The Foundations of Control and Cognition: The Every Good Regulator Theorem

Robert Stone[1] and Bradly Alicea[1,2]
rob.stone.academic@gmail.com, bradly.alicea@outlook.com
[1] Orthogonal Research and Education Lab, Champaign-Urbana, IL [2] OpenWorm Foundation, Boston, MA

Abstract

In this paper, we will attempt to bring the concept of the Every Good Regulator Theorem (EGRT) into the modern scientific discourse. The EGRT, first proposed 50 years ago by Conant and Ashby, will be connected to a broader set of concepts that are of interest to a modern, interdisciplinary set of issues. We will also engage in a discussion on how the EGRT relates to models of game theory and the mind in ways that recast so-called "good" regulation as a form of cognition. Yet we also turn the very notion of cognition on its head by relating this non-purposeful notion of generalized cognition to universal properties of regulation. In doing so, we propose specific models for very simple regulation (zeroth-order switching) and very complex regulation (cybernetic convolution architecture). To conclude, we reconsider the role of modeling itself in regulating a complex system, which may encourage people to reconsider the relevance of the EGRT and associated cybernetic approaches.

I. Introduction

In the history of scientific discovery, there have been examples of certain persons or approaches being considered 'out of step' with the dominant scientific or philosophical trends of the time. As such, they risk falling down a deep well in our cultural landscape, with their work's efficacy lost to subsequent generations. If their work has merit, it may be considered ahead of its time by future generations. The timing of a given theory or great idea is largely determined by cultural and cognitive biases that favor the dominant paradigm [1]. In other cases, ideas at the paradigmatic vanguard end up resurrected in a more pragmatic way. The acceptance of such ideas occurs either gradually or in one fell swoop at a later point in time [NOTE 1]. Let us keep this in mind as we discuss Ronald C. Conant and W. Ross Ashby's seminal work "Every Good Regulator Theorem" [2] (EGRT):

> "[The EGRT is]….a theorem is presented which shows, under very broad conditions, that any regulator that is maximally both successful and simple must be isomorphic with the system being regulated…….Making a model is thus necessary." [2]

The EGRT characterizes regulation with respect to cybernated control systems (see Definition A). In the case of Ashby and Conant [3], the EGRT developed within the context of several intersecting traditional fields. These include algorithmics, information theory, systems theory, and behavioral science. In such a context, models are exceedingly important. Given the reliance of the EGRT concept on inference and propositional thinking, there is an essential reliance on models (see Definition B). In fact, the EGRT exists at such a high level of abstraction that even with a high degree of specification may not be directly applicable in the real world [4].

However, there are certain advantages of cybernetic modeling that make their cross-contextual application useful.

II. Background

We shall explore the notion of modeling as phenomenology, and consider its implications. Systems engage in modeling not simply to purposely regulate their environments, but rather to reactively respond to input stimuli in a way that maintains higher-level states [5]. This ability to model becomes part of their structure at the most basic of levels, though it would be fair to say most modeling (in the way we will use the word) is the result of cognitive processes, although regulation in simple biological systems from bacteria to genetic material itself reveals the creation and interaction of models and modeling. We are more concerned with models created by brains, and most concerned with interactions of the brains of humans. These models are complex systems, and interact in ways that require us to consider different forms of reasoning, relationships between parts, and a transcendence of reductionist thinking [6].

The constructivist [NOTE 2] might argue that such metacognitive dynamics [7] would influence one's proposed scientific model. Like Shakespeare's Hamlet, however, the question of whether or not to model (or be) is one of survival, whether that survival be genetic or memetic. Rather than reviewing the EGRT proof in a step-by-step fashion [NOTE 3], let us discuss its potential significance in a variety of use-cases. In the process, we will be transcending the traditional boundaries of what is normally considered autonomic, a 'choice', or even what is 'cognitive'.

