

V. N. Alexander

AI, Stereotyping on Steroids and Alan Turing's Biological Turn

2019

<https://doi.org/10.25969/mediarep/13548>

Veröffentlichungsversion / published version

Sammelbandbeitrag / collection article

Empfohlene Zitierung / Suggested Citation:

Alexander, V. N.: AI, Stereotyping on Steroids and Alan Turing's Biological Turn. In: Andreas Sudmann (Hg.): *The democratization of artificial intelligence. Net politics in the era of learning algorithms*. Bielefeld: transcript 2019, S. 43–54. DOI: <https://doi.org/10.25969/mediarep/13548>.

Nutzungsbedingungen:

Dieser Text wird unter einer Creative Commons - Namensnennung - Nicht kommerziell - Keine Bearbeitungen 4.0/ Lizenz zur Verfügung gestellt. Nähere Auskünfte zu dieser Lizenz finden Sie hier:

<https://creativecommons.org/licenses/by-nc-nd/4.0/>

Terms of use:

This document is made available under a creative commons - Attribution - Non Commercial - No Derivatives 4.0/ License. For more information see:

<https://creativecommons.org/licenses/by-nc-nd/4.0/>

AI, Stereotyping on Steroids and Alan Turing's Biological Turn

V. N. Alexander

Introduction

Artificial Intelligence (AI) designers try to mimic human brain capabilities with “self-learning” neural networks trained by crowd-sourced selection processes or other “unsupervised” selection processes. Presumably, the logic of the input data is inscribed in the structure of the artificial network similarly to the way input shapes a human brain. Yet decades on, AI-trained chat bots and translation apps still fail to vault the low bar of the Turing Test. It is becoming clear that AI is not able to interpret signs within fluid contexts. Is biological computation qualitatively different from present-day machine computation? At the time of his death, Alan Turing was investigating how biological reaction-diffusion processes create patterns, which, in turn, constrain cellular responses and differentially trigger development. Similar mechanisms are now thought to provide the temporal and spatial constraints for ensembles of neurons allowing them to perform sensory binding and to form and recall memories. Had he lived to continue his work, Turing might have reoriented AI research to better address the challenge of creating contextual constraints, which may be what is needed to produce the unpredictable and almost miracle-like responses we call human judgement. As it is, organized statistically, current AI applied to human affairs is only good for stereotyping, which, of course, undermines the basic premise of individual democratic freedom.

Like an organism, a “smart” machine can seek an object, read a code, locate a pattern and make generalizations. Like an organism, a machine can even be designed to pursue self-preserving goals. However, we cannot say that machines currently possess humanoid intelligence. AI bots cannot understand people because they are not good with language. They do not get irony, new metaphors, metonyms, puns or jokes. Language is fundamentally allusive, not literal, as Turing once demonstrated in a letter to a friend:

Turing believes that machines think
Turing lies with men
Therefore machines do not think (1952b).

The fact that Siri cannot get this joke is not because there is not world enough and time to train the network; it is symptomatic of the essential difference between AI and Biological Intelligence (BI). Selection processes, such as those used to train AI networks, cannot evolve true intelligence. I can make such a bold statement because selection processes, such as those that neo-Darwinists have claimed have evolved animal intelligence, do not, in fact, do the job. The failure of AI chat bots is the proverbial dead canary indicating a much bigger problem, an oversimplified conception of the evolution and development of intelligent action. AI may be better than BI at mechanistic rule-bound actions like driving cars, but it is incapable of determining what humans mean or intend to do. The public should not be asked to trust AI, accepting that how it works is just a mystery. This paper aims to pull back the Wizard of AI's curtain, revealing that this allegedly superhuman intelligence is in fact just a tool, a very powerful one, that is being used by a few to control the many.

A Twenty-First Century Evolutionary Theory of Innovation

To theoretical biologists it is becoming clear that, although the natural selection of small random changes in genetic material plays a role in evolutionary processes, the outcome of such selection is the stabilization of a species and the reduction of diversity. Innovation, we now believe (See Turing, 1952a; Margulis & Sagan, 2002; Reid, 2007; Shapiro, 2011; Noble, 2016) is likely due to large, interrelated mutational events, like hybridization, gene duplication, lateral gene transfer, transposons, symbiogenesis and, importantly to this discussion, the thermodynamic self-organizing semiotic processes discovered by Turing. Such mechanisms tend to produce new ready-made tools (not randomly assembled stuff) whose functions can then be selected or not. This isn't your father's evolutionary theory.

