

To be published in French with title L'apprentissage de matières distinctes exige des organes distincts . In J. Bricmont & J. Franck (Eds) *Cahier n° ??: Noam Chomsky*. Paris: L'Herne, 2006

Learning Organs

C.R. Gallistel

Department of Psychology
and
Rutgers Center for Cognitive Science
at
Rutgers University, New Brunswick

Harvey (1628) revolutionized physiological thinking when he showed that the heart circulates the blood and that its structure suits it to perform this function. Before Harvey, the modern conception of an organ as something whose particular structure enables it to perform a particular function did not exist. Physiological thinking centered not on organs but on humors. Humors had properties and they had effects. Pathological conditions were thought to arise from an excess or deficiency in one or more of them. But they did not have specific and limited functions. Much less did they have a structure that enabled them to perform a specified function. Organs, by contrast, have specific and distinct functions and a structure that enables them to perform them.

Chomsky (1975) reconceptualized learning in his *Reflections on Language*. His reconceptualization is as radical in its implications for psychology and neuroscience as Harvey's work was for physiology. People generally conceive of learning as mediated by a small number of learning processes, none of them tailored to the learning of a particular kind of material. Examples are habituation, sensitization, and associative learning (see any textbook on learning, for example, Domjan, 1998, particularly pp 17-20, where he explicitly states and defends the general process view, which is taken for granted in associative theories of learning.) These learning processes have properties and they have behavioral effects, but they do not have specific computational functions, nor do they have structures that enable them to perform those computations.

In the associative view of learning, the brain is plastic. It rewires itself to adapt to experience. There are no problem-specific

learning organs, computing representations of different aspects of the world from different aspects of the animal's experience. Chomsky, by contrast, suggested that learning is mediated by distinct learning organs, each with a structure that enables it to learn a particular kind of contingent fact about the world. The non-contingent facts, the universal truths, are not learned; they are implicit in the structure of the learning organs, which is what makes it possible for each such organ to learn the contingent facts proper to it.

Chomsky's suggestion grew out of his recognition that learning was a computational problem—a view that is foreign to the associative conception of learning (Hawkins & Kandel, 1984; Hull, 1952; Skinner, 1950) and to most neurobiological conceptions of learning. In Chomsky's view, the learner must compute from data a representation of the system that generated the data. The example Chomsky had foremost in mind was the learning of a language. He conceived of language learning as the computation from the limited number of often fragmented and often agrammatical sentences a learner hears the grammar of the system that generates all and only the good (well formed) sentences in the language. The computational challenge this poses is so formidable that there is no hope of surmounting it without a task-specific learning organ, a computational organ with a structure tailored to the demands of this particular domain.

Whether this organ resides in a highly localized part of the brain or arises from a language-specific interconnection of diverse data-processing modules in the brain is irrelevant to whether it constitutes a distinct organ or not. Some organs are localized (for example, the kidney) while others ramify everywhere (for example, the circulatory system). The essential feature of an organ is that it has a function distinct from the function of other organs and a structure suited to that function, a structure that makes it possible for it to do its job.

Although Chomsky (1975) had language foremost in mind, he clearly understood that his computational conception of learning implied that other forms of learning must likewise be mediated by problem-specific learning organs. From a computational point of view, the notion of a general purpose learning process (for example,

associative learning), makes no more sense than the notion of a general purpose sensing organ—a bump in the middle of the forehead whose function is to sense things. There is no such bump, because picking up information from different kinds of stimuli—light, sound, chemical, mechanical, and so on—requires organs with structures shaped by the specific properties of the stimuli they process. The structure of an eye—including the neural circuitry in the retina and beyond—reflects in exquisite detail the laws of optics and the exigencies of extracting information about the world from reflected light. The same is true for the ear, where the exigencies of extracting information from emitted sounds dictates the many distinctive features of auditory organs. We see with eyes and hear with ears—rather than sensing through a general purpose sense organ—because sensing requires organs with modality-specific structure.

Chomsky realized that the same must be true for learning. Learning different things about the world from different kinds of experience requires computations tailored both to what is to be learned and to the kind of experience from which it is to be learned. Therefore, there must be task-specific learning organs, with structures tailored both to what they are supposed to extract from experience and to the kind of experience from which they are to extract it. For computational reasons, learning organs may be expected to differ between species of animals, just as do sensory organs. Pit vipers sense infrared radiation, whereas we do not, because they have a sensory organ, which we lack. We learn languages, whereas chimpanzees do not, because we have a language-learning organ, which they lack.