Simply put, the Every Good Regulator Theorem says that regulators (see Definition C) operate on approximations (e.g. models) of the thing they are regulating. This requires a mapping (see Definition D) of the system being regulated to the model. While one might consider the activities of encoding and translation to be inherently cognitive, genomic systems also perform biological control functions in the absence of cognition [8, 9]. In genomic systems, what matters is not intent (or goals), but accuracy (or fidelity) with noise tolerance. Accuracy of the approximated model influences the quality of regulation. Thus, the agency is not required of any single system component. Rather than an actively goal-oriented (see Definition E) criterion, what we observe here is passively goal-oriented system output. In terms of agency, it is the organism that plays the role of goal-seeker, while genes are (in the sense of agentive behavior) along for the ride. In evolutionary terms, genes do not need to evolve to meet a specified goal, but merely be fit enough for survival and transmission to subsequent generations. Indeed, to survive as a unit in an interrelated system, a regulating machine must construct an interactive model that includes inputs, outputs, and feedback.

Let us consider a couple cases of regulatory dynamics, which may be valuable in understanding the importance of this theorem. We can then move on to what this could mean for both further theoretical development and practical application. A good place to begin in cognitive science is game theory [10]. One of the most simple, and explanatory strategies in the Prisoners' Dilemma game is the tit-for-tat strategy [11], a strategy of simple reciprocity proven to be the most effective strategy in a series of famous Iterated Prisoners' Dilemma (IPD) game theory computer tournament 'experiments' conducted by Robert Axelrod and won by Anatol

Rapoport utilizing this strategy for his ‘player’ [12]. In this two-player, 2x2 game, the tit-for-tat strategy is simple: ‘Do unto others as they have done unto you’ after an initial good faith move of cooperation. The strategy is simply to copy your opponent's behavior (see Definition F). If the opposing agent cooperates, so does the tit-for-tat strategizing agent; if they defect, the tit-for-tat strategist follows suit. The intended outcome of the strategy is to move the exchange towards an equilibrium (though this is not the only possible outcome, nor is the strategy perfect).

Of specific interest here is that the mechanics of the strategy requires a model to be held in memory by the agent employing tit-for tat (a one-bit cooperate/defect model), regardless of the strategy employed by the other agent (whether that be a more sophisticated maximizing strategy, or random selections). While an economist might view this as free-riding behavior by one of the two agents, the selection of tit-for-tat by both players can produce a cooperative equilibrium, such as in the evolution of reciprocal altruism in biological systems. [NOTE 4] The EGRT suggests that the greater the memory of an agent, and the longer it has the opportunity to observe and integrate the moves of its opponent, the greater its potential for effective regulation.

Over time, this can lead to greater accuracy for the agent’s cognitive model and a more stable equilibrium game outcome. Further, this equilibrium state can be long-lasting, given extended memory capacity for more detailed models, and may evolve towards ‘a conspiracy of doves’, within a game of *homo lupus homini* [NOTE 5]. An agent with a greater memory capacity can also employ more elaborate (or deeper) strategies over time. This development of deeper strategies may also feedback into modifying its model of the external world [13]. Overall, the capability to regulate behavior of other players depends on the inferential and predictive capacities of each player’s model: in a highly complex competitive game environment, a good regulator has a superior model, or it will find itself regulated by a competing agent in the game, especially as the behaviors get more complex and follow leader-follower conditions [NOTE 6].

> “The theorem has the interesting corollary that the living brain, so far as it is to be successful and efficient as a regulator for survival, must proceed, in learning, by the formation of a model (or models) of its environment.” [2]

III. Further Considerations

Let us now consider a more complicated scenario where we might be able to uncover the universal components of the EGRT phenomenology. The context will be two people on a blind date (this can actually be a complicated scenario). If one has been in one of these (terrifying) contexts, then one can already see where we are going. The cognitive agents are continually competing to increase the efficacy of their models of the other agent, while also attempting to constrain the modeling of the other agent towards a compact image they prefer. Although rarely implemented successfully, winning strategies include accurately modeling the other actor and influencing the state of their mental model. This can include both elaborate, multi-step strategies, and simpler strategies, the complexity of which does not indicate their effectiveness. If the goal is a continuation of relations, the acquisition and intentional obfuscation of information occurs at appropriate times and in appropriate ways. This moves us beyond the assumptions of Conant [14], who made the connection between Shannon Information [15] and the concept of efficient regulation. Furthermore, this information has contextual value, and allows us to apply approaches such as game theory to understanding the informational aspects of cybernetic regulation [16]. As in most scenarios involving imperfect or asymmetrical information [17], superior models must have the opportunity for greater influence on a given interaction [18], and thus have control of the regulation to guide it toward a desired outcome.