AI is designed on the assumption that adaptive learning follows the random-change with gradual selection neo-Darwinian model of the 1950s. AI, like natural selection, makes generalizations based on a statistical definition of fitness: the most frequently reappearing patterns are selected. AI learns with repeated positive and negative reinforcement. BI can learn this way too, but it can also have eiphanies.

Algorithms Versus Semiotic Habits

To try explain why AI lacks of a sense of humor, I start by noting that while computers use *digital codes* and develop *algorithms* apart from contexts, living cells use *analog signs* and develop self-reinforcing *semiotic habits* within contexts. This paper will explore the differences between AI and BI from the perspective of Biosemiotics, a newly developing, transdisciplinary field related to the fields of Cybernetics, Complex Systems Science and Biochemistry. According to Biosemiotics, whereas a code requires a precise translation of one form into another, a sign can be translated into a variety of forms depending on the relative similarity and/or proximity of other signs and transducers. It may be that this flexibility of biological signs allows signal transduction to flow easily, to be communicated synchronously and coherently to neighboring cells, even if the signal is not quite the correct or conventional one. It may be that this difference between AI and BI can account for AI's failure to adequately translate signs in contexts.

If AI's self-learning algorithms seem to work well sometimes to predict human actions, this is because stereotypes are often true. AI is currently being used in US court systems to help determine sentences, exaggerating structural social prejudices in the data fed to the AI network. The likelihood that a criminal will re-offend is predicted by categorizing him or her as a type. The result is blacks get tougher sentences than whites with comparable data points (Angwin et al., 2016). AI is stereotyping on steroids.

AI is also being applied to the management of the public at large. According to Andrew Hallman, Deputy Director for Digital Innovation at the US Central Intelligence Agency, thanks to all the data collected on Internet users, the agency can now use Deep Learning to better "anticipate the development of social unrest and societal instability...three to five days out" (Konkel, 2016). This has me worried that a pre-crime unit is up and running. It cannot be true that sacrificing our privacy will keep us safe. Mass surveillance and Big Data collection can only serve the purpose of silencing dissent and maintaining the status quo, not anticipating actual crimes. Complex systems, like humans, tend to behave non-linearly: the ability to predict individual behavior does not improve in proportion to the increase in the amount of data that is used to make the predictions.

Although neural net designers use feedback and feedforward in an attempt to mimic non-linear biological processes, no creative mechanism is included, and the resulting intelligence resembles nameless, faceless bureaucracies that have accreted procedures for dealing with citizenry over many generations and which are not only conservative but which tend to narrow options more and more with each iteration.

The first step toward democratizing AI is to unmask this supposedly better-than-human judge. AI is no agent; it is a powerful and potentially useful tool

that should be in the hands of the many not just the few. A democratic digital society, Michael Kwet (2018) has cogently argued, requires uncompromised privacy, open-source software, as well as decentralized personal cloud systems that allow direct sharing of information. If collection of personal data is thus halted, courts and surveillance agencies will not be able to use AI to control individuals based on their memberships in or associations with various groups. I hope that my analysis of the differences between present-day AI and BI can convince the public to be more skeptical of the supposed wisdom or accuracy of AI predictions.

What Turing Knew about BI

Turing invented the most practical tool humans will likely ever wield. And yet his engineering successes were driven by an impractical desire to understand the nature of human intelligence. We follow his lead here as we try to understand how to best use computing tools in the twenty-first century.

In the 1950s, after proposing his model for a “self-learning” computer, Turing’s thinking began to take, what might be called in hindsight, a Second-Order Cybernetics or Artificial Life turn. He began conducting experiments and studies in mathematical biology. While Andrew Hodges (1983), Turing’s main biographer, saw his interest in plant and animal morphogenesis as a departure from his interest in mimicking intelligence, Jack Copeland (2004), who provides the definitive commentary on Turing’s science, points out that Turing made it clear that this new work was a further investigation of intelligent computation, even though his attention had fallen upon giraffe patterns, Fibonacci spirals and leaf generation.

Turing discovered the spontaneous processes by which unorganized systems organize themselves without interference, without external selection. C.H. Waddington (1940) had suggested to Turing that development simply falls into order somehow, flows down the path of least resistance. It was Turing who suggested that an instability, a chance pattern—not an inducer specifically designed for that function—could initiate the flow from less order to more order, from chaos to differentiation (or, to nod to Gregory Bateson, a difference that makes a difference). Ilya Prigogine, who won the Nobel Prize for related research, met with Turing in Manchester in 1952 and discussed the theory (Hodges, 1983: 587). Not until 1972 in a paper in *Physics Today* did Prigogine recognize Turing’s contribution.