While Chomsky argued that there must be other learning organs, he did not specify what the other domains might be wherein they operated. We are now in a position to do this. In doing so, we make clear the computational nature of learning and why this leads to learning organs, to mechanisms in the brain that carry out a particular kind of computation on a particular kind of data.

The Path-Integration Organ

Learning mechanisms enable us to acquire knowledge from experience. Among the most important kinds of acquired knowledge

is knowledge of where we are. A fundamental computation by means of which animals of many kinds maintain moment-to-moment knowledge of where they are is path integration.

When a sailor sets out from her home port, she keeps a log of her position relative to her home port, based on her estimates of her speed and direction of movement. Her log might look something like this:

<u>Time</u>	Est Speed (prev. hr)	Est <u>Direction</u>	<u>Northing</u> (knts/hr)	<u>Easting</u> (knts/hr)	<u>Change in position</u> (nm)	
					<u>N/S</u>	<u>E/W</u>
5:00	leaving port				<u>N/S</u>	<u>E/W</u>
6:00	4	W	0	-4	0	4.0W
7:00	4	W	0	-4	0	8.0W
8:00	4	W	0	-4	0	12.0W
9:00	4	W	0	-4	0	16.0W
10:00	6	W	0	-6	0	22.0W
11:00	8	NW	5.7	-5.7	5.7N	27.7W
12:00	5	N	5.0	0	10.7N	27.7W
13:00	5	N	5.0	0	15.7N	27.7W
14:00	6	ESE	-3.0	5.2	12.7N	22.5W
15:00	7	ESE	-3.5	6.1	9.2N	16.4W
16:00	5	ESE	-2.5	4.3	6.7N	12.1W
17:00	4	ESE	-2.0	3.5	4.7N	8.6W
18:00	4	ESE	-2.0	3.5	2.7N	5.1W
19:00	4	ESE	-2.0	3.5	0.7N	1.7W

At the conclusion of each hour she has recorded her estimated average speed during the past hour and her estimated direction of movement. She has then broken these estimates down into an estimate of how far she has progressed in a north/south direction and how far she has progressed in an east/west direction. For the first five hours, she sailed due west, so there was no change in her north south position. For the first 4 hours her speed was 4 knots, putting her 16 nautical miles west of her port at 9:00. In the next hour, her speed picks up. She covers 6 miles in that hour, putting her 22 miles west. A change in the wind at 10:00 forces her to change her course to northwest. Now, for every mile that she goes in the northwest direction, she moves 0.707 miles north and 0.707 miles west. Her speed picks up to 8, so after an hour of this, she is

$0.707 \times 8 = 5.7$ miles north of her home port and a further 5.7 miles west (making a total of 27.7 miles west). Now the wind forces her to sail due north for 10 miles, at which point she decides to head for home. Looking at her log, she sees that she is 15.7 miles north of her port and 27.7 miles west. To head for home she has to sail more east than south, so she sets a course to the east southeast. For every mile she covers on this course, she progresses half a mile to the south and .87 miles to the east. After 6 hours on this course at varying speeds, she calculates that she is about 0.7 miles north of her home port and 1.7 miles west.

The computation laid out in this table is the path integration computation. It is a computation performed on pairs of numbers. The pairs of numbers specifying for each hour the northing and easting (how fast and in what direction she has moved along a north/south axis and how fast and in what direction she has moved along an east/west axis) are her (estimated) velocities. It takes trigonometry to compute this number pair from her speed and course, which are the pair with which the computation begins. The number pairs in successive rows of the last two columns, which indicate her net change in positions, are obtained simply by adding up the hour-by-hour movements along each of the two directional axes.

It is easy to build physical devices that implement this computation. They do not, of course, work with numbers qua marks on paper; they work with physical quantities that could themselves be represented by such numbers, such as, for example, the bit patterns in the memory bank of a digital computer or the voltages in an analog computer. In the nervous system, such a computational mechanism works with neural activity proportional to the animal's speed and neural activity indicative of its direction. These neural signals take the place of the marks on paper (numbers) with which the human navigator works in computing her position. It is also easy to make the output of such a computation control the direction of a rocket or a robot, so no homunculus in the brain has to "see" and "interpret" these computations in order for their results to be manifest in the control of action. That the brain's behaviorally relevant activity must be understood in computational terms is the computational theory of mind, which is at the core of contemporary

cognitive science. Chomsky's work did much to establish this conceptual framework, the framework in which the brain's activity is conceived of in computational terms.

