Received: 25 May 2019 | Accepted: 17 July 2019

DOI: 10.1111/mila.12289

CRITICAL NOTICE

WILEY

Where meanings arise and how: Building on Shannon's foundations

Charles R. Gallistel

Rutgers University, New Brunswick, NJ

Correspondence
Charles R. Gallistel, 252 7th Avenue 10D,
New York, NY 10001.
Email: galliste@ruccs.rutgers.edu

Information theory provides a quantitative conceptual framework for understanding the flow of information from the world into and through brains. It focuses our attention on the sets of possible messages a brain's anatomy and physiology enable it to receive. The meanings of the messages arise from the inferences licensed by the brain's processing of them. Different meanings arise at different levels because different representations of the input license different inferences.

KEYWORDS

communication, information, levels of meaning, licensed inferences, set of possible messages

1 | INTRODUCTION

Shea (2018) presents and defends a version of what has emerged over the past few decades as the standard account of representation in the cognitive sciences, an account that I endorse: A representation is a homomorphic mapping between a represented system and a representing system, whose function is to enable the representing system to act effectively on and within the represented system.

A representing system has two essential components: (a) machinery for the rule-governed manipulation of symbols and (b) machinery for establishing the external referents of some of the symbols. The bit patterns in the memory of a computer are examples of symbols. The logical and arithmetic operations hard-wired into the processing unit(s) of a computer are an example of machinery for the rule-governed manipulation of symbols. The transducers that map from quantities in the world to bit patterns in computer memory are examples of reference-establishing machinery. The actions directed by a process-control computer in and on its environment depend on both the mapping from aspects of that environment to the symbols in its memory and on its manipulation of those symbols by its symbol-processing machinery. In order for those actions to have appropriate effects, the symbolic manipulations must be homomorphic to the relevant aspects of the environment (have the same mathematical form). The relevant aspects are determined by the mappings implemented by the transducers and multiple stages of posttransduction signal processing. The manipulations performed on the symbols must make sense given those referents.

Before we were surrounded by process-control computers, it was sometimes claimed that computers did not really represent anything; our impression that they did so depended on *our* assigning *meaning* to the bit patterns (Searle, 1980). The emergence of autonomous process-control computers going about what used to be our business without our intervention should put this canard to rest. I do not think we want to insist that a driverless car ferrying us around San Francisco does not really represent the locations of the pedestrians and their likely movements. It avoids collisions with them under a wide variety of challenging circumstances even though we do not mediate between the computer, where the representation resides, and the steering, brake, and accelerator, where its effects are observed.

The concept of a representation along the lines just sketched is central to any understanding of how most modern machinery works—down to the level of the toaster. That fact seems to me to naturalize representation. It is no more inescapably mentalistic than energy or mass. The concept is also central to how many contemporary cognitive scientists and neuroscientists think about how brains work. To claim that the brain of the frog does not represent the location of the object at which it is about to strike seems to me perverse because I do not believe that it is possible to construct a machine that does what the frog does without building in a component whose function is to represent the location of the target. In the frog, we even know where that component is; it is in the optic tectum. We even have a pretty good story about how the location of the target is computed and how it is represented by the firing of neurons in the optic tectum (Gallistel, 1999; Sparks, 1988).

2 | INFORMATION

Shea mentions Shannon's theory of information. However, like many of his philosophical predecessors (e.g., Dennett, 1986; Dretske, 1981), he makes no use of Shannon's conceptual framework. He does not use it because he is centrally concerned with semantics—what information-carrying symbols mean, what they are about, what their content is—and Shannon (1948) begins with the observation that the meanings of messages are irrelevant to understanding their communication. The theory of information that emerges from Shannon's penetrating analysis, therefore, says nothing about the meanings of messages. However, in what follows, I argue that an understanding of what messages mean must flow from an understanding of the implications of Shannon's analysis of communication.

2.1 | Mutual information instead of correlational information

Shea's treatment of information has two aspects: correlational information and unmediated explanatory information. The latter is correlational information that plays a directly causal role in the generation of a behavior. Correlational information is said to give rise to content. Shea takes a stab at formalizing correlational information in terms of the difference between a conditional probability and an unconditional probability, but he does not build on this. If he tried to build on it, he would discover hard-to-surmount conceptual and technical difficulties in applying a correlation, mathematically defined, to events that unfold asynchronously in real time.

