinivasan experiment were running across unfamiliar territory in which they could not observe any landmark with a known geometric relation to the point they sought. Under these circumstances, all forms of piloting are useless. (Piloting is estimating your position by reference to familiar landmarks, landmarks whose geometric relation to your goal is given on a map.) Nonetheless, the ants began their search almost at the point where their nest should have been (the fictive nest site). There is only one known process that could mediate this behavior, dead reckoning.

The conclusion that ants can navigate by dead reckoning is itself strong evidence that the brain of the insect is a symbol-processing organ. Dead reckoning is a classic example of a simple well-understood functioning isomorphism between a computational system and the spatial world. The entities in the represented world

are actual velocities or actual displacements. The symbols in the representing system are the numbers that the navigator uses to represent these velocities or displacements. In the case of the ant brain, the symbols are the signals that the nervous system uses to represent the ant's velocity or displacement. The relation in the represented world is that one's net displacement (one's current position relative to one's starting position) is the combined result of the step-by-step displacements that one made in moving from there to here. The representing system computes the net displacement by summing the symbols that represent the intervening displacements or integrating with respect to time the signal that represents the velocity. The current sum of the displacements (or the current value of the integral) represents one's present position. The representation of one's current position is in turn used to determine one's course and when to stop pursuing that course and start searching. It is this use of the integral or sum to control course setting that makes the procedure a functioning isomorphism.

The use of the current value of the sum or integral to determine the current direction of travel is demonstrated in a recent extension of the Wehner and Srinivasan experiment by Schmidt, Collett, Dillier, and Wehner (1992). They placed obstacles in the path of ants running homing courses to fictive (non-existent) nest sites across unfamiliar terrain. The obstacle forced the ant to deviate from its course. When the ant cleared the obstacle, the direction of the course required to reach the place where its nest should have been (but wasn't) was altered. That is, to reach the position of the fictive nest, the detoured ants had to set a new course. The closer the obstacle to the fictive nest site, the greater the course correction required to head toward that point. On clearing the obstacle, the ants did make course corrections in the right direction and of approximately the right magnitude. The closer the obstacle to the putative location of the nest, the greater the course correction the ants made.

The challenge posed by these findings to those who deny symbol-processing capacity in brains is to come up with a process that looks like dead reckoning but really is not. Can a neural net be created that does a non-symbolic or sub-symbolic form of dead reckoning? In such a net, there should be nothing that represents either the animal's velocity or its displacements (no velocity or displacement symbols), nothing that represents its current position (no position symbol), and no process equivalent to the summing of

displacements or the integration of velocity (no symbol processing). Still, the net should enable the ant to run a straight course across unfamiliar terrain for a distance equal to its distance from the nest, making suitable course corrections after enforced detours. Because dead reckoning is computationally well understood and experimentally well documented, it poses an interesting challenge to models that attempt to explain computation-like performance without positing explicit computation.

2.4. How do insects hold a course?

In conventional navigation, there is an important distinction between setting a course and holding a course. A course is the trajectory that takes you from where you are to your goal (or, at least, a sub-goal). Over modest distances (distances measured in hundreds rather than thousands of miles), the trajectory to be followed will traverse close to the minimum possible distance while maintaining a fixed compass direction, a fixed angle of progress relative to north. For this reason, *course* is used in two senses in traditional navigation--the trajectory followed (as in "plot a course") or the compass direction to be adhered to (as in "set a course"). When there is possibility of confusion, the latter is called the compass course. The compass course to your destination is simply the compass direction of your destination³. Thus, setting a course requires determining the compass direction of the destination from your current position, which in turn requires knowledge of the geometric relation between the destination and your current position. Holding a course, by contrast, requires knowledge only of the direction of north and the compass direction to be maintained. Thus, the problem in setting a course is to determine where you are relative to your destination, while the problem in holding a course is to know which way is north. Only if you know at each moment which way is north can you know whether you are on the compass course you are trying to maintain.

The problem in knowing from moment to moment which way is north is that as you move, the direction of the surrounding terrain features (landmarks) changes. If the tall tree lies due north when you start out toward the east, it won't when you have gone a few

³Over long distances, distance-minimization and a fixed compass direction are mutually incompatible properties of a course. The courses flown by airplanes change their compass direction steadily, which is necessary in flying most great-circle (distance minimizing) routes.

tens to a few hundred yards. Thus, maintaining your angle of march relative to the tall tree will not maintain your heading. This is the problem of the parallax motion (change in direction) of terrestrial landmarks. The farther away a landmark is, the less the change in its direction for a given distance of march and a given initial direction of the landmark. Also, the closer a landmark is to your course, the less it will change its direction as you march toward it. Thus, distant landmarks and landmarks that lie close to your course are the best landmarks to use in holding a course. However, in many terrains, no terrestrial landmark is sufficiently distant to serve as a convenient directional reference over the whole course.

2.4.1 The sun-compass mechanism. The sun is so distant that it has negligible parallax motion. And, in the absence of overcast, it is extremely salient. Thus, it is perhaps not surprising that ants and bees, in common with many other animals, hold a course by reference to the sun. The mechanism that enables them to do this is called the sun-compass mechanism. That ants hold a course by reference to the sun's direction is easily verified by finding an ant pursuing a straight course, blocking its direct view of the sun, and holding up a pocket mirror to reflect the sun's image onto the ant from another direction. The ant changes its course accordingly (Santchi, 1913)

The sun-compass mechanism has nothing to do with a magnetic compass. It is used in lieu of magnetic compass. A compass mechanism is any mechanism (e.g., a gyrocompass) that indicates a fixed direction with respect to the earth's axis of rotation, thereby making it possible to hold a compass course. Because a magnetic compass is so often used for this purpose, the term compass has come to mean magnetic compass in everyday usage.

What makes the sun-compass mechanism remarkable is that it compensates for the time-dependent change in the compass direction of the sun, even during periods when the animal cannot see the sun (Figure 7). To do this, the sun-compass mechanism requires an internal time signal (a *clock*) and a record of the sun's compass direction as a function of the time of day (a *solar ephemeris function*).

2.4.2 *The circadian clock and the solar ephemeris function as further examples of brain-world isomorphisms.* The rate at which the sun's compass direction changes varies dramatically as a function of time of day, season of the year, and latitude. Near the Arctic circle, at the summer solstice, the direction of the sun changes by about 15° every hour around the clock. On the equator, at the spring or fall equinox, the sun is due east from sunrise until noon. During this time, its *elevation*, the angle from the horizon straight up to the sun, increases steadily, but its azimuth remains constant. After noon (the moment when the sun passes through its *zenith* or maximum elevation), the sun is due west; thus, its compass direction changes by 180° in a matter of minutes. To use the sun to hold a compass course, an animal (or any navigating device) must know the time of day and the solar ephemeris. Because the solar ephemeris is locale- and season-specific, it must be learned.

Insects know the time of day because they, like most living organisms, have a *circadian clock*, neural tissue whose intrinsic characteristics are such that it completes one cycle of activity approximately every twenty-four hours. One function of the circadian clock is to time an animal's activities—for example, the onset of its foraging—so that they reliably occur at a certain phase in the earth's daily geophysical cycle. The circadian clock—an internal cycle whose period approximately matches the period of the earth's rotation and that functions to adapt an animal's behavioral cycles to the earth's cycle—is perhaps the simplest of all the well-documented functioning isomorphisms between brain processes and the world. The neural signal from the clock symbolizes the time of day, or, if you like, the momentary state of the earth's rotation.

In order for the circadian clock to symbolize correctly the time of day, there must be a process that synchronizes the neural clock with the earth's rotation. Otherwise, inevitable small differences between the rate at which the neural clock runs through its cycle and the earth's rate of rotation will cause the neural clock to drift out of phase with the day-night cycle. The process that adjusts the phase of the internal clock each morning at dawn in response to the rapid increase in light intensity is called *entrainment*, and it relies on light receptors specialized for the requirements of the entrainment mechanism (Foster, 1993). The function of the entrainment mechanism is to insure that the neural

clock is in phase with the day-night cycle. It insures that the brain's clock tells the right time.