The brains of most animals contain neural machinery that implements the computation shown in the above table. That machinery--wherever it is found, however dispersed it may or may not be, and however it may be realized--constitutes a learning organ. It has a particular structure. It takes sensory signals indicating speed and sensory signals indicating direction and combines them to get signals indicating velocity, then it adds up those velocities to get the net change in position. This structure has implicit in it a non-contingent purely mathematical truth about the world, namely, that position is the integral of velocity. That truth determines the structure of this learning organ in the same way that the truths of optics determine the structure of an eye. The function of this organ is to compute the animal's location from the sensory inputs it has generated in the process of moving to that location. Its structure suits it to perform that function and only that function. No one would suppose that an organ with this structure would be of any use in the learning of a language. Applying this learning organ to that learning problem would be like trying to hear with an eye or breath with the liver.

Parameter Setting in the Organ that Learns the Solar Ephemeris

The path integration computation requires a signal that indicates the compass direction in which the animal is progressing. Compass direction is direction relative to the axis of the earth's rotation, the north-south axis. Many animals use the sun's position for this purpose, which is remarkable because the sun does not have a fixed compass direction. Its compass direction varies continuously during the day. Worse yet, the manner in which it does so depends on the season of the year and how far the observer is to the north or south of the equator. In short, the compass direction of the sun at a given time of day is a contingent fact about the world. If, however, the observer can overcome this problem, then the sun has advantages as a directional referent. It can be seen from almost any vantage point, and, most importantly, it is so far away that movements of an extent realizable in one day have very little effect on its direction.

This is not true of earthly landmarks. The landmarks one can see are rarely farther away than the distance one can traverse in a small fraction of a day, which means that they change their compass direction as one moves. That makes them poor indicators of compass direction.

Human navigators also routinely use the sun as a compass referent; they have understood how to do so explicitly for centuries, and implicitly, probably for eons. The trick is to know the time of day and to learn the compass direction of the sun as a function of the time of day, that is, to learn the solar ephemeris.

Animals, including humans, have a built in clock, which solves the first part of the problem, knowing the time of day. The circadian clock mechanism is another example of the innumerable ways in which the enduring structure of the world they inhabit is reflected in the innate structure and functioning of the animal mechanism.

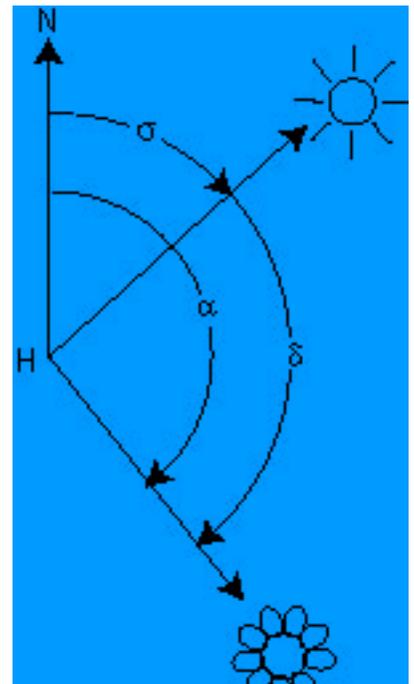
The second part of the problem is to know the compass direction of the sun for any given time of day. Bees and ants and girds (and probably many other animals as well) have an organ that learns this. Remarkably, bees, at least, learn the complete function from only a few observations, that is, without having seen the sun at most moments of the day. The learning of the solar ephemeris in bees illustrates the force and broad application of Chomsky's famous poverty-of-the-stimulus argument. Chomsky pointed out that the conclusions that humans draw about the grammatical and phonological structure of the languages they learn go far beyond what is justified by the limited samples from which they learn. Similarly, as we shall see, the conclusions that bees draw from limited observations of the sun's trajectory go beyond what is justified by its observations.

The structure of the organ that learns the solar ephemeris also illustrates a key idea in contemporary theories of language learning—the idea of parameter setting. As in Chomsky's thinking, the concept of a built in function with settable parameters explains the ability of a learning organ to know more than it has observed, that is, to overcome the poverty of the stimulus.

The key experiment here was done by Dyer and Dickinson (1994), following earlier related work of a similar nature by Lindauer (1957; 1959). Dyer and Dickinson used the famous dance of the returning bee forager (Frisch, 1967). The just returned forager dances on the vertical surface of the honeycomb, inside the hive, out of sight of the sun. The dance is in the form of a figure 8. The bee circles first one way, then the other. In the middle stretch, where the loops of the 8 converge, the dancing bee waggles as it runs (flicks its abdomen from side to side). The direction of the waggle run relative to the vertical indicates the direction of the food source relative to the sun. If the bee runs straight up while wagging, it indicates that one must fly toward the sun; if it runs horizontally to the right, it indicates that one must fly with the sun on one's left; if it runs straight down, it indicates that one must fly away from the sun; and so on. The number of waggles during each waggle run indicates the distance: the more waggles, the farther the food.