Does regulation even require what we would call cognition? This of course depends on our definition of cognition and regulation. However, let us consider that a bacterium does not have a "cognitive" or mental model of its environment, yet appears to have little trouble getting around and controlling some aspects of its landscape. The similarities between chemotactic sensation and mental models built upon multisensory stimuli serve as evidence for the universal character of the EGRT [NOTE 7]. In fact, Heylighen [19] has proposed that cybernetic regulation is a highly-generalized form of cognition. Yet do thermostats or other mechanical systems possess anything approaching what we consider cognition? While none of these has the cognitive capacity of a brain, they do have information processing capabilities from their physical or electronic structure, memory states, and crude models of how things 'should' be, towards which they regulate conditions.

Non-cognitive systems possessing these characteristics are obviously still capable of rudimentary communication, control, decision making, and regulation, at least in an abstract sense. We should also expect some degree of continuity that crosses the boundary of the cognitive and non-cognitive, since cognitive systems evolved from less intentional ones with more rudimentary forms of behavioral control. As an intermediate example between cognitive and non-cognitive systems, aneural cognition [20, 21] characterizes cognitive behaviors without a formal nervous system. In cybernetic regulatory terms, aneural cognitive systems possess the model of cognition without the underlying biological mechanism. Agent-based and chemical systems exhibiting so-called minimal cognition [22, 23] also exhibit this property. This may seem to obfuscate our argument, but instead supports the idea that models are the primary mode of regulatory action in complex systems.

> "...success in regulation implies that a sufficiently similar model must have been built, whether it was done explicitly, or simply developed as the regulator was improved." [2]

IV. Universal Properties of Regulation

One way to view cognition and the universal properties of regulation is by characterizing the overall scope of a regulatory system. According to this view, we are not only concerned with the overall extent of regulation, but also the boundary conditions of the cybernetic model. This addresses two of the issues we have discussed thus far: the nested nature of regulation according to the EGRT, and the phenomenology of modeling itself.

When models of the world get very large and complex, they tend to develop properties of a *cybernetic convolution architecture* (CCA). A box and arrow version of a generic CCA is shown in Figure 1. In this CCA, there is a heterogeneity of feedback modes, from indirect feedback to dual reciprocal feedback. While CCAs abide by the conditions of the EGRT, two additional laws drive their evolution and diversification. Firstly, the law of requisite variety plays a more direct role in how different components are connected to the main graph. Units with limited Shannon Information will tend to connect to units with higher Shannon Information. The informational difference between the two will determine the mode of regulation between each unit.

The second additional law involves open-ended evolution. Open-ended evolution [24] involves evolution that proceeds without purpose or a target goal. This stands in contrast with the aims of classical cybernetics [25] and even modern behavioral science [26]. Once again, we can see this in a model of an expanding regulatory system. In a way similar to an expander graph [27], the connections increase over time to form a denser and more connected graph. These connections form in response to local adaptive requirements that change and emerge over time rather than through a master plan. While Figure 1 shows a linear form of a CCA, these regulatory systems can appear to be hairball graphs in very complex regimes.

Based on the CCA example and these two additional principles, we suggest a corollary to the law of requisite variety: models of a system must be at least as complex as the system itself to capture the full dynamics of that system. While this generally disallows prediction or sometimes even explanation, it provides a generative framework for new and adaptive states. But not all models of complex behavior need to be massively complex themselves. At the other end of the spectrum, we consider the role of *zeroth-order regulatory models* in governing state transitions in complex systems such as a biological cell or an animal brain. A zeroth-order model provides a switch between two macrostates that is regulated by emergent features of the underlying regulatory model. Examples of this type of regulation include light switches, Heaviside step functions [28], and first-order phase transitions [29]. Figure 2 shows a demonstration of this mechanism and a model of the output.