The sun-compass mechanism makes more sophisticated use of the isomorphism between the neural clock and processes in the outside world. It records in memory the position of the sun relative to the terrain surrounding the animal's nest as a function of the time of day indicated by the brain's circadian clock. It records, for example, that the sun is over the tall tree when the internal clock reads 9:00 a.m. and that it is 70° clockwise from the tall tree when the internal clock reads 12 noon. This recorded information about the azimuth of the sun relative to prominent features of the terrain constitutes the solar ephemeris function, which is stored in the brain and used by the sun-compass mechanism. The information about the terrain surrounding the hive and its relation to the sun's trajectory is probably recorded during the brief survey flights that both bees and wasps make when they first leave a nest (Becker, 1958; Free, 1955; Manning, 1956; Tinbergen & Kruyt, 1938; Zeil, 1993b)

Dyer (1987) used the dance of foraging honey bees to obtain a direct behavioral demonstration of their knowledge of the solar ephemeris. He positioned a hive, to take one example, on the open south side of a east-west tree line. (Thus, the forest was on the north side of the hive--see Figure 7A.) On several successive days, he trained bees from that hive to forage at a feeding station 180 meters to the east (direction of training station from hive of 90°, range 180 m). The bees flew to the food with the trees on their left and returned with the trees on their right. One morning, when the sky was heavily overcast so that the bees could not see the sun, Dyer moved the hive to another similar tree line, also running east-west, but with the open (field) side to the north rather than the south (Figure 7B). When they flew from this location with the trees on their left, they flew west rather than east. He set up two test feeding stations, one 180 m to the east of the hive and one 180 m to the west. Both stations had very sweet sugar solutions, so that bees returning from them would dance vigorously to recruit new foragers.

When the bees previously trained to a feeding station in the east emerged from the hive at the test site under heavy overcast, they invariably flew west. That is, they followed a compass course of 270° rather than 90°. Thus, they all came to the feeding station in

the west (on the right in Figure 7B) rather than the station in the east. When they returned, however, they danced as though they had flown east (that is, to the feeding station on the left in Figure 7A) and as though the sun was where it would have been if the test tree line were the training tree line. This was, of course, 180° away from where the (invisible) sun in fact was during their test flights.




Figure 7. Schematic rendering of one of Dyer's (1987) experiments. **A.** The bees were trained with the sun visible, the food to the east of the hive, and a tree line to the north of the hive. Here we are looking to the south from above the forest. The hive and food table are nearly hidden behind the trees. Note that in flying to the food, at noon, the bee would have the forest on its left and the sun over the open meadow on its right. The hourly position of the sun is shown from 8:00 a.m. to 4:00 p.m (=16:00). **B.** The bees were tested with the sun hidden by heavy overcast (except in the late afternoon). In this new location, the tree line lay to the south of the hive. Feeding tables were positioned both east and the west of the hive. Here we are again looking south toward the hive, but now we are over open meadow looking at the forest behind the hive. Note that in flying to the food in the east at noon, the bee would have the (invisible) sun on its right as at the training location, but it would also have the forest (rather than the meadow) on its right. If the bee assumes that it is still at the training location and judges the sun's compass bearing by reference to the tree line, then its estimate of the sun's compass bearing will be wrong by 180°.

When Dyer first opened the hive at the test site at 10:30 a.m., a bee emerging from the hive and orienting toward the dark forest (hence south) would have the (invisible) sun 45° to its left. At that same time, back at the training site, for a bee likewise facing the forest (hence north), the sun would lie 135° to its right (see Figure 7). The bee at the test site flew away from the hive with the forest on its left and the open meadow on its right. If it were where it thought it was (back at the training site), this course would be 315° clockwise from the sun (45° to the left of it). In actual fact, however, the course flown was 135° clockwise from the (invisible) sun. The bee found exceptionally sweet nectar at the feeding station and returned to dance vigorously in the hive.

In the hive, straight up (vertical) symbolizes "toward the sun." The returning bee ran the waggle segment of their dance at an angle of 45° to the left of vertical, 180° away from the correct solar bearing of the food station just visited. As the day wore on, the waggle dance systematically changed its orientation to the vertical, rapidly at first, then more slowly, mirroring the rapid change in the compass direction of the sun around noon and the slowing of this change as the sun moves into the western sky in the later afternoon (Figure 7) By five in the afternoon, the waggle segment was oriented straight down, indicating that the food source was due east, dead away from the late afternoon sun. In fact, of course, it was due west, in the direction of the sun.

The bees recruited by these dances arrived at the same food station as the dancers returned from, because they, too, were mistaken about where the sun was relative to the tree line. What is necessary for the dance to succeed on a cloudy day when no one can see the sun is that both the dancer and the interpreter of the dance share the same belief about where the sun is at the moment (share the same ephemeris function). It does not matter whether their shared ephemeris function is correct. (It is, by the way, precisely the possibility that the process for binding a value to a symbol may bind the wrong value that makes it possible to have false beliefs.)

Another thing to note about these experiments is that the compass course the foragers fly is the sum of two angles, one of which varies as a function of time and one of which does not. The angle that does not vary with time is the compass direction of the food source, or, if you like, its solar direction at the time of the

bee's last visit, which may have been hours or days ago. The other angle is the current direction of the sun. This angle, whose value varies as a function of time, is specified by the ephemeris function. To hold its course, the bee must know both these values and combine them appropriately. In short, the data in Figure 8 imply that the nervous system of the bee can sum the values of angles.

Again, the existence of a time-compensating sun-compass mechanism in both vertebrates and invertebrates poses a challenge to those who believe there are no symbols and no conventional (that is, symbol-processing) computation in the nervous system. Is it

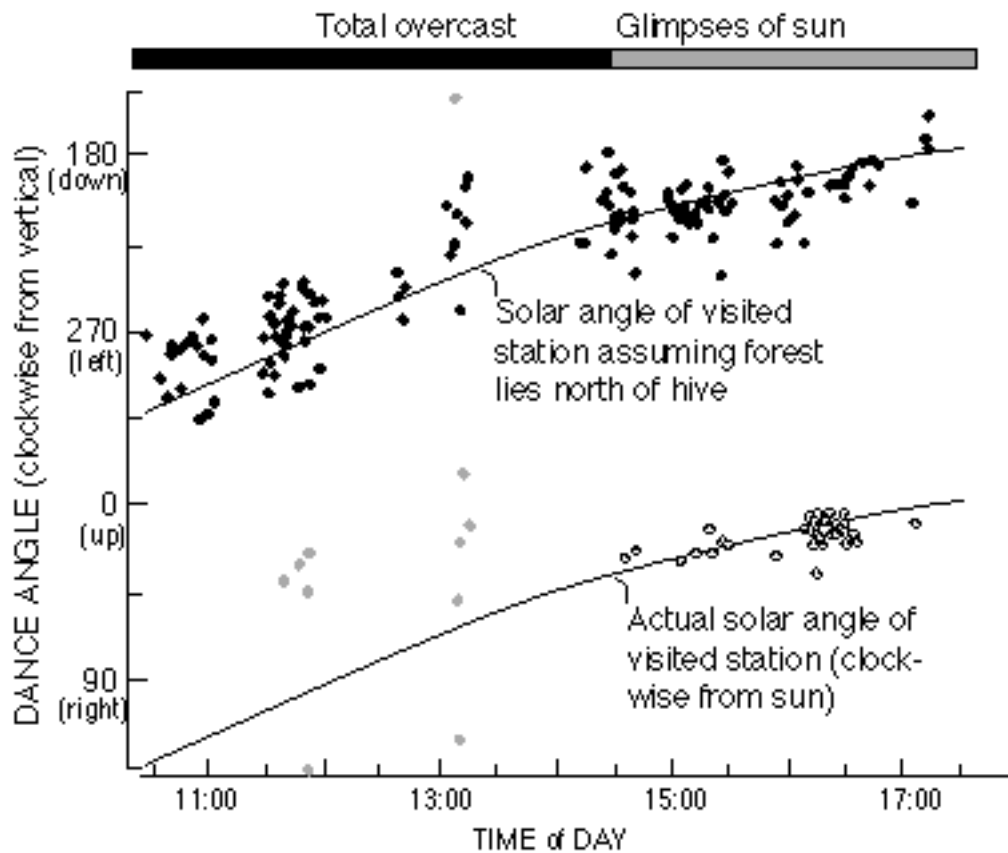


Figure 8. *Time-dependent control of the angle of the bee's recruiting dance by the internal ephemeris function. The bees had been trained to fly to a station lying to the east of the hive along an east-west tree line with the forest lying to the north of the hive. On the day these dance-angle data were gathered, the hive was moved to another similar looking tree line with the forest lying to the south of the hive. The previously trained foragers left the hive and flew west to a feeding station. The data give the angle of their recruiting dance on their return to the hive. For the first several hours after*

the hive was opened at 10:30 in the morning, the bees could not see the sun. Because they confounded the forest at this test site with the forest at the training site, their dance angle was consistently wrong by 180°. As the day progressed, they changed this angle in fairly close accord with the changing direction of the (invisible) sun if one allows for their consistent 180° error (by moving the solar angle curve up 180°). After 14:30, the sun was intermittently detectable (by the experimenter), and some of the dances began to show the correct solar angle (open circles). The gray data points are aberrant outliers not readily explicable by any assumptions about where the bees thought the sun was. Slightly modified from Dyer (1987, Fig. 3, p. 626). Used by permission of the author and publisher.

