Learning Without Error

Erez Braun and Shimon Marom

Learning is a process of growing success in a fixed environment. One can only speak about learning when behavior noticeably increases the efficiency with which information is processed so that desirable states are reached, errors are avoided, or a portion of the world is controlled[1]. The concept of error plays a key role in general learning theories, and reinforcement learning in particular. In these theories[2] a dedicated entity is invoked, whose function is to compare the state of the system to a desirable state, and to produce an error signal that drives the system to change. That such is the case in machine reinforcement-learning protocols, one cannot argue. But, what about learning in biological systems? We submit that, if not for deep philosophical reasons, the schematics portrayed for reinforcement learning in machines cannot hold for biological systems. For an error signal to be produced, states should be measured and compared; but since biological states are practically infinite object that are not local in time nor in space[3,4], the scales and standards required for measurements and comparisons do not exist. We provide examples for state-space immensity[5] at three levels of biological organization that are intimately related to the subject matter: molecules, cells and behavior. We then comment on the impacts of this immensity on learning and on the practice of experiments in biology.

As a representative of state space immensity at the molecular level, consider the number of possible states in which particular proteins, the voltage-gated ionic-channels family, may reside. These proteins are responsible for shaping the time-amplitude envelope of neuronal and synaptic signals, and their state space is described as highly relevant for neuronal activity (the interested reader is referred to the authoritative monograph by B. Hille[6]). Molecular biologists and physiologists have convincingly shown over the past two decades that the number of types of ionic channel proteins that a single neuron expresses at any given point in time is in the order of ten. Each type of channels is composed typically of five to ten sub-types that form combinatorial structures that have different functional consequences. Depending on that combinatorial structure, the channel proteins are extensively engaged in

E. Braun (✉)

Network Biology Research Laboratories, and Departments of Physics (E.B.) and Medicine (S.M.), Technion, Haifa 32000, Israel

e-mail: erez@physics.technion.ac.il

interactions with ligands and with other proteins, interactions that have significant functional consequences. Furthermore, recent estimates and measurements indicate that even a bare, stand-alone ionic channel protein, has intrinsic to it a large number of possible states[7]. Other examples may be given for molecular immensity within the domains of synaptic functions, second messengers and related signal transduction pathways.

The molecular immensity exemplified above does not disappear at higher levels of brain organization. Consider, for instance, a cortical neuron in its network: It is known that the number of synapses impinging upon a single neuron in the human cortex is well beyond one hundred. In fact, estimates are that each cortical neuron is affected by the states of thousands of other neurons (the interested reader is referred to the monograph by Abeles[8]). This means that, even for the most simple case, in which the spatial and temporal attributes of a single neuron are not considered, and even if one assumes that synaptic inputs come in packets of several hundred of correlated synapses, the number of possible single-neuron states is immense. Indeed, careful analyses of neural activity time series, recorded from neurons in wake animals, show that these neurons have no preference for any limited set of uniquely defined states. Rather, their activity characteristics are consistent with a continuum-of-states model. Strong claims were made, supported by exhaustive statistical analyses, pointing to the curious fact that the number of neuronal states, reflected in the activity of neurons, is practically limited only by the experimental conditions and the rituals of statistical analyses[9].