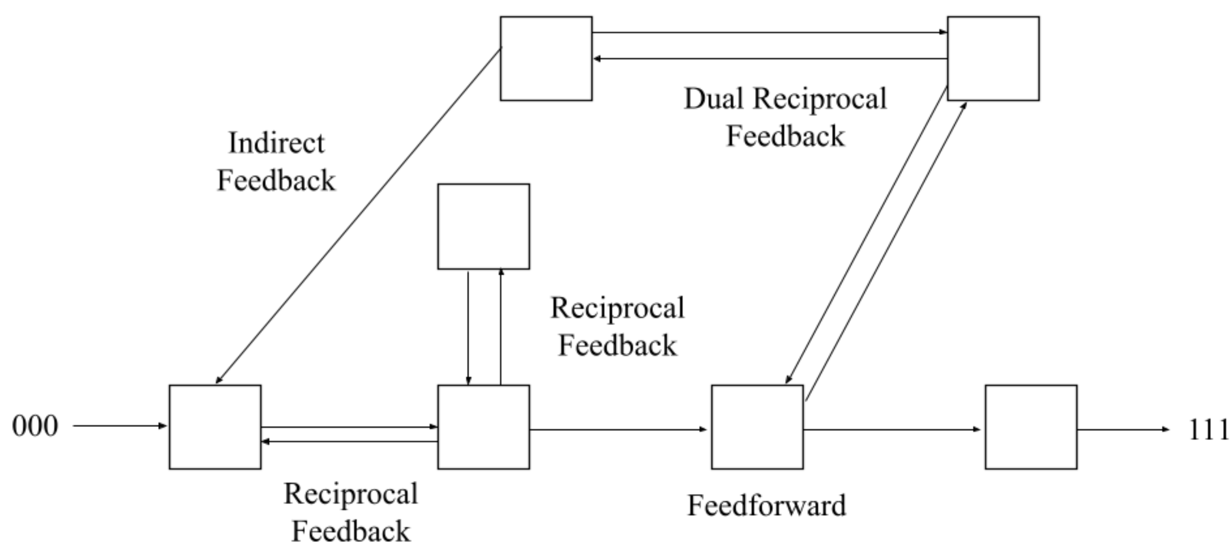

Figure 1. A linear example of the CCA model labelled with different types of regulation. 000 and 111 represent the input and output of the network, respectively.

In Figure 2, we see how switching operates in a simple, closed-loop feedback system. An input coming from a system is integrated with feedback in the form of a binary value (0 or 1). The feedback does not occur in the first time step, only after input is fed forward to an observer that monitors all current microstates of the system. A feedback signal of "0" allows the integrator to maintain the current input state, whereas a feedback signal of "1" shifts (or switches) the system to a new state (or dynamical regime). This state information is the output at every timestep of the model. In this case, good regulation is a simple dialogue between the integrator and observer to produce a switch in the macrostate. This type of supervised learning does not require goal-directedness or teleology of any kind, and thus is consistent with the principle of open-endedness.

Zeroth-order feedback becomes particularly useful in the context of our tit-for-tat game. The single bit of information content in a player's move corresponds to a binary switching event. To further place this in the context of Figure 2, we can turn our game into an iterated tit-for-tat, where there are many switching events over time. Using competing closed-loop feedback based switches, each player either decides to escalate (switch) or not to escalate (no switch). This yields the switching output shown in Figure 2. The observer embedded in each switching graph must decide to match their opponent, yielding the potential for synchronization of the two player's outputs. In such cases, this yields a game-theoretic equilibrium. We can also apply our leader-follower conditions to this game: when one player always initiates the interaction, they may constrain the second-run player, essentially forcing them to escalate and always switch, even in cases where a suboptimal equilibrium is reached. Playing this game with a CCA that governs switching is expected to often yield suboptimal behavior, regardless of leader-follower behavior.