possible to have a neural net (or some other non-symbol-processing device) that generates behavior like that shown in Figure 8 without having an internal clock to symbolize the time of day? Is it possible for a device to generate behavior like that without having anything in it that specifies position as a function of the time indicated by its internal clock? What does the non-symbolic realization of an ephemeris function look like? Can a device that is incapable of summing the values of two angles hold a remembered compass course by reference to the solar ephemeris over many hours during which the sun is not visible? What is the process that looks like summing the values of angles but is not really?

2.5 How do insects recognize a three-dimensional configuration from different angles and distances of view?

In order for a configuration of landmarks in the vicinity of a goal to be useful in setting a course toward the goal, the configuration must be recognizable when the animal is not at the goal. That is, the animal must be able to match the images that this configuration casts on its retinas when viewed from other vantage points to the images cast by the configuration when viewed from the goal. When bees leave a food source, they make a survey flight similar to the survey flight they make when first leaving their hive (Lehrer, 1993). It is likely that the purpose of this survey flight is to record the compass directions and distances of the landmarks surrounding the food source (Collett & Baron, 1994). In any event, it has been convincingly shown that bees and wasps locate food sources by reference to surrounding landmarks (Cartwright & Collett, 1983; Cheng, Collett, Pickhard & Wehner, 1987; Collett, 1992). Experiments in which insects find goals by reference to nearby

landmarks have repeatedly shown that the insects approach the goal from different directions on different occasions (Cartwright & Collett, 1983; Collett & Land, 1975)

Experiments have also shown that ants recognize familiar landmarks they pass along their route, and they adjust their course so as to pass them on the correct side (Collett, Dillmann, Giger, & Wehner, 1992) . When placed anywhere from 5 meters to 15 meters in front of a familiar landmark along their habitual route, they immediately begin to deviate to the appropriate side of it (i.e., to the left or right depending on how they are accustomed to pass it). However, they will pass that same familiar landmark at a distance of anywhere from 2 to 7 meters depending on how far lateral to it their compass course from the release site runs. That is, while they try to pass a landmark on the correct side, they do not adjust their beam distance from it (their distance from it when it lies directly to one side). These findings imply that insects can recognize a landmark despite changes in the size of its retinal image by factors as great or three- or four-fold.

All of these findings indicate that insects do recognize landmarks from different distances and different angles. Thus, they do not simply match the current retinal image against a template of fixed size and shape painted on the retina in a fixed position. When they are still far from the goal and the shape and size of the terrain image is very different from what it is in the immediate vicinity of the goal, they somehow overcome the discrepancy between the current and remembered image, so that they can use the landmarks to instruct them as to what course to set to reach their goal . The question arises, How do they recognize the landmark configuration from a vantage point different from the vantage point at which the configuration was committed to memory?

This is a classic question in pattern recognition. Pattern recognition processes are no better understood in insects than they are in mammals. No one imagines that we understand all the different algorithms that might in principle be used to solve this problem. Nonetheless, one may recognize at least two broad approaches to the problem, terrain matching and view matching. The terrain-matching approach requires the computation of a three dimensional representation of the terrain during the animal's initial experience of it. The terrain map is committed to memory (along with perhaps also some two-dimensional views, hereafter called

snapshots). In recognizing the mapped terrain on a subsequent occasion, the perceptual system computes a three-dimensional representation of the terrain currently being viewed. The recognition process--the search for a match between what is currently viewed and a record of what has previously been viewed--searches for a terrain match. If this is how the nervous system of an insect solves the problem, then another function of the map is to mediate landmark recognition.

Notice that I assume that the cognitive map has the same dimensionality as the terrain it represents, namely three. Before describing the alternative approach to landmark recognition--view matching--I digress briefly to comment on the dimensionality of the cognitive map. Conventional maps are two-dimensional because the paper on which they are drawn constrains them to be so. The third dimension of a terrain (elevation) is represented, if at all, by an artifice such as the iso-elevation contour lines on the topographic maps that hikers carry. Any hiker who has strayed from the beaten path knows that two dimensional maps are useless for route planning, let alone terrain recognition. You need to know all three dimensions of the terrain, both to recognize it when you see it and to navigate through it. In other words, you need to know the shape of the environment. It takes training to extract the shape of the environment from a topographic map. The only maps that give the shape of the environment in a readily apprehended form are plastic terrain-relief maps, which are made from molds that reproduce the actual shape of the terrain (albeit with an exaggeration of elevation to make it more salient). The mold is itself produced by a computer-controlled molding process using the three-dimensional representation of the terrain in computer memory. The on-board computer in a terrain-recognizing guided missile also has a three-dimensional terrain map (elevation as a function of latitude and longitude). When spatial information is symbolized electronically rather than on paper, the recording medium does not constrain the dimensionality of the encoding. I assume that the neurobiological recording medium also does not constrain the dimensionality of the spatial information encoded in the brain, and that the brain has a three-dimensional terrain map analogous to the ones in the computers of map companies and guided missiles.

The other approach to terrain recognition, view matching, tries to match previously recorded two-dimensional retinal images (snapshots) to the images currently on the retina and to deduce

from the character of the mismatch between the images the direction in which the animal should move. This latter approach is favored by most specialists in insect navigation (Collett, 1992; Dyer, Berry, & Richard, 1993; Menzel, et al., 1990; Wehner, 1992; Wehner & Wehner, 1990) . The intuition that motivates models of this class is that the insect nervous system is not capable of the complex computations required to derive a terrain map from visual input. The nervous system of the insect is assumed to rely more on feats of memory than on complex computations. In general, in conventional computation, there is a trade-off between the sophistication and complexity of the algorithms used (computational complexity) and the amount of memory used, because almost any computation (including the elementary operations themselves) can in principle be done by table lookup. Table lookup is computationally simple but it requires prodigious amounts of memory.

View-matching models of landmark recognition in insects have been developed by Cartwright and Collett (1983, 1987) and most recently elaborated on by Collett (1992) . These view matching models do, however, require some computation. And, as they have been elaborated, they have made greater and greater use of what amounts to the computation of geometric features of the terrain that generated the images being matched. In particular, they make essential use of both the compass directions and the distances of terrain features corresponding to components of the images. Compass directions and distances are properties of the three-dimensional terrain, not the two-dimensional retinal projection of that terrain. Thus, these models do make essential use of the computed information about the three-dimensional shape of the terrain that generates the retinal images. However, it is possible that they do so by storing in memory many different orientation and distance-specific views, then searching their memory for a view that matches the view they currently have. Again, notice that the prodigious use of memory can substitute for complex computation. Thus, an interesting implication of this approach is that the nervous system has an enormous memory capacity; gigabytes of bit-mapped image information are assumed to reside in brains the size of pin heads. This is not necessarily as implausible as it may sound, because we have no idea what the physical basis of memory in the nervous system is. If the values of variables (in this case the values of retinal pixels) are recorded by changes in, for example, the

conformations of individual molecules, then megabytes of information can be stored in individual nerve cells.

In view-matching models of landmark recognition, the insect is presumed to know by dead reckoning approximately where it is and its compass heading. Thus, it is presumed to know approximately which previously recorded landmark images it should seek to match the current image against. This knowledge is called route-based knowledge of landmarks. The idea is that the position and heading signal from the dead-reckoning mechanism summons up stored images of landmarks that should be seen at that point along the route. The images actually falling on the retina are then compared to these stored images. Again, there is a close analogy to conventional three-dimensional navigation. Navigators routinely use their dead-reckoned chart position to tell them what terrain features they should be on the lookout for.