What about overt behavior? Unlike the case of neurobiology, where the space of possible brain states is defined, the language is agreed upon and the problem is of identifying relevant states, in psychology there seems to be no consensus on the actual definition of the state-space itself[10]. This alarming situation is further complicated by the fact that whereabouts precise analysis is feasible, experimental findings indicate that behavior cannot be decomposed to a fathomable number of uniquely defined states. Consider, for example, an experimental psychology field that is (arguably) most approachable for scientific analysis, namely- the study of memory. In 1885, Ebbinghaus reported his seminal introspective study of human memory. He demonstrated that while the retention of nonsense syllables decreases as time elapsed from the initial learning increase, the rate at which that decrement occurs monotonically slows down, that is to say, not fixed. Ebbinghaus intuited that the mathematical form of the forgetting function is logarithmic. Unlike the exponential function commonly used to describe relaxation data, which would imply a fixed rate of forgetting in that case, Ebbinghaus's logarithmic function implies a rate of forgetting that depends on the time elapsed since the learning. Thus, if one accepts the physical-chemistry principle equating a single state with a single rate, Ebbinghaus's interpretation implied no unique memory states because no unique rates are found[11]. Notwithstanding a transient belief in the short-term memory theory during the early 1970's, cognitive psychologists now have ample evidence in favor of Ebbinghaus's interpretation. In fact, re-analysis of Ebbinghaus's data, taken together with a host of new data, indeed confirm that memorizing and forgetting cannot be described by a fathomable set of uniquely defined rates[12], and therefore no fathomable set

of uniquely defined states are involved. Thus, we see that in the behavioral level of organization the problem is twofold. First and foremost, behavioral states are ill defined, leading to a conceptual alienation. Second, behavioral entities that are approachable for precise analysis reveal an immense space of possible states.

The examples given above show that immensity of state space reigns at every level of biological observation, from molecules to behaviour. We conclude that the scales and standards required for the identification of a state and hence an error cannot be constructed in a biological system. This is the case already at the molecular level and definitely at higher levels (cells, organs, organisms and populations), which are aggregates of molecules. No error assignment can be exercised under such conditions. While the above examples are from the discipline of neural sciences, the picture is practically identical in many other biological systems (e.g. immune system, genetic networks and development).

Another possibility to construct scales and standards is by attributing functionality to a biological system. In that case error would be defined in terms of deviations from the "standard" functionality. While functionality is the hallmark of biological systems, it is an alien concept for the other natural sciences. In that respect, biology is closer to engineering sciences and man made machines. However, in contrast to engineering, biology is an historical science; presently observed configurations reflect accumulation of accidental events over evolutionary time scales, selection processes, multi-functionality, an immense number of entailments between functions, redundancies and overlap within and between levels of organization. These facts, taken together with the above-mentioned immensity of degrees of freedom preclude the possibility of adapting a given functionality as a standard for error detection. As a result, assigned functionality reflects the point of view of a given observer rather than that of the "designer". For instance, going back to the channel protein mentioned above: It is known that the protein can function only within a narrow voltage range around a set point determined by the gradient of ionic concentrations across the cell membrane. From the point of view of electrical functionality one may assume that the ionic concentration has been optimized to support the activity of the ionic channel proteins. This functionality of ionic concentrations, however, cannot serve as standard for error detection because the set point is similar in practically all cell types, including those that are not generators of electrical signals. The fact is that the ionic concentration gradients are key determinants of many, unrelated physiological processes within the same cell; hence the requirement for evolutionary selection. In order to understand the design principles behind ionic concentrations one needs to uncover the entire historical path and the set of all the involved interactions. Indeed, understanding the functional relations between the components of the biological system is possible only from the evolutionary point of view.

From the above we are again forced to conclude that lack of standard makes the concept of error irrelevant in the context of biological systems. This conclusion introduces serious problems in attempts to understand the biology of learning in its wider sense. For whatever definition we use for learning, a measure is required for the gap between present and desired configuration; in other words, a measure of the error is

required. Under these conditions such an error measure is translated to driving force that modifies the configuration of the system, aiming at error reduction. Lack of scales and standards thus presents us with a challenge to understand learning in the biological context. How does the immune system learn to identify pathogens? How does a colony of ants learn to construct nests, or to form patterns of tracks towards food sources? How do bacteria swim upstream food gradients given a single bacterium small scale? In fact, how do we, human beings, learn? What are the mechanisms allowing for adaptive reconfiguration of the immense number of entities involved in learning?

The common feature in all learning phenomena is the existence of exploration in configuration space. What makes learning in biological systems unique is the fact that, unlike man-made machines, the driving force for the exploration does not scale with the gap between present position and desired one. Rather, it is only dictated by local measures, irrespective of its distance from target.