The compass direction of the source from the hive (α in Figure 1), equals the compass direction of the sun (σ in Figure 1) plus the solar bearing of the source (δ in Figure 1). Thus, when the experimenter, who observes the dance through a window on the hive constructed for that purpose, observes the direction of the waggle run, he can infer from it and from the compass direction of the source, the direction in which the dancer believes the sun to be.

Figure 1. *Diagram of the angular (directional) relations. H = hive; N = north. σ = the compass direction of the sun. This direction is locally constant; it is the same regardless of where the bee is within its foraging range. α = the compass direction of the flower from the hive. δ = the solar bearing of the flower from the hive, which is the direction that the bee must fly relative to the sun in order to get from the hive to the flower.*



In Dyer and Dickinson's experiment, they raised bees in an incubator without a view of the sun. Then, they allowed them to forage around their artificial hive, but only in the late afternoon, when the sun was sinking in the west. The bees learned to fly back and forth between the hive and a feeding station, which was located some tens of meters away to (for concrete example) the west of the hive. When these foragers arrived back in the hive, they did a dance in which the waggle run was directed more or less straight up, indicating that, to get to the food, one should fly toward the sun.

Dyer and Dickinson allowed their bees to forage only in the late afternoon, until a morning came when the sun was completely hidden by a heavy overcast. Then, for the first time, Dyer and Dickinson allowed their bees to forage in the morning.

One might suppose that bees do not do their dance when overcast prevents their seeing the sun. How can you tell another bee how to fly relative to the sun if neither you nor the other bee can see the sun? But bees do dance when they cannot see the sun; other bees do attend to the dance; and they do fly in the direction indicated by the dance, even when they cannot see the sun as they fly? How can this be?

The key thing to realize is that once you have formed a compass-oriented map of the terrain within which you are navigating, you do not need to see the sun in order to set a course by the sun, provided you know the time of day and the sun's solar ephemeris, its compass direction at different times of day. If I tell you to fly away from the sun at 9:00 in the morning, a time at which you know the sun lies more or less due east, then I am in effect telling you to fly west. If you have a map of the terrain and you know the direction at which the sun must be relative to that terrain at that time of day, then you can fly in the indicated direction by reference to appropriate terrain features. You fly toward the landmarks that lie to the west of your hive on your map, that is landmarks that lie in the direction opposite the direction in which you know the sun must currently be.

The key to this scheme is the mechanism that learns the solar-ephemeris, which, to repeat, is the compass direction of the sun as a function of the time of day. Learning the solar ephemeris—learning

where the sun is at different times of day in relation to the terrain around the hive—is what makes the bee's map of that terrain compass oriented.

Returning now to the Dyer and Dickinson experiment: They allowed the bees to forage in the morning only when the sky was overcast, so that they could not see the sun. These bees, who had only seen the sun above the terrain lying to the west of the hive, flew to the food, and when they returned, they did the waggle dance, telling the other bees how to fly relative to the (unseen) sun. The astonishing thing is that the waggle run was now directed straight down rather than straight up. They were telling the other bees to fly away from the sun. This implies that they believed that the sun in the morning was positioned above the terrain in the direction opposite the food, even though, in their experience, it had always been positioned above the terrain in the direction of the food. Their experience of the sun, which was confined to the late afternoon, gave no grounds for their belief about where it is in the early morning (the poverty of the stimulus). Nonetheless, they had an approximately correct belief. They believed it to lie over the terrain to the east of the hive. For those who are uncomfortable with the attribution of beliefs to bees, this above can be reworded in terms of the compass direction of the sun specified by nerve signals in the brains of the dancers. These signals specified an eastward solar direction in the morning, even though the experienced sun had always been westward.

These fascinating results suggest the following model for the mechanism that learns the solar ephemeris (cf Dickinson & Dyer, 1996): Built into the nervous system of the bee is a genetically specified neural circuit that takes as its input signals from the bee's circadian clock and gives as its output a signal specifying the compass direction of the sun. If we imagine this signal to be the firing rates of two pairs of compass neurons in the brain, a N/S pair and an E/W pair, and if we imagine that different firing rates correspond to different compass directions, then the function in question is the one that determines what the firing rates of the compass neurons at different phases of the bee's circadian clock.