In the original Cartwright and Collett model, the snapshot (stored retinal image) against which the current retinal image was compared had no fixed position on the retina. The snapshot behaved as if it were pasted onto the wheel of an internal magnetic compass such as that sometimes found mounted on the windshield of a car. If, at the goal, where the snapshot was taken, the compass direction of the imaged landmarks was east, then the snapshot was pasted onto the east side of the bee's internal computational compass wheel. Pasting the snapshot onto a computational compass "revolving" behind the retina makes the retinal discrepancy between the current image and the snapshot independent of the bee's orientation. As the bee turns to point in different directions, the image of the world moves on its retina, but, by means of computational mechanisms in the bee's nervous system, the snapshot undergoes a similar motion relative to the retina. Thus, what the image-matching process is really doing is comparing the compass direction of a landmark in the snapshot with the compass direction of what it presumes to be that same landmark in its current view. Cartwright and Collett did not attempt to model the mechanism that kept the snapshot correctly oriented, but such a mechanism in a conventional image-processing computer would require massive rapid computation. More recently, Collett and Baron (1994) have suggested that the bee stores retinal snapshots made when the bee was pointing in each of several different directions. Which snapshot the bee uses in seeking a match to the current image depends on which direction it is currently pointing.

Note again that the massive use of memory can obviate complex computation.

In the Collett Cartwright and model, the comparison of compass-oriented (and distance filtered, see below) snapshots to the current retinal image generates movement vectors that reduce the discrepancies in the compass directions. If the current image lies to the left of the snapshot, the bee moves to the right. If the current image is smaller than the snapshot, it moves ahead.

In the Cartwright and Collett models, not all components of the snapshot and the current image are processed at any given distance from the landmarks. In approaching the goal, the bee is assumed to use several different snapshots, which were obtained using distance filters, so that one snapshot contains only the images of landmarks close to the goal, while another contains only images of landmarks far from the goal. The snapshot with only the more distant landmarks on it is retrieved when the bee is far from the goal, the snapshot with only near landmarks when it is close to the goal. The image currently on the retina is also assumed to be filtered on line, so to speak, so that only the far components of the current image are compared to the snapshot of far landmarks and only the near components to the snapshot of near landmarks. That is, in conformity with standard navigational practice, the model relies on the compass directions of landmarks far from the goal when the bee is far from the goal and the compass directions of landmarks near the goal when the bee is near the goal.

Filtering on the basis of distance is a computationally intensive operation. Moreover, it extracts a property of the three-dimensional terrain--the distance from the vantage point of a terrain feature corresponding to some component of the retinal image. In filtering the snapshot and the current image, the Cartwright and Collett model makes the computations that would be required to derive a three-dimensional representation of the terrain. Thus, although the landmark recognition process is far from understood, current formally specified models of it all assume that it is based on computing geometric features of the terrains that generate the images on the retina, in particular the distances and compass directions of terrain features. The most interesting difference between the approaches is that the terrain-matching approach makes more limited use of memory at the price of more complex

computation, while the view-matching approach reduces computational complexity by making prodigious use of memory.

2.6 Do insects have an integrated map?

On any given day, a foraging bee may visit several different sites. It may even visit several sites in a single foraging flight. Orchid bees may visit as many as 40 different orchids at widely scattered sites along a 20 kilometer long "trap line" in the course of a single foraging flight (Janzen, 1971) . Thus, the brain of a bee must contain at any one time distance and direction information about many different sites, as well as information about the distances and directions of the landmarks close and far from those sites. Moreover, this route information cannot be just information about the routes from the hive to individual sites. Some of it must be information about the route from one site to another, because bees are known to visit more than one site without returning to the hive between visits.

When it is realized that the insect brain contains at any one time a great deal of learned positional information, the following questions arise: How integrated is the information about different routes and their associated landmarks? Is each route a separate file, so that the positional information encoded in connection with that route is only usable when the bee is following that route, like the strip maps that automobile clubs provide to their members? Alternatively, is all positional information—both the information derived from dead reckoning and the information about landmarks computed during survey flights—recorded in a single, common coordinate framework? If so, then the insect brain forms a integrated map of its environment (as opposed to a series of route-specific strip maps)⁴.

An integrated map makes it possible to do course-setting and holding computations that draw on information obtained in connection with different routes. Thus, evidence that insects can use distance and direction information obtained while following one route in order to follow some other different route or to compute a new route is evidence in favor of a integrated map. Three sorts of experimental findings are relevant:

⁴For a computational model of how a integrated map may be constructed, see Gallistel (1990))

2.6.1. Bees can, at least sometimes, use their knowledge of the distance and direction of Site A from their hive and the distance and direction of Site B from their hive to compute the compass course from A to B. Gould (1986) did experiments that formulated the integrated cognitive map question in precisely these terms (see also Gould & Gould, 1988, and Gould, 1990). He worked with bees from a hive in a field flanked by woods (Figure 9). The bees were accustomed to foraging in this field. He set up a feeding station near the hive until he had a squad of individually marked bees coming to forage at it, then he moved it in small steps down a flowerless path in the woods to the west northwest of the hive. (The entrance to the path is not shown in Figure 9.) He moved it in steps along this path to the small flowerless clearing in which the Feeding Station in Figure 9 is shown (B). He observed that the bees generally arrived at this station from the direction of the hive and departed from it also in that direction. Thus, the bees had been induced to visit by direct flight from the hive a site they were unlikely to have visited in the course of their foraging in the field. Gould then captured these same bees as they emerged from the hive en route to the Feeding Station. He transported them to a release site toward the south end of the field (A in Figure 9), released them, and observed the direction in which they departed from the release site.

The clearing (B) could not be seen from the release site (A). At A, the bees could see that part of the field border over which they had to fly to reach the clearing. This was approximately the same segment of tree line over which they had to fly en route from the hive. However, because of the change in point of view, the view of the forest edge bore no readily discernible resemblance to the view of that same edge seen from the hive (see Gould, 1990, pp 95-97). The course the bees were on when captured was 270° . The compass direction of Site B from Site A was 330° . The average departure direction of the released bees was within a few degrees of this (the correct) course and highly significantly different from the course they were on when captured. All of the released bees arrived at the feeding station, and they did so in about the same amount of time it normally took them to cover a similar distance from the hive to the feeding station along their trained course (about half a minute).

When Gould captured bees as they returned to the hive from various and unknown directions and released them at Site A, they departed on a course for the hive, rather than on a course for the feeding station in the clearing. Thus, bees released at Site A could

set a course for different destinations depending on what their destination was when captured. Similar results were obtained when the role of the sites was reversed, that is the bees were released at Site B to find a feeding station at Site A. This is strong evidence that bees can set a novel course computed from information obtained on different routes at different times. That is, it is evidence in favor of a integrated map and the associated computational capacity.

Unfortunately, Gould's (1986) results are not consistently replicable, for reasons that remain to be determined. Dyer (1991) partially replicated it. In one of his experiments, his feeding station was down in a quarry and his release site was up in a field where the bees had been accustomed to forage prior to their training to the feeding station. The bees captured as they left the hive bound for the quarry site and released up in the field, set a course for the