Turing argued that reactions that diffused away from the point of instability result in the so-called morphogenic fields that differentially determine gene action, as described by Waddington. While biologists were interested in what this meant for embryology, what Turing was after was knowledge of how neurons might similarly differentiate and self-organize. In a 1951 letter to neurophysiologist J. Z. Young, Turing remarks, “The brain structure has to be one which can

be achieved by the genetical embryological mechanism” (qtd in Copeland, 2004: 517). Although we do not know if Turing thought development was analogous to learning, it turns out that it is. His mother Sara Turing may not have been incorrect when she opined that her son had been on the verge of an “epoch-making discovery” when he died (qtd in Hodges, 1984: 624).

AI Compared to BI

When Turing first designed his self-learning network computer, he had assumed, tacitly following neo-Darwinism, that humans make *random* guesses when they do not know a procedure for solving a problem. In 1948 in “Intelligent Machinery,” he claims, “training a human child depends largely on a system of rewards and punishments” for good and bad guesses respectively (Copeland, 2004: 425). Turing designed a chess-playing program with optional moves that could be tried at random. If a move ultimately led to failure, it would not be reinforced. Turing’s neural network was designed to start out unorganized and become organized with appropriate “interference,” mimicking education. Similar kinds of *connectionist* approaches are used today in most self-learning algorithms. The “instruction table,” as Turing called a program, is embodied in the network as it is altered by reward and punishment. Feedback can be administered by a programmer or crowd-sourced on the Internet. Although this approach is called self-learning or self-organizing, as Turing noted, such approaches still require “interference” from the outside.

The newest phase of AI is referred to as “unsupervised” learning. For example, a visual recognition network was exposed to millions of random unlabeled images on the Internet. It eventually detected some common patterns of, you guessed it, cat faces, acquiring pathways and biases in unknown ways, hundreds of levels deep (Le et al., 2013). Programmers did not tell the network what to find, but the network can now be used to find cat faces. These new unsupervised networks are not so dissimilar to Turing’s 1948 notion of a self-learning network. The main difference is the point at which the programmer interferes, during the training process to target a pre-specified pattern or after the network as detected a pattern that is of interest to the programmer. In the latter case, the unit of selection is the entire network, not individual connections within the network.

Animals most often learn in “unsupervised” situations, especially non-human animals, and are less often taught, intentionally rewarded and punished. It is the monkey see, monkey do approach. But interference, or selection, is still at work. Experiencing a procedure over and over, actually changes neuronal connections. Neurons that fire together wire together, as Donald Hebb (1949) so famously not-

ed. Learning by rote, strengthening connections over time, is statistical in nature. What happens the *most*—whether it is “right” or not—gets selected and reinforced.

Repetition is one way neurons develop connections, but not the only.

Humans (and probably other animals too) can recall details better in contexts, if they are associated with things arbitrarily similar or arbitrarily nearby. Rhymes, rhythms, tones and other poetic devices, such as metaphor and metonymy tend to aid memory, even if the connections are not repeatedly reinforced. The semi-otic habits of neuronal groups may be initiated by rare stochastic resonances (i.e., purely coincidental patterns) which lead to self-organization. A source of unpredictability in human logic and language use, this poetic type of sign action in and among neuron cells dominates subconscious processes. Subjects under hypnosis experience cross-modal perception—they begin to hear colors, for example—which indicates that when conscious perception is bypassed, the poetic workings of the subconscious are more observable (Alexander and Grimes 2017). People with synesthesia are better able to recall arbitrary facts because numbers or letters can be associated with unique colors, textures and shapes (Harvey 2013). Connections based on arbitrarily similar/proximate factors cannot be reduced to statistical description; the *number* of factors is not as relevant to outcomes as the *qualities* of the factors vis-a-vis other factors.