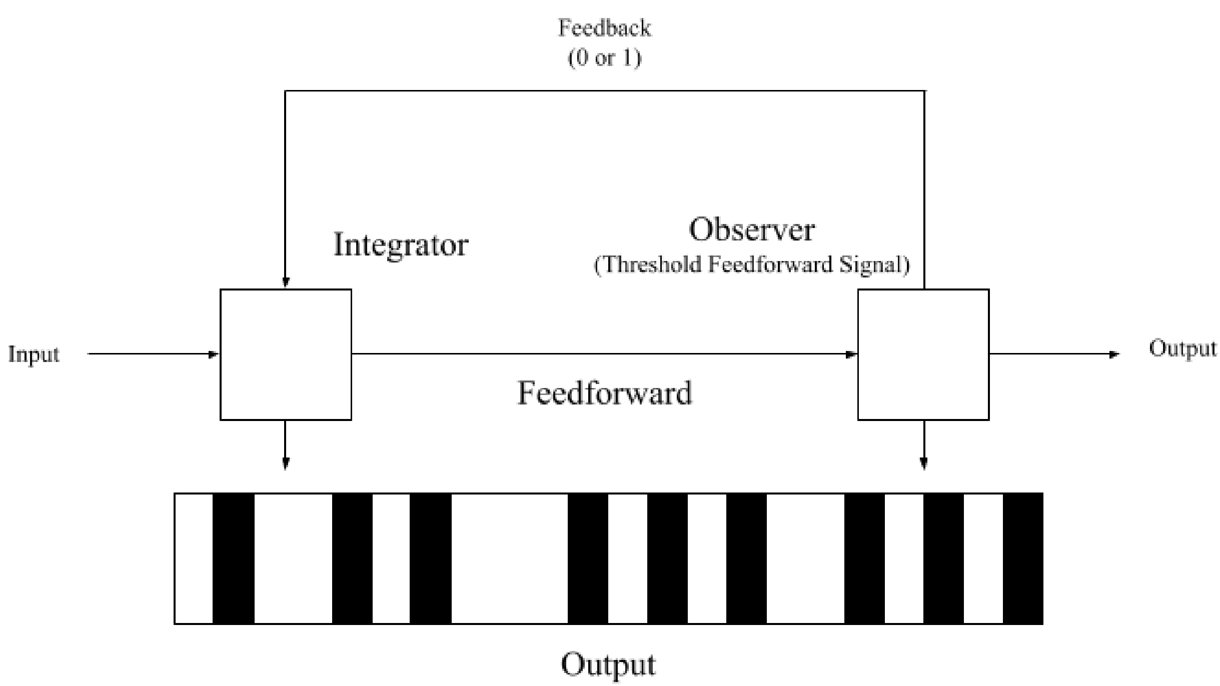

Figure 2. A demonstration of zeroth-order regulation in the form of switching between states.

V. Conclusion

Earlier, we had touched upon the history of scientific discovery, and contextual model building. A scientific theory is simply a model, and its value lies in its efficacy and repeatability (thus its' trustworthiness and ability to aid in regulation). Theoretical models have tended, historically, to shift from informal, conceptual models towards formal mathematical ones (as in Comte's Philosophy of Science). As a given model acquires more data, that data creates ever-more accurate model revisions with higher fidelity. The overall capacity to aid regulation increases via feedback. Thus, the model's value to humans increases. However, as noted by the example of non- paradigmatic 'ahead of their time' thinking, scientific thought does not exist in a vacuum, and the landscape conditions need to be aligned so that the model can prove fruitful. Consider how we are witnessing an explosion in robust formal mathematical and/or computer models either aiding or besting human cognitive efforts [30, 31]. Informational revisions of the model often occur faster than the landscape conditions change, so adaptive cross-contextual models may prove more successful in dynamic situations, such as ones which are developed by human thought and human cultural systems.

This ability to cross the boundary between cognitive and non-cognitive with models may challenge either our informal, colloquial conception of cognition or the universality criterion of the formal EGRT. As both features of cognition and more universal mechanisms, information processing, memory, communication, and selection can occur without any kind of cognitive superstructure. Perhaps the context of what we call "cognition" is too limiting. What about human cognition then is truly universal, and what is unique to a certain set of mechanisms and representational models? For example, are models of so-called cellular decision-making [32] an unduly anthropomorphic representation of cellular differentiation and metabolism, or is it drawing upon a common set of universal properties that can only be abstracted from the system by an appropriate model?

By reconsidering the EGRT, we reaffirm the role of cybernetic regulation in the natural world. Part of this reassessment involves thinking about the relationship between nature and our own conceptual models in new ways. We also need to consider a wider range of influences, from social complexity to theoretical biology. Bringing cybernetics into the mainstream of scientific inquiry also involves proposing more contemporary models of regulation that draw from complexity theory and can be implemented on large datasets. We hope this work will lead to new ways of thinking about the intersection of complexity and cybernetic models, even in ways that commingle the models and their phenomenology. This combination of causality, retrocausality, and interactivity may guide us towards more accurate, repeatable, and otherwise 'good' models.