The first difficulty with correlation is that Pearson's r cannot be computed for dichotomous variables. One has to compute Phi, Pearson's coefficient of correlation. Computing Phi requires constructing a contingency table. That brings in the second conceptual difficulty, which is that the numbers which go into a contingency table—hence the obtained values of Phi—depend on an arbitrary theorist-imposed discretization of the flow of time (Gallistel, Craig & Shahan, 2014, 2019; Gibbon, Berryman & Thompson, 1974; Hallam, Grahame & Miller, 1992; Hammond, 1980). Thus, correlational information is not well defined.

Shea's account would be strengthened by replacing correlational information with Shannon's mutual information. Mutual information is the information shared between two variables. It tells us how much information about the value of one variable may be gained from knowing the value of the other. Measuring/computing mutual information does not require discretizing the flow of time, and it does not assume a form for the relation between the variables (Cover & Thomas, 1991; Kinney & Atwal, 2014). It is an intuitive, generally applicable measure of how much can be learned about a remote (distal) variable by processing a more proximal variable that carries information about it (for a detailed treatment of its computation and use in understanding neural communication within a brain, see Rieke, Warland, de Ruyter van Steveninck & Bialek, 1997).

2.2 | Source, signal, and receiver

More generally, we can build a clearer, more intuitive, quantitative, and scientifically useful theory if we build on Shannon's foundations. Shannon's analysis identifies three key quantities in the flow of information from the world into the brain: *the information available* from a source, *the mutual information* between that source and information-bearing signals that derive from the source and impinge on sensory receptors, and *the communicated information*—that is, the information extracted from those signals by the receiver. The information available from a source sets the upper limit on how much it is possible for a receiver to learn about that source. The mutual information between source and signal sets the upper limit on how much the receiver can learn about that source from that signal. The communicated information is how much information the receiver in fact extracts from the signal.

When applied to the flow of information from the world into brains, the brain is Shannon's receiver. On closer consideration, a brain is a sequence of receivers and sources because the proximal stimulus—the world variable (e.g., light pattern, sound pattern) that acts on sensory receptors—generates signals in sensory receptors that then pass through a sequence of interpretive stages of neural processing. Each stage is a receiver of the signals from the preceding stage and the source of the signals it sends to succeeding stages.

2.3 | Sets of possible messages

The most important aspect of Shannon's analysis for our understanding of how meanings arise in brains is the starting point of his analysis: the assumption that there is a set of possible messages. This assumption is foundational: In the absence of a set of possible messages, information is undefined because information is the average surprisal of a message in that set, where the surprisal of the ith message is $\log(1/p_i)$. The less probable a message is, the more surprising it is because $\log(1/p_i)$ goes to infinity as p_i goes to 0. The average surprisal is the entropy of the distribution of probabilities on that set.

When we apply Shannon's conceptual framework to the flow of information from the world into the brain, the first and most basic question is: What determines the set of possible messages? The world is indifferent to the informational requirements of a brain. It is the brain that needs information from the world to program effective actions in and on it. Thus, it is the brain that stipulates what the possible messages are; they are all and only the messages it is equipped to receive. What I find missing from Shea's analysis is the insight this simple fact provides regarding the fundamental role of the brain in creating meaning. A brain creates meaning when it sets up an apparatus for receiving and processing information about an aspect of the experienced world.

The meaning of a message arises from: (a) which message it maps to within a brain-stipulated set of possible messages, and (b) the computational operations that the brain has stipulated are applicable to the messages in that set. Together, these determine the (possibly infinite) set of inferences that the brain may draw from a message. In stipulating the sets of possible messages, the brain sets up the loci at which the meanings of messages arise. The genetic program underlying brain development does the stipulations necessary to get the brain's many representational systems up and running.

In this story, the set of potential inferences that the brain may draw from a message *is* the meaning of the message. These inferences need not be truth preserving. However, if the inferences actually drawn are often or always not truth preserving, the representation will be a very poor one. It will more often misrepresent than represent the reality within and on which the brain acts. A common consequence of systematic misrepresentation in a brain is that it ceases to rely on the malfunctioning representation (Gallistel, 2012).