Formalizing Biosemiotics

Could a computer model the way nature organizes itself by linking things arbitrarily similar/proximate? Turing discovered non-linear equations that can produce computer-generated zebra stripes, invagination, metachronal waves and other natural emergent patterns. Although for years Turing’s work went unproven and many believed the similarities between the patterns generated by his equations and those found in nature were merely coincidental, Sheth (2012) and Raspopovic (2014) have finally shown that a Turing mechanism does indeed describe the process whereby fingers are created in developing embryos. It has taken some time for biologists to identify the actual chemical signals that correspond to kinds of relationships Turing imagined would have to obtain if self-organization were a mechanism for differentiation and development. Turing’s equations are complex, but suffice it to say that they involve variables for diffusion rates, reaction rates, and the ways in which these rates change. Reactions typically involve a number of morphogens, for example, X and Y react to produce Z; Z and A react to produce 2Y. The first reaction depletes Y; the second increases Y. To put it differently, the process might involve an activator that can catalyze its own production *and* that of its own inhibitor, which, in some cases, might diffuse away rapidly, setting the stage for traveling wave patterns to emerge. There is contradiction or

paradox in these processes, which are both self-creating and self-constraining, a bit like Turing's syllogism introduced at the beginning of this paper.

Let me try to elucidate the biosemiotic elements of these types of processes with a very simplified visual model with only two elements. To illustrate biological computation, I use shapes with material qualities as symbols because the binding of biological signals and receptors (sign readers) is often shape dependent. Let us say we have a molecule type \mathbb{L} and molecule type $\overline{\mathbb{T}}$. They can be turned in various directions, e.g., $\overline{\mathbb{T}}$ and \mathbb{L} . Neither \mathbb{L} s or $\overline{\mathbb{T}}$ s interact with themselves. So

$$1. \mathbb{L} + \mathbb{L} = \mathbb{L}\mathbb{L} \text{ and}$$

$$2. \overline{\mathbb{T}} + \overline{\mathbb{T}} = \overline{\mathbb{T}}\overline{\mathbb{T}}.$$

$\overline{\mathbb{T}}$ s and \mathbb{L} s together in certain orientations also result in no change: for example,

$$3. \mathbb{L} + \overline{\mathbb{T}} = \mathbb{L}\overline{\mathbb{T}} \text{ and}$$

$$4. \overline{\mathbb{T}} + \mathbb{L} = \overline{\mathbb{T}}\mathbb{L}.$$

But when $\overline{\mathbb{T}}$ s and \mathbb{L} s meet in other ways, they can interact and undergo change, e.g., a $\overline{\mathbb{T}}$ can turn into an \mathbb{L} . Transformations depend on whether the open horizontal part of the $\overline{\mathbb{T}}$ meets with the open or closed horizontal part of the \mathbb{L} . For example,

$$5. \mathbb{L} + \mathbb{L} = \mathbb{L}\mathbb{L}$$

$$6. \overline{\mathbb{T}} + \overline{\mathbb{T}} = \overline{\mathbb{T}}\overline{\mathbb{T}}$$

These are the simple local rules that limit interactions. In the contexts of [5] and [6], we may say that the \mathbb{L} is metaphorically like an \mathbb{L} , and an $\overline{\mathbb{T}}$ is metaphorically like a $\overline{\mathbb{T}}$.

Because the molecules are always in thermal motion, the way they happen to meet up is random. Statistically speaking, the production of new \mathbb{L} s or new $\overline{\mathbb{T}}$ s is equally likely. One might think that together these reaction scenarios would tend to average out, maintaining a random mixture, but, as Turing found in a similar experiment, instead, differentiation can occur. In our experiment, a clump of, say, \mathbb{L} s happens to form in one area, as they might since randomness is not perfectly non-repetitive. No new $\overline{\mathbb{T}}$ s will be produced in an \mathbb{L} clump because a $\overline{\mathbb{T}}$ is required to produce more $\overline{\mathbb{T}}$ s. Even more \mathbb{L} s may be produced at the edges of the clump when \mathbb{L} s happen to come in contact with \mathbb{L} s in the appropriate orientation.

The clump is self-increasing. No external interference is required. We may say that the material qualities of these \mathbb{L} and $\overline{\mathbb{T}}$ signs (i.e., the relative similarity and proximity of the signs) lead to the collective activity, an emergent spot pattern. \mathbb{L} s can interpret (respond to, interact with, translate) \mathbb{L} s and produce more of themselves, more \mathbb{L} s.

A soup of this mixture would yield some \mathbb{L} clumps and some $\overline{\mathbb{T}}$ clumps, floating in random mix of both \mathbb{L} s and $\overline{\mathbb{T}}$ s. If \mathbb{L} = black and $\overline{\mathbb{T}}$ = white, black and white spots will appear on a gray background, as on the coat of an Australian cattle dog. The

actual process forming animal coat patterns is much more complicated, but this serves as a simple visual illustration of spontaneous self-organization that occurs throughout nature, especially in the brain.