> "Now that we know that any regulator (if it conforms to the qualifications given) must model what it regulates, we can proceed to measure how efficiently the brain carries out this process. There can no longer be a question about whether the brain models its environment: it must." [2]

Notes:

1) A classic example of this paradigmatic effect is Ignaz Semmelweis and his evidence for the germ theory of disease. While others had asserted this philosophically in previous centuries, he provided some of the first practical evidence for these ideas; however, this did not become part of the dominant medical paradigm until after his death. Similar historical examples would be Boltzman's commitment to atomism, or facets of Giordani Bruno's thoughts on cosmology.

2) In this essay, we will explore a number of topical domains, including some for which cybernetics is not typically the dominant paradigm. As Marshall McLuhan [33] might put it, our view of cybernetics is that the model is both the medium and the message. In this case, a model of process and feedback is something that occurs both in the mind of the observer and amongst the interactions of non-autonomous objects. Stuart Umpleby refers to this as a form of "second-order science" [34], where we consider the observer and their accompanying biases as part of the system. Our vision is a bit less unified than that provided by second-order science, yet nevertheless addresses a conceptual gap in systems science. Whereas the former (something occurring in the mind) may be clearly intentional, the latter (something occurring in nature) is definitely not. The properties of universality inherent in the modeling paradigm (cybernetics) may allow us to understand both of these types of dynamic phenomena in a common context.

3) For a deeper dive into the EGRT, Daniel Scholten's work [35] provides an introduction to the underlying concepts, as well as a rigorous analysis of the mathematical proofs featured in Conant and Ashby [2].

4) There are scenarios, when employing tit-for-tat, for the game to evolve towards instability and 'mutually assured destruction' scenarios when noise and errors come into play. Furthermore, some strategies (often permutations on tit-for-tat) may be contextually superior to it; however these strategies often require more memory, and advance knowledge (or a model) or the agent's tit-for-tat strategizing. This suggests that tit-for-tat is maximal for strategies that require finite memory resources. The purpose of this analysis is to show that even the simplest regulation strategies, in the simplest contexts, require a model, and that regulations' efficacy may only increase from more accurate modeling.

5) A Roman proverb "Man is a wolf to man", can be contrasted with the concluding hope of Dawkins' "The Selfish Gene" [36], wherein humans may create a more cooperative and altruistic society results from their ability to model future events and outcomes more effectively over time. Consider Hobbes 'De Cive', where Hobbes' conclusions suggest that mistrust (and warfare) between groups is a result of incomplete and noisy agent modeling due to their lack of memory. In such a scenario, only shorter-term strategies are possible. This is why the greater context of *homo lupus homini* is that a stranger is considered to be non-human (a wolf) by another man, unless that other person knows and trusts him. In this case, trust merely is an indication of an effective model with repeatable outcomes and good error control. The importance of this aside is in the universal nature of modeling and the EGRT in both micro- and macro- contexts.

6) The phenomenon of "being regulated" by someone else can be summarized nicely in the form of a Stackleberg equilibrium [19]. A Stackleberg equilibrium occurs in a two-player game in

which one player becomes the leader (first-mover), and the other becomes the follower (second-mover). The first-mover has the upper hand simply by constraining what strategies the follower can employ. In this scenario, a first-mover can also take advantage of being able to infer the second-mover's strategy. In a tit-for-tat game, the second-mover would copy the first-mover's strategy, and could quickly converge the game to an equilibrium point. This form of first-mover regulation has been used to model biological development, specifically the partitioning of space in early embryos [37] and the formation of neuronal networks (connectomes) [38].

7) An alternative way of thinking about comparisons between models (the commonalities between human thought and autonomous behavior in bacteria) is to assume that any mechanism is possible, but that the *controllability* of interacting parts is what is most important. High degrees of controllability allow for recovery from perturbations and the production of optimal outputs. Heylighen [39] further defines controllability as the relative homogeneity of outcome. If a system encounters repeated stimuli or motivating actions that are similar and produces a uniform output, the system is highly controllable. If this output is highly variable given the same scenario, the system possesses low controllability. Controllability is also dependent upon the degree of connectivity inside of the controller (black box) [40]. If most of the variables that constitute the regulating system are randomly connected, then there is a low degree of controllability. A more orderly set of connections indicative of hierarchical systems of high-order complexity will be more controllable. Despite the hierarchical nature of a human brain, the behavior of a bacterium would be more controllable than human cognition, as the degrees of freedom representing a human mind are much greater than the combinatorial complexity of a bacteria's behavioral repertoire. Correspondingly, a much larger CCA is required for regulation of the human mind.