3 | EXAMPLE 1: THE SEQUENCE OF MEANINGS IN THE INTERPRETATION OF THE DANCE

At several points, Shea discusses the dance of the returned honeybee forager. This fascinating and much studied phenomenon provides an excellent example for illustrating both mutual information and for my account of how a sequence of different meanings arise as signals pass to successive interpretive stages in a brain—in this case, a small brain (the honeybee brain has about 1,000,000 neurons, which is, by some estimates, only 50% more than the number of neurons in one voxel of a human fMRI).

A forager bee returning from a good source of nectar runs a repeated figure-eight pattern on the vertical surface inside the hive—in the dark, out of sight of the sun. The angle of the eight's tilt away from the vertical approximately equals the *current solar bearing* of the source from the hive. In the segment of the eight where the two circles superpose, the dancing bee makes the waggle run. The angle of this run relative to vertical is the tilt of the eight. The duration of the waggle run is approximately proportional to the *range* of the source from the hive. Thus, this signal comes in polar form: The parameters of the dance specify the *bearing* and *range* (or in less technical language, the *direction* and *distance*) of the nectar source in a *solar frame of reference* (i.e., the bearings are specified relative to the sun's *azimuth*). The mutual information between the dance parameters and the current solar bearing and range of the source depends on the scatter plots relating the angle and duration of the dance to the current solar bearing and range of the source. By analyzing scatter plots of the parameters of repeated waggle runs (gravitational angle and duration) versus the known bearing and range of the source, Schürch and Ratnieks (2015) find that there are approximately three bits of mutual information between the angle of the waggle run and the solar bearing of the source and approximately four bits between the duration of the waggle run and the range of the source.

The relevant *neura*l signals, the signals generated in the sensory neurons of dance followers by *proximal stimuli* during a waggle run, are not a simple reflection of the angle of the waggle runs observed. Judging by recent evidence from Drosophila (Kamikouchi et al., 2009; Sun et al., 2009), the sensing of the direction of the gravitational pull depends on mechanoreceptor (strain sensitive) organs, which are specialized for that purpose and located in the antennae. Antennae are highly mobile. The direction of gravitational pull on the sensing organ depends on how the antenna is oriented, not on how the bee itself is oriented. The two antennae generally have different orientations at any given moment, neither of which corresponds to the gravitational angle of the dance attendant itself. Thus, extracting the bee's orientation to gravity from the signals generated in the organ for sensing the direction of gravitational pull is already a nontrivial computational problem. There is probably an early stage of sensory processing devoted to solving just that problem.

Sensing the gravitational direction plays a role in a great many different aspects of insect behavior, so the meaning of the representation of the relevant sensory input at the initial stage of sensory processing has nothing specifically to do with the angle of the waggle run relative to hive vertical, let alone with the solar bearing of the nectar source. That is why it is wrong to assign high-level meanings to low-level sensory processing stages (see Shea, Section 6.2 on *Indeterminacy*).

The relation of proximal stimuli and the sensory signals they produce to a property of the distal world is almost always complex. A complex relation between a signal and its import is characteristic of modern communication systems. This complexity is even greater when it comes to the one-sided communication between the inanimate world and a brain. That is why "Computation is communication limited and communication is computation limited" (Cover & Thomas, 1991, p. 4). Truer words were never written. The relevance of this fact to a philosophical understanding of representation in brains is that, when we delve into what the brain actually has to do in order to get information in a useable form, we come to see how meaning is created by the computations required to put the information in that form.

Because the flow of information from the world into the brain requires extensive and complex computation simply to describe in a useful way the signals created by proximal stimuli in sensory neurons, the meanings of the messages in early stages of sensory processing are generally to be understood in terms of technical and abstruse characterizations of proximal stimulus configurations as, for example, when neurons in V1 are understood as spatiotemporal Gabor filters. The utility of the low-level descriptions of proximal stimuli computed at the initial stages of sensory processing resides in the manner in which they facilitate the computation of higher-level, more intuitively intelligible meanings, the meanings that refer to behaviorally relevant properties of the distal stimuli.

I assume that it is these higher-level messages that Shea has in mind when he speaks of unmediated explanatory information. However, because he treats correlational information as a brain-independent fact about the relation between source variables and signal variables and because he treats unmediated explanatory information as correlational information, his analysis tends to make one think that the world forces its messages on a passive brain. This does not put the focus where it belongs, which is on the brain itself.