In “The Chemical Basis of Morphogenesis” (Turing, 1952a) and “A Diffusion Reaction Theory of Morphogenesis in Plants” (Turing & Wardlaw, 1952), Turing demonstrates that non-linear equations can describe the way patterns form spontaneously from unorganized material. He shows that genes do not need to fully *specify* the complex structure of the organism. The coding genes mainly provide the templates for making the materials, in the right order and in the right amounts, but do not contain the instructions for how to put the materials together.¹ They do not have to. The laws of physics and chemistry and the qualities of the materials (such as that of π s and \mathbb{L} s) act as the transformation rules and the constraints that help self-organize the gene-produced materials. As a computer programmer, Turing would have had great admiration for Nature’s ingenuity and economy. She did not have to physically record the procedure for development in the DNA. Instead, Nature availed herself of spontaneous self-organizing programs.

Biosemitic adaptation is possible in this system if, for example, a \mathbb{L} happens to bind with a new molecule, L, *as if it were* an \mathbb{L} . (L thus functions as a mistaken sign of \mathbb{L} .) All new signs discovered by biological systems must function to an already existing sign-reading system. *They cannot be purely random* as with neo-Darwinian theory. The outcome of an L and \mathbb{L} binding might be a new molecule that will differentially trigger cells affected by this new combination.

Waddington (1940) had provided Turing with the *epigenetic landscape* as a visual metaphor for the physical forces that guide development (or cellular responses generally) which inspired Turing’s theory. Waddington had argued that before a cell has differentiated, it is in a state of instability, like a ball sitting atop a mountain with various valley features down below. Turing realized that any slight fluctuation might push it toward one valley pathway or another from this point of instability. These ideas became known as the “catastrophe theory” of early biosemiotician René Thom (see Favareau, 2009: 337-376). Waddington guessed that alternative pathways might be “competing” autocatalytic reactions that used some of the same molecules for different processes. L might bind to \mathbb{L} and trigger one pathway or \mathbb{L} might bind with \mathbb{L} and trigger a different pathway.

The selection process of self-organization is based on the *formal* properties of the elements, qualities, not just the *number* of the elements as with statistical selection. Turing discovered the process whereby differentiating waves, morphogenetic fields, emerge spontaneously without external selection. *This* type of computation is truly self-learning.

¹ See Keller (2002) for a history of the understanding of gene action.

Emergent Brain Patterns

Throughout much of the twentieth century, brainwaves were believed to be superfluous, like the sound an engine makes without contributing to the operation of the engine. Now we must consider that these waves may be a type of emergent program for organizing the actions of neurons. In thorough reviews of the literature, Kelso et al. (1991), Uhlhaas et al. (2009) and De Assis (2015) report that many neuroscientists understand the mechanisms underlying working memory and attention in terms of emergent brain waves that synchronize distant neurons, creating virtual neuronal assemblies (De Assis, 2015; Postle, 2006). It appears that waves may provide “the ‘contexts’ for the ‘content’ carried by networks of principal cells” and “the precise temporal structure necessary for ensembles of neurons to perform specific functions, including sensory binding and memory formation” (Buzsáki & Chrobak, 1995). In addition, emergent wave patterns may also define what data gets attention, that is, consciousness (see Thompson & Varela, 2001), which, in turn, affects further sensory processing.

This signal propagation theory of learning, using self-organizing signs (not codes), may help explain how people are able to form and use fluid adaptable categories and deal with complex changing environments. Local fluctuations allow stochastic resonance (as with the L_s and L_s^l s), the similarity and proximity of possible states, which in turn allows sameness to spread, instant organization. Natural selection cannot “see” to select these local interactions (it does not need to since these interactions just flow spontaneously to the lowest energy state). What can be selected for fitness are the effects of the global patterns that emerge from the local interactions (Cf. Rocha, 1998).²

No Artificial Neural Networks or Deep Learning networks are designed to imitate the fluid interplay between self-organization and natural selection. AI designers are more committed to strictly selectionist, aka connectionist, approaches. Although learning can be accomplished this way, it produces automatons, as does standardized curriculums and relentless testing, reward and punishment.

Even with the latest celebrated update (Levis-Kraus, 2016), Google Translate is still bad with puns, jokes and poetry. Psychologists Jung-Beeman et al. (2004) suggest that insight—understanding literary themes and metaphors and getting jokes—requires synchronizing distant brain areas instantly via gamma waves. To design computers that can get allusive language, that understand people, one might need a more fluid medium for traveling waves to emerge. Atomic switch networks as per Stieg et al. (2014) seem promising; they have been used to create emergent patterns that imitate simple natural systems. Experimental chemical

2 Likewise, contrary to the selfish gene hypothesis, natural selection cannot “see” the genes *per se* only their products.

reaction-diffusion computers have been around for more than a decade (Adamatzky et al., 2005), but although they create emergent patterns, they do away with more permanent connections. Our brains seem to use both.