For example, let us consider learning in the neural system. We know, from every day experience, that it is widely accepted to use evaluative-concepts in descriptions of learning in psychology; i.e., "appropriate" behaviour is "rewarded", "right" actions are "positively reinforced" and "wrong" are "negatively reinforced" (i.e. "punished"). Such language usage implies that in learning, surely, error measures are used and therefore standards are required. At the beginning of the twentieth century, the modern terminology of learning was established. Rules of association by simultaneity and temporal sequences involved in instrumental conditioning were defined. It became clear that the concepts of reinforcement, reward, punishment, are extremely useful in describing and controlling behaviours. Attempts to understand how the concept of reward is realized in a biological world that lacks standards were made by eminent psychologists such as Hull[13] and Guthrie[14] over fifty years ago and even earlier by Freud[15] and James[16]. The resulting learning theories, which may collectively be referred to as Drive Reduction theories, stress the effect of reward on the driving stimulus. Specifically, the reward acts to reduce the stimulus that drives the exploration process. This reduction is based on local cues and precludes the acquisition of new stimulus-response entailments. Sharpening the stimulus-response entailment, in turn, is achieved through a selection process. Such description of learning in neural systems classifies the operation of the brain as a Darwinian process, similar to the other above-mentioned biological systems. That is, no separate neural rewarding entity is postulated or needed for shaping behaviour. In fact, one may find texts that explicitly reject mapping of evaluative behavioural concepts to defined brain entities, suggesting that the concept of error does not belong to the neural system itself, but rather to the larger complex that contains the environment, the system and the observer. Here is, for instance, what Guthrie said in his presidential address to the American Psychological Association in 1946:

> Psychologists who think in terms of punishment and reward have almost uniformly neglected to note how the animal at the time responded to the punishment or to the reward, and the role this played in subsequent behavior. The resulting generalization is inevitably an attempt to link the intentions of the experimenter (intentions to reward or punish) with good or bad behavior on the part of the animal. Punishment and reward are, objectively viewed, stimuli acting on the animal's sense organs, and their effect must be mediated through the

animal's nervous system and appear in muscular contraction or glandular secretion. Since levers and loops and mazes are not innervated, the operations of these devices are incidental to the actual learning which the living animal performs.[17]

In spite of the above, the prevailing trend is to describe biological processes, including brain functions, in mechanistic terms of error detection, non-local driving forces and optimization processes[18]. Indeed, parts of biology may be described in such terms; the price being loss of the global, system point of view. It is acknowledged that mechanistic approaches to biology reflect present technological forefronts and are effective for specific applications (e.g. medical treatments). A most notable example is the interaction between the impressive technological developments of computing machines and brain research. While, as mentioned above, the general learning theories of the early 1920's explicitly advocated avoidance of attempts to map evaluative concepts to specific brain structures, nowadays such mapping dominates neuroscience. This shift reflects the dominance of the computer paradigm in brain research. Most algorithms used for effective machine learning are supervised ones, where an additional source of information, of knowledge of the error, dictates the drive and directs the learning process. Interestingly, when cognitive psychology, heavily relying on computational theories, entered the arena and practically removed behaviourism and general learning theories from the scene, it brought with it the error-based algorithmic computational approach.