3.1 | The limits on the amount of information available

One does not, for example, take away from Shea's analysis the fact that it is the brain, not the world, that sets an upper limit on the amount of information that the brain can obtain from a source. The amount of information it can, in principle, obtain about a quantitative variable, such as the gravitational angle of a waggle run, is determined by the precision and dynamic range of its representation because a probability distribution is defined only over the number of *distinguishable* possibilities in the support for the distribution.

The precision with which a brain represents a quantity is the inverse of the Weber fraction, denoted w, which is the standard deviation of the variation in the representation divided by the mean. On the assumption that the brain scales the dynamic range of its measurement apparatus to whatever is required in a given context and the further assumption that the mean is near the middle of that dynamic range, the number of distinguishable values in the brain's representation of a quantity is $2/w$. This number sets the upper limit on how much information the brain can have about the value of a sensory variable because the maximum possible entropy for a distribution over n possibilities is $\log n$.

Weber fractions for abstract quantities such as direction, distance, and duration generally fall between .06 and .25, which implies somewhere between 32 and 8 distinguishable possibilities. This range of distinguishable possibilities translates into between five and three bits of information, respectively, when all possibilities are equally probable. Thus, if the brain sets up a memory buffer for storing the angle of an observed waggle run and gives that buffer a three-bit capacity, the maximum amount of information it can retain about the bearing of a nectar source from observing a waggle run is three bits. If it sets up another buffer for storing the duration of a waggle run and gives that buffer a four-bit capacity, then the maximum amount of information it can retain about the range of a source from observing a waggle run is four bits.

Buffers with those capacities make sense given the results of Schürch and Ratnieks (2015) computations of the mutual information between the bearing and range of nectar sources and the angles and durations of the waggle runs of foragers returned from those sources. There is no point in implicitly assuming that you have more information about a behavior-relevant quantity than your method of measuring it permits you to have. Indeed, it is dangerous to do so (Gallistel, 2017). Well-designed systems avoid doing so, and brains are well-designed systems.

3.2 | Levels of meaning

I take it that, what Shea calls, "unmediated explanatory information" is a message whose meaning determines some aspect of an action. How to determine what those meanings are is an empirical question. In the case of the bee dance, it is one for which we have relevant experimental results. They suggest that Shea underestimates the level of abstraction at which observing a dance influences the subsequent behavior of a bee forager. I draw this conclusion from the size and scope of the inferences that foragers have been shown to draw from the dances they have observed.

Superficially, the dance indicates the current solar bearing and range of the source. Shea interprets these indications as flying instructions: Fly with the sun 20° on your left. von Frisch (1967) already understood that this was not the meaning of the dance message. If it were, it would be of little use because the sun's azimuth changes by more than 40° during an hour at midday in summer in Germany. Thus, if interpreted as flying instructions, the dance would lead other foragers astray in short order.

However, one of the first things a new forager learns is the solar ephemeris, the position of the sun relative to the local terrain, as a function of the time of day, as indicated on its circadian clock (Gallistel, 1998). The stored solar ephemeris, together with the time-of-day signal from the circadian clock, enables the bee brain to add the angle of the waggle run to the sun's azimuth at the time the dance is observed. This simple computation yields the *compass bearing* of the source from the hive, its bearing in a frame of reference established by the local landscape, not by the sun. The change in meaning—the increase in the level of abstraction at which the meaning is appropriately described—broadens the scope of the inferences that may be drawn from it. The solar bearing is only useful for a few minutes after it is communicated; the compass bearing is useful all day long. Given it and the current azimuth of the sun (from the stored ephemeris function), the brain may compute how to fly relative to the sun in order to arrive at the source at any time of day.

A message that conveys a compass bearing and a range may be interpreted/described in two different ways: (a) as the rhumb line from the hive to the source—a flying instruction, a description of a location in terms of what you have to do to get there from the hive, and (b) as specifying the location of the source relative to the hive—its map coordinates; hence, its relation to everything else on the map. These are examples of two different descriptions with the same referent but different meanings. Both descriptions use two-element vectors—strings composed of two numbers—but the numbers in these strings have different referents. The rhumb line numbers describe a destination in polar coordinates; the first number in the string refers to the compass direction, and the second refers to the length of the line segment connecting the start to the destination. These numbers are directly useful (unmediatedly explanatory) only for flight-control actions originating at the hive. The map coordinates describe the source location in Cartesian coordinates; the first number refers to the directed (signed) distance of the source location to the north or south of the hive; the second refers to the directed distance to the east or west. The Cartesian description is directly useful for dead reckoning, map construction, and course planning because Cartesian coordinates combine additively. Given that the range and bearing of Oxford from London is 83 km 289°, and the range and bearing of Cambridge from Oxford is 93 km 57°, you have to first convert to map coordinates to get the north–south and east–west differences between Cambridge and London—and then convert these differences to polar form to get the bearing and range of London from Cambridge. Therefore, for preserving locational information on a map, the Cartesian description is more useful. It is a higher-level description from which a lower-level description may be computed whenever necessary. It directly licenses a broader range of inferences.