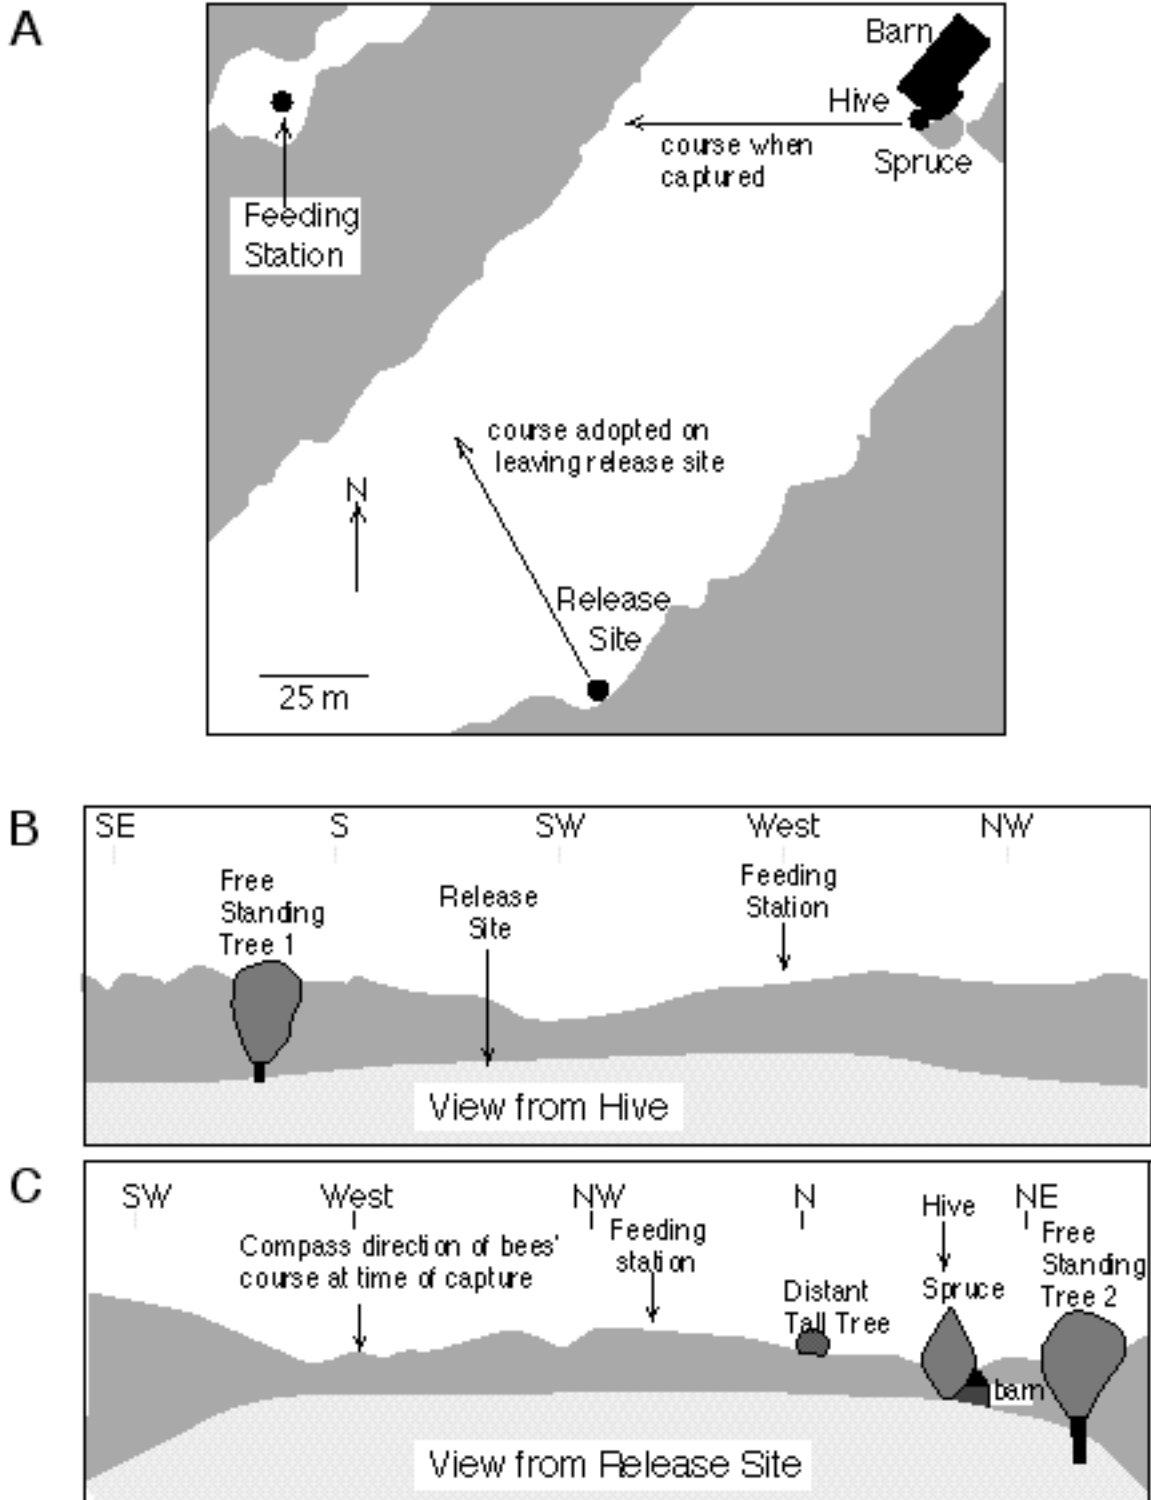


Figure 9. Map of the locale for the capture and release experiments of Gould (1986). Gray areas indicate woods. [Redrawn with minor modifications from Fig. 5 of Gould (1990) by permission of the author and publisher.]

quarry site. This replicates Gould's basic finding in his comparable experiment. In both cases, these bees could see from the release site a distant forest line that was close to the course they had intended to follow. Also, in both cases, this tree line looked very different from their release site, and they had to approach it from a very different angle, so it is difficult to see how this landmark could be recognized from the new vantage point by view matching, or how, in the absence of an integrated map the bee could determine what course to set relative to this landmark. Thus, both Gould's result and this similar result of Dyer seem to imply a terrain map.

However, when Dyer trained bees to feed at the site up in the field and then released them down in the quarry, they did not depart in the direction of the field site. Some departed on a course back to the hive, while the majority departed on the same compass course they were on when captured. This fails to replicate the comparable Gould experiment (the one in which he released his bees in the clearing in the woods). In these two cases, the bees could not see any distant landmarks at the release site. The failure of the bees released down in the quarry to replicate Gould's result from his bees released in the clearing is itself only a partial failure, because Dyer's bees, although, they departed in the wrong direction, mostly ended up at their originally intended destination (as did Gould's).

In another experiment, Dyer (1991) first trained bees to a feeding station in the quarry, then moved the station in steps along the route from the quarry to the pasture site. At no time in the course of this training did these bees fly the full course from the quarry site to the pasture site, although they very likely flew small segments of it when searching for the displaced feeding table. After the feeding station had reached the pasture site, the bees were captured on emerging from the hive and released at the quarry site. They almost all set a course for the pasture site from the quarry site. If information gained while flying a particular route to and from the feeding station were compartmentalized so that it could be used only when flying that route, then the bees in this experiment should not have been able to set a course for the pasture site from the quarry site. In fact, Menzel, et al. (1990) reasoned that just this sort of result—the one Dyer obtained—would indicate the use of a cognitive map. However, when Menzel, et al. ran an experiment of similar design, they failed to obtain Dyer's result. Their bees flew off on the compass heading they were on when captured rather than setting a new and appropriate course from their release site.

2.6.2. Bees can use the terrain surrounding a release site to set a course that is novel for that site. The last mentioned result—that released bees sometimes fly off on the heading they were on when captured rather than setting a new and appropriate course from their release site—was obtained in several capture-and-release experiments by both Menzel et al. (1990) and Wehner, Bleuler, Nievergelt, and Shah (1990), as well as in some of Dyer's (1991) experiments. One important finding in all three cases was that the courses set by bees released under overcast skies (so that they could not see the sun) did not differ from the courses set by bees that could see the sun when released (Dyer, 1991; Wehner & Menzel, 1990).

These findings have important implications regarding the compartmentalization (or lack thereof) of the positional information associated with different routes. First, they imply that the bees can use what they have learned about the landmarks at the release site to determine which way is north, because in the absence of the sun, bees use the terrain to determine which way is north (as they did in the Dyer experiment portrayed in Figures 7 and 8). Second, they imply that bees can recognize a terrain even when they do not expect to see it, that is, even when it is not on their planned route. Third, they imply that when they have recognized a terrain they can compute from it the landmarks toward which they should fly (that is, the direction in which they should fly) in order to fly the compass course they were on when captured. In these experiments under overcast skies, the course the bees set when released was not the course they were accustomed to fly when viewing the terrain at the release site. Thus, these repeatedly replicated findings are evidence for the non-compartmentalized use of information about terrains remote from the hive.