The model further assumes that the general form of the function relating the output of the bee's circadian clock to the firing

rates of its compass neurons is also genetically specified. It has the property that the firing rates for phases of the circadian clock 12 hours apart indicate opposing compass directions. One can imagine a genetically specified dynamic biochemical process in the brain, closely analogous to the one that implements the bee's circadian clock. This mechanism leaves two things to be specified by the bee's experience: 1) the terrain views associated with a given set of firing rates (this anchors the compass signal to the local terrain); 2) the kinetics of the change in the firing rates, how they change as the circadian clock mechanisms runs through its cycle. The terrain views are analogous to the words in a lexicon. The kinetics of the ephemeris function are analogous to the grammar of a language. Observing the sun at a few different times of day suffices to determine the values of the free parameters in the innately specified ephemeris function. Fixing the values of those parameters adapts the universal function to the locally appropriate variant.

The interest of this model for present purposes is that it is an example of learning by parameter setting. The genes specify the general form that the relation between compass direction and time of day must take; in mathematical terms, they specify the function that is to be fitted to the data from experience. What is left for experience to specify are the values of a few parameters in the function (that is, the levels of a few physical variables in the mechanism that specifies where the sun is at different stages of the circadian cycle). What this learning mechanism does is allow the bee to get quickly to an approximately correct representation of the local solar ephemeris. It solves the poverty of the stimulus problem. It enables the bee to estimate where the sun will be at times when it has never seen the sun, just as knowing the grammar of French enables one to speak and understand French sentences one has neither heard nor spoken before.

Modern theories of language learning in the Chomsky tradition are similar in spirit. The basic, high-level form of the grammar of any conceivable human language is assumed to be given by genetically specified language learning machinery. This genetically specified general structure for any human language is what linguists call the universal grammar. Much research in linguistics is devoted to trying to specify what that general form is.

The general form has in it a number of parameters whose value must be specified by experience. These parameters are generally thought to be binary parameters (the language must either do it this way, or it must do it that way). Learning the language reduces to learning the parameter settings. When the child has got all the parameters set, it knows the language. It can produce and understand sentences it has never heard for the same reason that the bee can judge where the sun is at times of day when the bee has never seen the sun. As in the bee, the built in structure of the learning organ, which is what makes learning possible, has as a consequence that the inferences the learner has drawn from limited experience go beyond what is justified by that experience.

This is a radically different conception of the learning process from the one that has dominated thinking in Western empiricist philosophy and psychology and neuroscience for centuries. It remains deeply controversial. Indeed, neuroscientists have not even begun to contemplate its implications. The idea that there may be no such thing as associative learning is an unthinkable thought for most neuroscientists. Nonetheless, this conception is now very much, so to speak, on the scientific table. Chomsky did more than anyone else to put it there.

References

- Chomsky, N. (1975). *Reflections on language*. New York: Pantheon.
- Dickinson, J., & Dyer, F. (1996). How insects learn about the sun's course: Alternative modeling approaches. In S. Wilson & e. al. (Eds.), *From animals to animats. Vol. 4* (pp.??-??). Cambridge, MA: MIT Press.
- Domjan, M. (1998). *The principles of learning and behavior*. Pacific Grove: Brooks/Cole.
- Dyer, F. C., & Dickinson, J. A. (1994). Development of sun compensation by honeybees: How partially experienced bees estimate the sun's course. *Proceedings of the National Academy of Sciences, USA, 91*, 4471-4474.
- Frisch, K. v. (1967). *The dance-language and orientation of bees*. Cambridge, MA: Harvard Univ. Press.

- Harvey, W. (1628). *Exercitatio anatomica de motis cordis et sanguinis in animalibus* (R. W. w. r. b. A. Bowie, Trans. Harvard Classics ed.). Frankfurt am Maine.
- Hawkins, R. D., & Kandel, E. R. (1984). Is there a cell-biological alphabet for simple forms of learning? *Psych. Rev.*, *91*, 375-391.
- Hull, C. L. (1952). *A behavior system*. New Haven, CT:: Yale University Press.
- Lindauer, M. (1957). Sonnenorientierung der Bienen unter der Aequatorsonne und zur Nachtzeit. *Naturwissenschaften*, *44*, 1-6.
- Lindauer, M. (1959). Angeborene und erlernte Komponenten in der Sonnenorientierung der Bienen. *Zeitschrift für vergleichende Physiologie*, *42*, 43-63.
- Skinner, B. F. (1950). Are theories of learning necessary? *Psychol. Rev.*, *57*, 193-216.