A recent experiment by Randolf Menzel's group showed that bees extract map coordinates from the dances they witness and use them for route planning (Menzel et al., 2011). They trained two platoons of individually labeled bees to two different feeding locations (Figure 1). Through the analysis of an in-hive video tape, they knew that the platoon trained to the FT30° site in Figure 1 had repeatedly observed the dances of the bees in the platoon trained to the FD30° site.

The day after these dance observations, the first platoon found no nectar when they arrived at the FT30° site. Figure 1 shows the flight made by two of these disappointed bees on leaving FT30°. One set a course more or less straight to the FD30° location, which it knew about only by hearsay (red trajectory). The other started back to the hive, but after it had flown a couple of hundred meters on that course, it changed its mind and set a course from where it then was directly to the FD30° location (cyan trajectory). This latter flight is even more important in its implications than the first. It shows that bees can set a course to a known location of motivational interest from any arbitrary location on their map. They can do this even when the coordinates of the destination are known only by hearsay and the location from which the course is set is where the bee happened to find itself when it changed its mind about where it wanted to

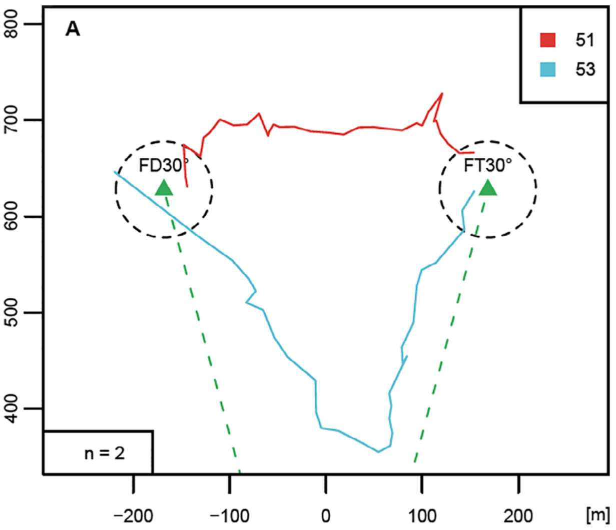

FIGURE 1 Radar tracks of the flights of two disappointed bees who arrived at their habitual feeding source (FT30°) after observing the dances the previous day of bees foraging at another site (FD30°) only to discover that their habitual site had no nectar. The dashed green lines are the rhumb lines from the hive to the two sites (From Menzel et al., 2011, used by permission of the author and publisher). See the online version for a colour figure

go. The number of different possible locations within the several square kilometers over which a bee forages is, for practical purposes, infinite and so, therefore, is the set of courses it can set. Thus, the number of different inferences that may be drawn from a single dance message, when coded as map coordinates, is indefinitely large.

Results from two other experiments expand still further the range of inferences to which a message conveyed by the dance may contribute (Gould & Gould, 1995; Tautz et al., 2004; but see Wray, Klein, Mattila & Seeley, 2008). These results, if confirmed by further work, imply that, before deciding to follow the flight instructions provided by a dance, the receiver of the information consults its map to see whether there is a reasonable prior probability of a nectar source at the location specified by a dance. In both of these experiments, the feeding station was placed on a small boat and moved out into a small lake in small stages so that the platoon of foragers trained to that station continued to visit it even when it was in the middle of the lake. In both cases, the experimenters observed that, when the station was in the middle of the lake, the dances of that platoon ceased to recruit additional foragers. In one case, as the boat approached a small island, recruits reappeared—but at the point on the island shore nearest to the boat, not at the boat itself. The recruited bees reinterpreted the meaning of the message on the basis of other information on their map.