2.6.3. Bees and wasps find their way home when released in familiar territory faster and more often than when released in unfamiliar territory. Another firmly established finding is that bees captured and released within their foraging territory find their way back to the hive fairly quickly no matter what direction they head off in when released. By contrast, when they are released on territory they have not had a chance to survey, they come back much more slowly or not at all. This sort of finding goes back to Romanes (1885). He put a hive in the basement window of a house with a large flower garden on one side and a lawn leading to a beach on the other. When he released foragers from this hive anywhere in

the garden, they soon appeared back at the hive (often faster than he could run there from the release site), but when he released them on the beach or on the flowerless lawn (where they presumably did not forage), they did not return to the hive, even though many of these release sites were closer to the hive than release sites in the garden.

Becker (1958) , moved a hive into territory unfamiliar to the bees and allowed forager bees with experience at other hive locations to make a single wide-ranging but short survey flight lasting only 3-5 minutes. She captured these bees as they returned from their survey and then released them at various distances south or north of the hive. These bees had never fed at any site in this area, so they had no route-based memories. All 40 bees released 250 meters south or north of the hive came back to the hive—24 of them in less than 5 minutes, and 14 of those in less than 2 minutes. In a control experiment, she brought other hives of bees into the same unfamiliar territory and kept them for a day or two with a flight cage in front of the hive. The cage restricted their flights to within 3 meters of the hive but gave them a view of the sun and of the terrain to the south of the hive (the experience they need to build their ephemeris function so that they can use their sun compass). When these bees—which had never made a survey flight more than 3 meters from the hive—were captured and released 200 meters north or south from the hive, less than 20% made it back to the hive. The average time for those 20% to get there was about an hour.

Janzen (1971) removed twelve orchid bees from their hive in the Costa Rican rain forest and released them at arbitrarily chosen sites 14-23 kilometers away. Seven returned on the day of their release. One, released 20 kilometers from its hive, was back at the hive in 65 minutes.

Almost all of the bees released in the experiments by Dyer (1991), Menzel, et al. (1990). and Wehner, et al., (1990) also soon appeared either at their original destination or back at the hive, regardless of the course they departed on immediately after their release (personal communications from R. Menzel and R. Wehner). The persistent finding that bees released anywhere in familiar territory, whether at a feeding site or not, find their way home fairly rapidly, even if they are initially misoriented, would seem to imply that when they recognize that they are not getting where they intend

to go, they are capable of setting a course for home from wherever they then find themselves. This in turn seems to require an integrated cognitive map.

Menzel (personal communication 11/3/93) has suggested that the bees find their way back to the hive by a spiral search. This suggestion is inconsistent with some well documented aspects of these results. First, it clearly matters whether the bees are released in familiar territory. When they are released in unfamiliar territory, even when it is close to the nest, they often fail to find their way home (Becker, 1958; Romanes, 1885) or take a long time doing so. If they found their hive by a spiral search, it should not matter whether they knew the territory or not, because the assumption is that a spiral search is what the bee does when it does not know which way to go. Second, they often come back in a time that allows decidedly limited opportunity for circling. Bees can fly at 20-30 kilometers an hour, so the orchid bee that made it home from 20 kilometers out in 65 minutes could not have found its hive in that time by a spiral search. It must have been on course for the hive within minutes of its release. Moreover, it was not released at a feeding site, which implies that it could determine where it was even when it was not at a privileged spot. In other words, it implies an integrated map of its foraging range. (In fact, Janzen did these experiments for the purpose of estimating the foraging range of the orchid bee.)

Finally, there is the evidence from the Thorpe experiment with the digger wasps that walk home after displacement (Figure 1). They were able to set a course for their nest from wherever they were released. Sometimes the bees wandered around after their release, presumably in search of recognizable landmarks, but the area covered by these initial, tortuous search trajectories never included the nest area, because it was too far away. The winding search pattern always gave way at some point to a straight march and that march was always aimed close to the nest, even though the nest could not have been visible when that march began.

In short, the often replicated finding that bees and wasps find their way to the nest in a fairly short time when released anywhere in familiar territory--and that bees and wasps unfamiliar with the release territory do not!--is evidence that bees can use what they know about terrain far from their hive to set a course for the hive. They can use this terrain knowledge even when they did not expect

to find themselves in that particular terrain and do not have a terrain-appropriate signal from their dead-reckoning mechanism. On the face of it, this implies that the information they have acquired about the positions of landmarks far from their hive is not compartmentalized. It is integrated into a common system of coordinates, and it available to be used for course setting and holding computations different from those performed when the information was acquired. If the use of information about landmarks is tied to the flying of preplanned routes, then it is not clear how bees and wasps get home so quickly when they find themselves in unexpected (but familiar) places.

3. Concluding Observations

The intuition that the nervous system does not carry out complex computations, at least not in the symbolic mode used for conventional computation, is deeply rooted. The foundation of this intuition is, at least in part, the fact that we do not yet understand how, at the cellular and molecular level, the nervous system computes. We do not know what cellular or molecular process is isomorphic to any of the fundamental elements of computation—adding, subtracting, multiplying, dividing, comparing, storing, and retrieving. We do not know what processes belong to the basic instruction set of the nervous system—the modest number of elementary operations built into the hardware of a computing device to enable it to carry out computations of arbitrary complexity.

The elementary computational operations assumed in many neural net theories are only three: weighting (multiplication), summing, and thresholding. Some of these theories do not have any explicit provision for the storing and retrieving of the values of variables, either because they do not believe that neural nets can properly be described as working with the values of variables (symbols) or because the values are imagined to be represented by the patterns of synaptic weights.

It is often assumed that synaptic transmission is isomorphic to the multiplicative weighting and subsequent summation at the heart of most neural net models. However, it must be appreciated that we do not understand how and over what range synaptic transmission is in fact capable of mimicking either multiplication or linear summation. To say that a process is isomorphic to multiplication is to say that there is a process wherein: (1) one can identify two input

variables and one output variable, and (2) when the output variable is plotted as a function of the two input variables, the resulting surface looks just like the surface you get when you instruct a graphics program to plot $z=x*y$. A similar statement applies to addition. That is, the surface obtained by plotting the output of a process or circuit as a function of the two inputs to the process or circuit should look like the addition surface over some nontrivial portion of this surface (say, two orders of magnitude along both input axes). I know of no case in which this simple test has been successfully applied to any process or circuit in the nervous system. Thus, it must be recognized that we do not know with any certainty how the nervous system implements even the operations that are assumed in neural net models, let alone the fuller set of operations that are taken for granted in any computational/ symbolic model of the processes that mediate behaviors.

The evidence we have that the nervous system really does compute comes from behavior, not from neuroscience. There is nothing surprising in this. The initial evidence for all behaviorally important processes in the nervous system—including nerve trunk (axonal) conduction and synaptic transmission—came from behavioral observations. It is therefore extremely important for the future of neuroscience that we be clear about the extent to which the behavioral evidence does or does not imply computation in the nervous system. If we are persuaded by the behavioral evidence that the nervous system really does compute, then that establishes an agenda for behavioral neuroscience, namely, to discover the processes that enable it to do so. What processes enable it to store the value of a variable? What processes enable it to retrieve that value when it is needed in a computation? How does it multiply together two such values? How does it determine which if either of two values is bigger? And so on. On the other hand, if the nervous system really does not compute, if it only appears to do so, then this agenda makes no sense.

In assessing the strength of the behavioral evidence, it makes sense to focus on computational problems whose solution space is thought to be well understood. Language learning, scene segmentation, and pattern recognition, which are among the problems often focused on in comparing symbolic versus non-symbolic models, are not well understood computationally. Navigation is. Thus, it makes sense to focus on animal navigation. Moreover, some insects make excellent subjects for the study of

navigation, and there is a rich, constantly expanding literature on insect navigation.

I have argued that it is difficult to imagine how any system that did not store and retrieve the values of variables and apply to them all of the operations in the system of arithmetic could do sun-compass course holding, dead reckoning, and the determination of distance from parallax, to name just three of the experimentally well documented processes in insect navigation. A system that stores and retrieves the values of variables and uses those values in the elementary operations that define arithmetic and logic is a symbol-processing system. The behavioral data seem to imply that the nervous system is one such system.