History teaches us that the duality of mechanistic and Darwinian aspects will continue to drive the biological research. Jerne describes this historical pattern in a text written in 1967:

> Looking back into the history of biology, it appears that wherever a phenomenon resembles learning, an instructive theory was first proposed to account for the underlying mechanisms. In every case, this was later replaced by a selective theory. Thus the species were thought to have developed by learning or by adaptation of individuals to the environment, until Darwin showed this to have been a selective process. Resistance of bacteria to antibacterial agents was thought to be acquired by adaptation, until Luria and Delbrück showed the mechanism to be a selective one. Adaptive enzymes were shown by Monod and his school to be inducible enzymes arising through the selection of pre-existing genes. Finally, antibody formation that was thought to be based on instruction by the antigen is now found to result from the selection of already existing patterns. It thus remains to be asked if learning by the central nervous system might not also be a selective process; i.e., perhaps learning is not learning either.[19]

This duality presents a challenge to experimental biologists. Setting up experimental designs aimed at exposing mechanistic aspects of a given biological system is a natural extension of prevailing paradigms in engineering and physical sciences. However, uncovering the Darwinian aspects of biological systems requires new experimental concepts. The experimental design, in such cases, should allow the observed system to control its driving forces based on interactions with the environment. In other words, standards reflecting the experimental constraints should be eliminated. Results from such experiments will enable development of comprehensive understanding the unique aspects of biology, and may serve as a basis for a paradigm shift in engineering.

Notes

1. Klaus Krippendorff's Dictionary of Cybernetics (URL= http://pespmc1.vub.ac.be/ASC/LEARNING.html)
2. E.g., Sutton and Barto (1998)
3. Rosen (1991)
4. A term introduced in Elsasser (1987)
5. Elsasser (1987)
6. Hille (1992)
7. E.g., Toib et al. (1998); Ellerkmann et al. (2001); Gilboa et al. (2005)
8. Abeles (1991)
9. E.g., Teich et al. (1997)
10. E.g., the alienation between the concepts of cognitive psychology and those of psychodynamics.
11. Interestingly, the modern concept of scale-free distribution does fit Ebbinghaus's description; scale-free distributions are often interpreted as indicating (practically) continuum of states. See also note 7 and 9 above.
12. E.g., Wixted and Ebbesen (1997)
13. Hull (1943)
14. Guthrie (1946)
15. Freud (1895/1966)
16. James (1890)
17. Guthrie (1946).
18. E.g., Hollerman and Schultz (1998)
19. Quarton et al. (1967), p. 204

References

Abeles, M. (1991). Corticonics: Neural Circuits of Cerebral Cortex. Cambridge: Cambridge University Press.

Ellerkmann, R. K., V. Riazanski, et al. (2001), J Physiol 532(Pt 2): 385–97

Elsasser, W. M. (1987). Reflections on a Theory of Organisms: Holism in Biology. Baltimore: The Johns Hopkins University Press.

Freud, S. (1895/1966). Project for a Scientific Psychology. The Standard Edition of the Complete Psychological Works of Sigmund Freud. Pre-Psychoanalytical Works and Unpublished Drafts. J. Strachey. London: The Hogarth Press. 1: 283–360.

James, W. (1890). Principles of Psychology. New York: Dover Publications, Inc.

Gilboa, G., R. Chen, et al. (2005), J Neurosci 25(28): 6479–89.

Guthrie, E. R. (1946). "Psychological facts and psychological theory." Psychol Bull 43: [Presidential address of the APA, Evanston, Illinois, 1945].

Hille, B. (1992). Ionic Channels of Excitable Membranes, Sunderland, MA: Sinauer.

Hollerman, J. R. and W. Schultz (1998). Nat Neurosci 1(4): 304–309.

Hull, C. (1943). Principles of Behavior. New York: Appleton-Century-Crofts.

Quarton, G. C., et al. (1967) eds., In the Neurosciences: A study program (Rockefeller University Press), pp. 200–205

Rosen, R. (1991). Life Itself. New York: Columbia University Press.

Sutton, R. and A. Barto (1998). Reinforcement Learning: An Introduction. Cambridge, MA: MIT Press.

Teich, M. C., C. Heneghan, et al. (1997), J Opt Soc Am A Opt Image Sci 14(3): 529–546.

Toib, A., V. Lyakhov, et al. (1998), J Neurosci 18(5): 1893–1903.

Wixted, J. T. and E. B. Ebbesen (1997), Mem Cognit 25(5): 731–739; and references therein.