3.3 | Differing homomorphisms

The sets of messages with navigational import are vector spaces with different formal properties. Bearings are points in a one-dimensional circular vector space, the space occupied by points on a unit circle centered on the navigator. Ranges are points in a one-dimensional unipolar linear vector space. Locations are points in a vector space with two bipolar dimensions. In every case, however, the messages (points) in these spaces may be represented by strings composed of either just one number or of two numbers, depending on the dimensionality of the space. Moreover, it is possible to devise forms for the numbers themselves that facilitate the computational operations by which the meanings of those messages—the meanings of the vector symbols—arise.

Addition is fundamental to all of the navigationally relevant meanings considered. However, in the case of bearings, the addition must be homomorphic to the combining of directed (signed) angular distances on the unit circle. In the case of locations, on the other hand, it must be homomorphic to the linear addition of directed distances. Gallistel (2017) explains how both are accomplished in computers that use the two's complement fixed-point binary incarnation of the rational numbers. The two's complement incarnation enables the machine to use the same addition machinery that works for unipolar (unsigned) numbers to add signed numbers. For linear addition, the system keeps the carry digit, if any, that is generated by the addition of left-most (most significant) binary digits. For circular addition, it drops the carry digit, in which case addition maps all possible pairs of n-bit number symbols back into the set of n-bit number symbols in a way that is homomorphic to the way that directed angles combine. These formal aspects of symbol-manipulating machinery underscore the point that the meanings of symbols arise from their referents, from the formal characteristics of the symbols themselves (Aristotle's formal cause), *and* from how the symbols are processed. The physical form of the symbols and the way in which that form interacts with the physical form of the computational machinery must together be such that the results of a computation are homomorphic to the properties of the world to which the symbols are assigned referents by the reference-assigning machinery. The complexity of the preceding sentence is an indication of why it has been hard to understand how meanings arise in physically realized representational systems. It is not easy to explain how and why a symbol in such a system comes to carry information *about* something out there in the world because the answer requires considering four things: the referent, the form of the symbol, how that form is processed, and how the referent and the processing are related (the nature of the homomorphism).

4 | EXAMPLE 2: COLOR

4.1 | The ubiquity of vector space representations in brains

Vector spaces are natural when it comes to representing the spatial aspects of the world, and the historical development of vector algebra and matrices is rooted in the numerical representation of geometric truths. However, this form of representation—vector spaces as sets of messages—has spread far beyond geometry. It is difficult to exaggerate the importance of vector spaces and the operations defined on them in the operation of computing machines. It is becoming increasingly clear that brains also represent many nongeometric facts about behaviorally important aspects of the world as points in vector spaces (Theves, Fernandez & Doeller, 2019), facts such as odors (Stevens, 2015, 2016), faces (Chang & Tsao, 2017; Stevens, 2018), and colors (Chang, Bao & Tsao, 2017). The extraordinary potential of this form of representation became manifest long ago in the evolution of brains.

4.2 | Source statistics and efficient encodings

Spectral differences are represented in a vector space with three bipolar dimensions (Cornsweet, 1970). Because spectra are infinite dimensional and unipolar, there seems to be a disconnect between the formal properties of the referents (spectra) and the computational properties of the messages into which they map (color sensations). Reducing an infinite-dimensional input space—the space of all possible visible spectra—to a three-dimensional color space involves a radical loss of information because an infinity of different spectra map to any given location in color space (different spectra that produce the same color sensation are called metamers).

Shannon comes to our aid in understanding this seemingly profligate waste of spectral information. He stressed that an efficient communication system must be tuned to the statistics of the source. The source of most spectra in our environment are the surfaces off which ambient light reflects. The reflectance spectra of surfaces have huge internal covariances, that is, the amount of light reflected at one wavelength strongly predicts the amount reflected at neighboring wavelengths. The redundancy among the intensities of nearby components in a reflectance spectrum greatly reduces the amount of information actually available from real-world reflecting surfaces. It appears that, when this redundancy is taken into account, the three-dimensional color space captures a large part of the information available about differences in spectral reflectance *profiles* in the natural environment (Boker, 1997; Maloney, 2003). Again, the brain appears to stipulate an efficient encoding of possible messages about the spectral reflectance statistics of the world in which it operates, as Shannon's analysis would lead us to expect.