4. Suggestions for Further Reading

- Cartwright, B. A., & Collett, T. S. (1983). Landmark learning in bees: experiments and models. *Journal of Comparative Physiology*, *151*, 521-543. The most explicit description of retinal-image-matching models for piloting in bees.
- Collett, T. S. (1992). Landmark learning and guidance in insects. *Philosophical Transactions of the Royal Society of London, Series B*, *337*, 295-303. A recent review.
- Collett, T.S. & Baron, J. (1994) Biological compasses and the coordinate frame of landmark memories in honeybees. *Nature*, *368*, 137-140. The most recent results on how bees can use the compass bearing of a landmark to locate a food source.
- Gallistel, C. R. (1990). *The organization of learning*. Cambridge, MA: Bradford Books/MIT Press. More on the definition of a representation (Chap 2) and on animal navigation (Chaps. 3-6)
- Harth, E. (1986). Does the brain compute? *Behavioral and Brain Sciences*, *9*, 98-99 An emphatic statement of the position that the brain does not compute.
- McClelland, J.L., Rumelhart, D.E., & Hinton, G.E. (1987) the appeal of parallel distributed processing. In Rumelhart, D. E., & McClelland, J. L. (eds). *Parallel distributed processing*. (pp. 3-44) Cambridge, MA: MIT Press. See pages 31-40 for the "as though" argument—the argument that the brain appears to compute but really does not, or at least does not do so in the symbol-

processing sense of computation. Other chapters in this book that are in part relevant to issues raised in the present chapter are Chap. 4 "PDP Models and General Issues in Cognitive Science" (by Rumelhart & McClelland) and Chap. 6 "Information Processing in Dynamical Systems: Foundations of Harmony Theory," (by P. Smolensky) which compares and contrasts symbolic and subsymbolic modes of computation.

Müller, M., & Wehner, R. (1988). Path integration in desert ants, *Cataglyphis fortis*. *85*, 5287-5290. A formal model of the dead reckoning process that accounts for systematic errors in it.

Séguinot, V., Maurer, R., & Etienne, A. S. (1993). Dead reckoning in a small mammal: The evaluation of distance. *Journal of Comparative Physiology. A.*, *173*, 103-113. Dead reckoning in a mammal.

Srinivasan, M. V. (1992). Distance perception in insects. *Current Directions in Psychological Science*, *1*, 22-26.

Wehner, R., & Srinivasan, M. V. (1981). Searching behavior of desert ants, genus *Cataglyphis* (*Formicidae*, Hymenoptera). *Journal of Comparative Physiology*, *142*, 315-338. A lengthy set of experiments on dead reckoning in the ant

Wehner, R. (1992). Homing in arthropods. In F. Papi (Ed.), *Animal Homing* (pp. 45-144). London: Chapman & Hall. Another recent review, arguing against the hypothesis that insects have integrated maps.

5. Problems and Questions for Further Thought

1. Intuitions about what is computationally complex and what is not play a large role in discussions of cognitive maps. How might we define computational complexity? Is there some definition that makes the computations required to construct an integrated cognitive map (see Gallistel, 1990, pp. 106-110) less complex than the computations described by Cartwright and Collett (1983) in their view-matching model?

2. Give as many examples as you can of functioning isomorphisms. In each example, you must identify two distinct systems, a represented system and a representing system. You must specify what the entities (symbols) are in the representing system and what

the corresponding entities are in the represented system (the things symbolized). Finally, you must specify at least some of the relations in the represented system and the corresponding relations or operations in the representing system. Finally, state at least one if-and-only-if statement linking the two systems. Hint: Whenever anything is measured, it is almost invariably with a view to setting up an isomorphism between some system of physical things and the number system. Second hint: Ever since Descartes, higher mathematics has made extensive use of the capacity of one branch of mathematics to represent another branch.

3. Sketch a graph of the solar ephemeris at a location only a few degrees north of the equator around the time of the summer solstice and another around the time of the winter solstice. Also sketch the solar ephemeris at midsummer at the latitude of Berlin. IN each case only show that portion of the ephemeris during which the sun is visible.

4. Write a computer program that simulates the sun-compass mechanism.

5. If you sail due west at 5 knots (nautical miles per hour) for 6 hours and then southeast at 4 knots for 12 hours, what is your course to get back to where you started and how far is it?

6. Write a computer program to perform dead reckoning, that is, to compute position as a function of time from a signal indicating velocity as a function of time. Do separate programs, one that specifies position in Cartesian (rectangular) coordinates, the other that specifies position in polar coordinates.

7. How many values (numbers) must be stored to specify the position and orientation of a tree line in a rectangular system of coordinates with the nest at the origin?

8. Suppose a rectangular coordinate system defined by the observer's body. The y axis is the line through the middle of the head bisecting the distance between the eyes (the line in the sagittal plane). The x axis is the horizontal line through the middle of the head perpendicular to the y axis. The unit of distance is one stride. Suppose that the equation that represents a tree line in this framework is $y = 5x + 11$. Suppose that the observer is 200 strides from his tent and facing due east. What is the equation of the same line in the coordinate system whose y and x axes are, respectively,

the north-south line and east-west lines through the tent and whose unit of distance is also the stride?

9. How are complicated curves represented in Postscript™? (Postscript is a computer language specialized for representing arbitrary 2-dimensional shapes and doing computations on them? It is used in some illustration programs and to control laser printers. Complicated curves are curves with several inflection points.) What are the advantages of representing curves this way as opposed to by what is called bit mapping them? How many computations are required to reposition a curve in Postscript versus to reposition the bit-mapped version of the same curve?

10. Does the straight line between New York and Los Angeles on a standard (Mercator projection) map represent the shortest route? Does its angle represent the compass heading you should follow if you want to fly a constant-compass-angle course to Los Angeles from New York?

6. References

- Becker, L. (1958). Untersuchungen über das Heimfindevermögen der Bienen. *Zeitschrift für vergleichende Physiologie*, 41, 1-25.
- Cartwright, B. A., & Collett, T. S. (1983). Landmark learning in bees: experiments and models. *Journal of Comparative Physiology*, 151, 521-543.
- Cartwright, B. A., & Collett, T. S. (1987). Landmark maps for honey bees. *Biological Cybernetics*, 57, 85-93.
- Cheng, K., Collett, T. S., Pickhard, A., & Wehner, R. (1987). The use of visual landmarks by honey bees: bees weight landmarks according to their distance from the goal. *Journal of Comparative Physiology*, 161, 469-475.
- Collett, T. S. (1992). Landmark learning and guidance in insects. *Philosophical Transactions of the Royal Society of London, Series B*, 337, 295-303.
- Collett, T. S., & Baron, J. (1994). Biological compasses and the coordinate frame of landmark memories in honeybees. *Nature*, 368, 137-140

- Collett, T. S., Dillmann, E., Giger, A., & Wehner, R. (1992). Visual landmarks and route following in desert ants. *Journal of Comparative Physiology. Series A*, 170, 435-432.
- Collett, T. S., & Land, M. F. (1975). Visual spatial memory in a hoverfly. *Journal of Comparative Physiology*, 100, 59-84.
- Dyer, F. C. (1987). Memory and sun compensation by honey bees. *Journal of Comparative Physiology. Series A*, 160, 621-633.
- Dyer, F. C. (1991). Bees acquire route-based memories but not cognitive maps in a familiar landscape. *Animal Behaviour*, 41, 239-246.
- Dyer, F. C., Berry, N. A., & Richard, A. S. (1993). Honey bee spatial memory: Use of route-based memories after displacement. *Animal Behaviour*, 45, 1028-1030.
- Foster, R. G. (1993). Photoreceptors and circadian systems. *Current Directions in Psychological Science*, 2, 34-39.
- Free, J. B. (1955). The adaptability of bumblebees to a change in the location of their nest. *Journal of Animal Behaviour*, 3, 61-65.
- Frisch, K. v. (1967). *The dance-language and orientation of bees* . Cambridge, MA: Harvard Univ. Press.
- Gallistel, C. R. (1990). *The organization of learning* . Cambridge, MA: Bradford Books/MIT Press.
- Gould, J., & Gould, C. G. (1988). *The honey bee* . New York: Freeman.
- Gould, J. L. (1980). The case for magnetic sensitivity in birds and bees (such as it is). *American Scientist*, 68, 256-267.
- Gould, J. L. (1986). The locale map of honey bees: do insects have cognitive maps? *Science*, 232, 861-863.
- Gould, J. L. (1990). Honey bee cognition. *Cognition*, 37, 83-103.
- Hull, C. L. (1930). Knowledge and purpose as habit mechanisms. *Psychological Review*, 37, 511-525.
- Janzen, D. H. (1971). Euglossine bees as long-distance pollinators of tropical plants. *Science*, 171, 203-205.