8) This unit (a system) is constrained by boundary conditions that close the system off from some outside influences or information but not others in a controlled manner. These boundaries often consist of physical laws such as gravity or surface tension or barriers both physical and social. A systemic unit is also composed of subsystems selectively interconnected to each other. The systemic unit is itself a subsystem of a higher-level unit in a systems hierarchy.

Definitions:

A. System: A set of interrelated variables that operate together as a single, bounded, unit. For a more detailed definition, please see Note 8.

B. Model: The collection of representational mappings of a system that creates an approximation of that system as a more memory compact (yet incomplete) representation.

C. Regulator: A reactive or adaptive component that constrains the system's behavior (output) to a range of values. Examples of first-order regulators range from the throttle of an internal combustion engine to a counterweight on a scale.

D. Map: The functional relationship between the system and an approximate representation. An isomorphic map consists of every element in the system mapped to a corresponding element in the model.

E. Goal: An end-point upon which a system's behavior should converge. In goal-directed systems (e.g. prey acquisition in animal behavior), this can be done adaptively. In stochastic systems, the system will converge upon the goal due to behavioral constraints (e.g. minimization of energy in a thermodynamic system).

F. Behavior: The output of a system and its dynamics (over time). A system's behavior can range from the collective output of all subsystems to holistic, aggregate behaviors.

References:

[1] Kuhn, T. Structure of Scientific Revolutions. University of Chicago Press (1962).

[2] Conant, R.C. and Ashby, W.R. Every good regulator of a system must be a model of that system. International Journal of Systems Science, 1(2), 89–97 (1970).

[3] Ashby, W.R. Introduction to Cybernetics. Chapman and Hall (1962).

[4] Fishwick, P. The Role of Process Abstraction in Simulation. IEEE Transactions on Systems, Man, and Cybernetics, 18(1), 18-39 (1988).

[5] Brooks, R. Intelligence Without Representation. Artificial Intelligence, 47, 139-159 (1991).

[6] Klir, G.J. Facets of Systems Science. Springer, Berlin (1991).

[7] Kornell, N. Metacognition in Humans and Animals. Current Directions in Psychological Science, 18(1), 11-15 (2009).

[8] Ertel, A. and Tozeren, A. Human and mouse switch-like genes share common transcriptional regulatory mechanisms for bimodality. BMC Genomics, 23(9), 628 (2008).

[9] Gormley, M. and Tozeren, A. Expression profiles of switch-like genes accurately classify tissue and infectious disease phenotypes in model-based classification. BMC Bioinformatics, 9, 486 (2008).

[10] Gintis, H. Game Theory Evolving. Princeton University Press (2000).

[11] Imhof, L.A., Fudenberg, D., and Nowak, M.A. Tit-for-tat or Win-stay, Lose-shift? Journal of Theoretical Biology, 247(3), 574–580 (2007).

[12] Axelrod, R. and Hamilton, W.D. The Evolution of Cooperation. Science, New Series, 211(4489), 1390-1396 (1981).

[13] Liberatore, P. and Schaerf, M. Belief Revision and Update: Complexity of Model Checking. Journal of Computer and System Sciences, 62(1), 43–72 (2001).

[14] Conant, R.C. (1969). The Information Transfer Required in Regulatory Processes. IEEE Transactions on Systems Science and Cybernetics, SSC-5(4), 334-338.

[15] Shannon, C.E. A Mathematical Theory of Communication. Bell System Technical Journal, 27, 379–423 (1948).

[16] Novikov, D.A. (2016). Cybernetics: from past to future. Springer, Berlin.

[17] Rasmussen, E. Games and Information: an introduction to game theory. Blackwell Publishing (2006).

[18] Simaan, M. and Cruz, J.B. On the Stackleberg Strategy in Non Zero-Sum Games. Journal of Optimization Theory and Applications, 11(5), 533-555 (1973).