The focus on efficient encodings of the information actually available may also help us understand why the axes of the brain's representation of color are bipolar (signed, like north/south and east/west). Spectral points are unipolar; there is no such thing as a negative intensity at any wavelength of light. If the brain is focused on the spectral *profiles* that are common in the natural world, then those profiles may well be mutually exclusive: The existence of one profile precludes the existence of another profile that is its inverse (low where it is high and high where it is low).

Profiles do not, however, have direction and distance. Thus, although the brain's computational apparatus could compute the direction of one color from another and/or the distance—the bearing and range of yellow from green—it does not do so because there is no useful global homomorphism between these properties of the message space and the real-world attributes to which the messages refer.

The brain does compute bearings and ranges for spatial locations, and people have conscious access to the results and a feeling of necessary truth. If you ask people why you cannot go both north and south at the same time, they will say that it is because they are opposing directions. (And most will think it a stupid question.) If you ask people why there are blue-greens and purples—mixtures of blue and green or blue and red, perceived as such—but there are no yellow-blues, they are baffled. The answer is the same as for the north–south case—yellow and blue are opposite ends of the same axis in the brain's representation—but that part of the formal system is not used in computations on color messages. I consider this further evidence that meanings arise from a combination of the referents of symbols and the *referent-appropriate* computational operations defined on those symbols, the operations that produce appropriate inferences.

5 | CONCLUSION

The brain creates meanings when it creates the sets of possible messages without which information cannot flow from the world into the brain. Each such creation involves (a) the setting up of the set of possible symbols for those messages and the sensory and computational mechanisms that establish reference between aspects of the world and symbols in that set and (b) the specification of the computational operations appropriate for messages in that set. Together, these establish the inferences that may be drawn from a message. The meaning simply is the set of inferences that may be drawn. These inferences do not necessarily preserve consciously apprehended truth.

ORCID

Charles R. Gallistel https://orcid.org/0000-0002-4860-5637

REFERENCES

Boker, S. M. (1997). A measurement of the adaptation of color vision to the spectral environment. *Psychological Science*, *8*(2), 130–134.

Chang, L., Bao, P. & Tsao, D. Y. (2017). The representation of colored objects in macaque color patches. *Nature Communications*, *8*(1), 2064. https://doi.org/10.1038/s41467-017-01912-7

Chang, L. & Tsao, D. Y. (2017). The code for facial identity in the primate brain. *Cell*, *169*, 1013–1028.

Cornsweet, T. N. (1970). *Visual perception*. New York, NY: Academic.

Cover, T. M. & Thomas, J. A. (1991). *Elements of information theory* (2nd ed.). New York, NY: Wiley Interscience.

Dennett, D. C. (1986). *Content and consciousness* (2nd ed.). London: Routledge.

Dretske, F. (1981). *Knowledge and the flow of information*. Cambridge, MA: MIT Press.

Gallistel, C. R. (1998). Symbolic processes in the brain: The case of insect navigation. In Scarborough, D. & Sternberg, S., (Eds.), *Conceptual and methodological foundations. Volume 4 of an invitation to cognitive science* (2nd ed. 1–51). Cambridge, MA: MIT Press.

Gallistel, C. R. (1999). Coordinate transformations in the genesis of directed action. In B. M. Bly & D. E. Rumelhart (Eds.), *Cognitive science: Handbook of perception and cognition* (2nd ed., pp. 1–42). San Diego, CA: Academic.

Gallistel, C. R. (2012). Extinction from a rationalist perspective. *Behavioural Processes*, *90*, 66–88. https://doi.org/10.1016/j.beproc.2012.02.008

Gallistel, C. R. (2017). Finding numbers in the brain. *Philosophical Transactions of the Royal Society. Series B*, *373*, 2017119. https://doi.org/10.1098/rstb.2017.0119

Gallistel, C. R., Craig, A. R. & Shahan, T. A. (2014). Temporal contingency. *Behavioural Processes*, *101*, 89–96. https://doi.org/10.1016/j.beproc.2013.08.012

Gallistel, C. R., Craig, A. R. & Shahan, T. A. (2019). Contingency, contiguity and causality in conditioning: Applying information theory and Weber's law to the assignment of credit problem. *Psychological Review*, *126*, 761–773.

Gibbon, J., Berryman, R. & Thompson, R. L. (1974). Contingency spaces and measures in classical and instrumental conditioning. *Journal of the Experimental Analysis of Behavior*, *21*(3), 585–605.