- Lehrer, M. (1993). Why do bees turn back and look? *Journal of Comparative Physiology. Series A*, 172, 549-563.
- Manning, A. (1956). Some aspects of the foraging behaviour of bumblebees. *Behaviour*, 9, 164-201.
- Menzel, R., Chittka, L., Eichmüller, S., Geiger, K., Peitsch, D., & Knoll, P. (1990). Dominance of celestial cues over landmarks disproves map-like orientation in honey bees. *Zeitschrift der Naturforschung*, 45c, 723-726.
- Morris, R. G. M. (1981). Spatial localization does not require the presence of local cues. *Learning and Motivation*, 12, 239-260.
- Romanes, G. J. (1885). Homing faculty of hymenoptera. *Nature*, 32, 630.
- Rumelhart, D. E., & McClelland, J. L. (1986). *Parallel distributed processing*. Cambridge, MA:: MIT Press,
- Rumelhart, D. E., & McClelland, J. L. (1987). Learning the past tenses of English verbs: Implicit rules or parallel distributed processing? In B. MacWhinney (Ed.), *Mechanisms of language acquisition* (pp. 195-248). Hillsdale, NJ: Erlbaum Associates.
- Santchi, F. (1913). Comment s'orient les fourmis. *Revue Suisse de Zoologie*, 21, 347-426.
- Schmidt, I., Collett, T. S., Dillier, F.-X., & Wehner, R. (1992). How desert ants cope with enforced detours on their way home. *Journal of Comparative Physiology. Series A*, 171, 285-288.
- Skinner, B. F. (1938). *The behavior of organisms*. New York: Appleton-Century-Crofts.
- Sobel, E. C. (1990). The locust's use of motion parallax to measure distance. *Journal of Comparative Physiology A*, 167, 579-588.
- Sternberg, S. (1966). High speed scanning in human memory. *Science*, 153, 652-654.
- Srinivasan, M. V. (1992). Distance perception in insects. *Current Directions in Psychological Science*, 1, 22-26.

- Thorpe, W. H. (1950). A note on detour behaviour with *Ammophil pubescens* Curt. *Behaviour*, 2, 257-264.
- Tinbergen, N., & Kruyt, W. (1938). Über die Orientierung des Bienenwolfes (*philanthus triangulum* Fabr.). III. Die Bevorzugung bestimmter Wegmarken. *Zeitschrift für vergleichende Physiologie*, 25, 292-334.
- Wehner, R., Bleuler, S., Nievergelt, C., & Shah, D. (1990). Bees navigate by using vectors and routes rather than maps. *Naturwissenschaften*, 77, 470-482.
- Wehner, R., & Menzel, R. (1990). Do insects have cognitive maps? *Annual Review of Neuroscience*, 13, 403-414.
- Wehner, R., & Srinivasan, M. V. (1981). Searching behavior of desert ants, genus *Cataglyphis* (*Formicidae*, Hymenoptera). *Journal of Comparative Physiology*, 142, 315-338.
- Wehner, R., & Wehner, S. (1990). Insect navigation: Use of maps or Ariadne's thread? *Ethology, Ecology and Evolution*, 2, 27-48.
- Zeil, J. (1993a). Orientation flights of solitary wasps (*Cerceris*; Sphecidae; Hymenoptera). I. Description of flight. *Journal of Comparative Physiology. Series A*, 172, 189-205.
- Zeil, J. (1993b). Orientation flights of solitary wasps (*Cerceris*; Sphericidae; Hymenoptera). II. Similarities between orientation and return flights and the use of motion parallax. *Journal of Comparative Physiology. Series A*, 172, 207-222.

7. Answers to problems

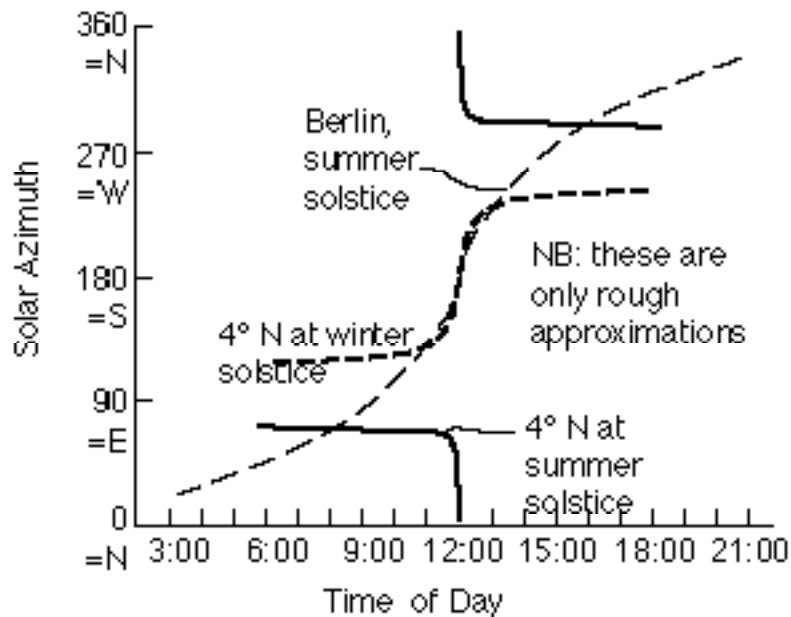
1. Thought question.
2. Examples of isomorphisms

Example 1: Algebra is isomorphic to geometry. Points in geometry correspond to vectors in algebra. Placing two line segments end to end in geometry corresponds to vector addition. A linear equation in algebra corresponds to a line, and so on for the equations for other basic geometric objects,

such as conic sections. Two lines intersect if and only if there is a common solution to their corresponding equations.

Example 2: When numerical weights are generated by a good scale properly used then: The numbers correspond to the inertial masses of the objects weighed. The operation of combining inertial masses corresponds to the arithmetic operation of adding numbers. The inertial mass of two objects combined equals the inertial mass of a third object if and only if the sum of the weights of the two objects equals the weight of the third object.

3. Solar ephemeris functions (next page)



4. Thought question

5. Your course is 357 (7° west of north) and the distance is 34.2 miles.

6. Thought question. For flow diagrams of the computations, see Gallistel (1990, pp. 74-74).

7. Four because it takes two points to define a line and two values to define a line (in the plane). You might think that it could be done with only two, namely, the parameters, a & b , of the linear equation $y = ax + b$; however, this equation cannot represent a line parallel to the y axis. If you thought about the question in terms of representing each successive point in the tree line, then you would answer infinity. Notice that some methods of coding a line are literally incomparably or infinitely better than others. However, these better methods always require computation. That is, they save memory but require computation in order to interpret and use the code that is layed down in memory.

8. $y = 5(x + 200) + 11$ or $5x + 1011$

9. By the parameters of the cubic equations that define successive segments of the curves. These parameters in turn specify the positions of points on the plane (called the control points of the segment). There are 8 parameters (= 4 control points) for each segment of the curve. The advantage is that you save memory and it is much easier to compute useful properties of these curves, such as where they intersect, when are they parallel, what is the coding for the repositioned curve. To compute a repositioning in Postscript, all you have to do is compute the new coordinates of the control points for each segment. To compute the repositioning of a bit mapped image, you have to compute the new coordinates for every point on the line.

10. No. Yes.

8. Key words for Index

symbol-processing system
value
computational complexity
non-symbolic
representation
binding
write operation
isomorphism
cognitive map
compass direction
position
angle
azimuth
ephemeris function

symbol
computation
neural net
sub-symbolic
compositional
homunculus
read operation
function/al/ing isomorphism
integrated map
three-point fix
distance
solar bearing
solar ephemeris
allocentric

egocentric
dead reckoning
apparent size
parallax
scalar
dimensionality
velocity
sun compass
circadian
elevation
zenith

vantage point
translated
triangulation
piloting
vector
displacement
course
clock
time signal
terrain recognition
computation