[19] Heylighen, F. Principles of Systems and Cybernetics: an evolutionary perspective. CiteSeerX, doi:10.1.1.32.7220 http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.32.7220 (1992).

[20] Gershman, S.J. and Balbi, P.E.M., Gallistel, C.M., and Gunawardena, J. (2021). Reconsidering the evidence for learning in single cells. eLife, 10, e61907.

[21] Baluska, F. and Levin, M. On Having No Head: Cognition throughout Biological Systems. Frontiers in Psychology, 7, 902 (2016).

[22] Godfrey-Smith, P. Individuality, subjectivity, and minimal cognition. Biological Philosophy, 31, 775–796 (2016).

[23] Hanczyc, M.M. and Ikegami, T. Chemical Basis for Minimal Cognition. Artificial Life, 16, 233–243 (2010).

[24] Taylor, T., Bedau, M., Channon, A., Ackley, D., Banzhaf, W., Beslon, G., Dolson, E., Froese, T., Hickinbotham, S., Ikegami, T., McMullin, B., Packard, N., Rasmussen, S., Virgo, N., Agmon, E., Clark, E., McGregor, S., Ofria, C., Ropella, G., Spector, L., Stanley, K.O., Stanton, A., Timperley, C., Vostinar, A., and Wiser, M. Open-Ended Evolution: Perspectives from the OEE Workshop in York. Artificial Life, 22(3), 408-423 (2016).

[25] Rosenblueth, A., Wiener, N., and Bigelow, J. (1943). Behavior, Purpose and Teleology. Philosophy of Science. 10(1), 21.

[26] Gomez-Marin, A. and Ghazanfar, A.A. The Life of Behavior. Neuron, 104, 25-36 (2019).

[27] Hoory, S., Linial, N., and Wigderson, A. Expander graphs and their applications. Bulletin of the American Mathematical Society, 43(4), 439–561 (2006).

[28] Weisstein, E.W. Heaviside Step Function. MathWorld: a Wolfram web resource. Retrieved November 29, 2021. https://mathworld.wolfram.com/HeavisideStepFunction.html

[29] Schroeder, M. R. Fractals, chaos, power laws: minutes from an infinite paradise. W.H. Freeman, New York (1991).

[30] LeCun, Y., Bengio, Y., and Hinton, G. Deep Learning. Nature, 521, 436-444 (2015).

[31] Ferrucci, D., Brown, E., Chu-Carroll, J., Fan, J., Gondek, D., Kalyanpur, A.A., Lally, A., Murdock, J.W., Nyberg, E., Prager, J., Schlaefer, N., and Welty, C. Building Watson: an overview of the DeepQA project. AI Magazine, Fall (2010).

[32] Kobayashi, T.J., Kamimura, A. Theoretical aspects of cellular decision-making and information-processing. Advances in Experimental Medicine and Biology, 736, 275-291 (2012).

[33] McLuhan, M. Understanding Media: the extensions of Man. McGraw-Hill (1964).

[34] Umpleby, S. Second-order science: logic, strategies, methods. Constructivist Foundations, 10(1), 16-23 (2014).

[35] Scholten, D. A Primer for Conant and Ashby's Good Regulator Theorem. Wayback Machine (archive.org), Retrieved November 29, 2021.
http://web.archive.org/web/20160308074837/http://www.goodregulatorproject.org/images/A_Primer_For_Conant_And_Ashby_s_Good-Regulator_Theorem.pdf

[36] Dawkins, R. The Selfish Gene. Oxford University Press (1976).

[37] Stone, R., Portegys, T., Mikhailovsky, G., and Alicea, B. (2018). Origins of the Embryo: self-organization through cybernetic regulation. BioSystems, 173, 73-82.

[38] Alicea, B. (2020). Raising the Connectome: the emergence of neuronal activity and behavior in *Caenorhabditis elegans*. Frontiers in Cellular Neuroscience, 14, 524791.

[39] Heylighen, F. Objective, Subjective and Intersubjective Selectors of Knowledge. Evolution and Cognition, 3, 63-67 (1997).

[40] Pu, C-L., Pei, W-J., and Michaelson, A. Robustness analysis of network controllability. Physica A, 391, 4420-4425 (2012).