Gould, J. L. & Gould, C. G. (1995). *The honey bee* (2nd ed.). New York, NY: W. H. Freeman.

Hallam, S. C., Grahame, N. J. & Miller, R. R. (1992). Exploring the edges of Pavlovian contingency space: An assessment of contingency theory and its various metrics. *Learning and Motivation*, *23*, 225–249.

Hammond, L. J. (1980). The effect of contingency upon the appetitive conditioning of free operant behavior. *Journal of the Experimental Analysis of Behavior*, *34*, 297–304.

Kamikouchi, A., Inagaki, H. K., Effertz, T., Hendrich, O., Fiala, A., Gopfert, M. C. & Ito, K. (2009). The neural basis of drosophila gravity-sensing and hearing. *Nature*, *458*(7235), 165–171.

Kinney, J. B. & Atwal, G. S. (2014). Equitability, mutual information, and the maximal information coefficient. *Proceedings of the National Academy of Sciences*, *111*(9), 3354–3359. https://doi.org/10.1073/pnas.1309933111

Maloney, L. T. (2003). Surface colour perception and environmental constraints. In R. Mausfeld & D. Heyer (Eds.), *Colour perception: Mind and the physical world*. New York, NY: Oxford University Press.

Menzel, R., Kirbach, A., Hass, W.-D., Fischer, B., Fuchs, J., Koblofsky, M., … Greggers, U. (2011). A common frame of reference for learned and communicated vectors in honeybee navigation. *Current Biology*, *21*, 645–650.

Rieke, F., Warland, D., de Ruyter van Steveninck, R. & Bialek, W. (1997). *Spikes: Exploring the neural code*. Cambridge, MA: MIT Press.

Shea, N. (2018). *Representation in cognitive science*. Oxford: oxford University Press.

Schürch, R. & Ratnieks, F. L. W. (2015). The spatial information content of the honey bee waggle dance. *Frontiers in Ecology and Evolution*, *3*, 22. https://doi.org/10.3389/978-2-88919-765-1

Searle, J. R. (1980). Minds, brains and programs. *The Behavioral and Brain Sciences*, *3*, 417–457.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell Systems Technical Journal*, *27*(379–423), 623–656.

Sparks, D. L. (1988). Neural cartography: Sensory and motor maps in the superior colliculus. *Brain, Behavior & Evolution*, *31*, 49–56.

Stevens, C. F. (2015). What the fly's nose tells the fly's brain. *Proceedings of the National Academy of Sciences (USA)*, *112*, 9460–9465.

Stevens, C. F. (2016). A statistical property of fly odor responses is conserved across odors. *Proceedings of the National Academy of Sciences*, *113*(24), 6737–6742. https://doi.org/10.1073/pnas.1606339113

Stevens, C. F. (2018). Conserved features of the primate face code. *Proceedings of the National Academy of Sciences*, *115*(3), 584–588. https://doi.org/10.1073/pnas.1716341115

Sun, Y., Liu, L., Ben-Shahar, Y., Jacobs, J. S., Eberl, D. F. & Welsh, M. J. (2009). TRPA channels distinguish gravity sensing from hearing in Johnston's organ. *Proceedings of the National Academy of Sciences*, *106*(32), 13606–13611. https://doi.org/10.1073/pnas.0906377106

Tautz, J., Zhang, S. W., Spaethe, J., Brockmann, A., Si, A. & Srinivasan, M. V. (2004). Honeybee odometry: Performance in varying natural terrain. *PLoS Biology*, *2*, 915–923.

Theves, S., Fernandez, G. & Doeller, C. F. (2019). The hippocampus encodes distances in multidimensional feature space. *Current Biology*, *29*, 1226–1231. https://doi.org/10.1016/j.cub.2019.02.035

von Frisch, K. (1967). *The dance-language and orientation of bees*. Cambridge, MA: Harvard University Press.

Wray, M. K., Klein, B. A., Mattila, H. R. & Seeley, T. D. (2008). Honeybees do not reject dances for "implausible" locations: Reconsidering the evidence for cognitive maps in insects. *Animal Behaviour*, *76*, 261–269.

How to cite this article: Gallistel CR. Where meanings arise and how: Building on Shannon's foundations. *Mind & Language*. 2020;1–12. https://doi.org/10.1111/mila.12289