Maybe we will eventually use reaction-diffusion to create more humanoid AI, but we already have eight billion human computers coupled together on the Internet, like so many neurons ready to organize. The potential for spectacular evolution of knowledge is at our finger tips, if only we were in control of AI algorithms rather than controlled by them. With more information about the nature of AI compared to BI, we could make better choices with regard to how little or much we are willing to let AI think for us.

References

- Adamatzky, A., Costello, B., & Asai, T. (2005): *Reaction-diffusion computers*. Amsterdam: Elsevier.
- Alexander, V. N. & Grimes, V. A. (2017): Fluid biosemiotic mechanisms underlie subconscious habits. *Biosemiotics*, 10(3), 337-353.
- Angwin, J., Larson, J., Mattu, S., Kirchner, L. (2016): Machine Bias. *ProPublica*. Retrieved from <http://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing> (accessed July 9, 2019)
- Copeland, J. (2004): *The essential Turing*. Oxford: Oxford University Press.
- De Assis, L. (2015): Neural binding, consciousness, and mental disorder: Complexity as a common element. *Journal for Neurocognitive Research*, 57(3-4), 110-121.
- Favareau, D. (2009): *Essential readings in Biosemiotics: Anthology and commentary*. Amsterdam: Springer.
- Gilbert, C. & Sigman, M. (2007): Brain states: Top-down influences in sensory processing. *Neuron*, 54(5), 677-696.
- Harvey, J. (2013): Sensory perception: Lessons from synesthesia: Using synesthesia to inform the understanding of sensory perception. *Yale Journal of Biology and Medicine*, 86(2), 203-216.
- Hebb, D.O. (1949): *The organization of behavior*. New York: Wiley & Sons.
- Jung-Beeman, M., Bowden, E.M., Haberman, J., Frymiare, J. L., Aradmbel-Liu, S., Greenblatt, R., Reber, P. J., & Kounios, J. (2004): Neural activity when people solve verbal problems with insight. *PLOS Biology* 2(4), 500-510.
- Keller, E. F. (2002). *The century of the gene*. Cambridge: Harvard UP.
- Kelso, J., Bressler, S.L., Buchanan, S., DeGuzman, G.C., Ding, M., Fuchs, A., & Holroydet, T. (1991). A phase transition in human brain and behavior. *Physics Letters A* 169(3), 134-144.
- Konkel, F. (2016): The CIA says it can predict social unrest as early as 3 to 5 days out. *Defense One*. Retrieved from <http://www.defenseone.com/technology/2016/10/>

- cia-says-it-can-predict-social-unrest-early-3-5-days-out/132121 (accessed July 9, 2019)
- Kwet, M. (2018): Break the hold of digital colonialism, *Mail & Guardian*. 29 Jun 2018. Retrieved from <https://mg.co.za/article/2018-06-29-00-break-the-hold-of-digital-colonialism> (accessed July 9, 2019)
- Le, Q. V., Ranzato, M., Monga, R., & Ng, A. Y. (2013): Building high-level features using large scale unsupervised learning. *Proceedings of the 29th International Conference on Machine Learning*. Edinburgh.
- Levi-Kraus, G. (2016): The great A.I. awakening. *New York Times Magazine*. 14 Dec.
- Margulis L. & Sagan, D. (2002): *Acquiring genomes: A Theory of the origin of species*. New York: Basic Books.
- Noble, D. (2016): *Dance to the tune of life: Biological relativity*. Cambridge: Cambridge University Press.
- Postle, B. (2006): Memory as an emergent property of the mind and brain. *Neuroscience* 139(1), 23-38.
- Prigogine, I., Nicolis, G. & Babloyantz, A. (1972): Thermodynamics of evolution. The functional order maintained within living systems seems to defy the Second Law; nonequilibrium thermodynamics describes how such systems come to terms with entropy. *Physics Today* 25(11), 23.
- Raspopovic, J., Marcon, L., Russo, L, & Sharpe, J. (2014): Digit patterning is controlled by a Bmp-Sox9-Wnt Turing network modulated by morphogen gradients. *Science* 345(6196), 566-570.
- Reid, R. G. B. (2007): *Biological emergences: Evolution by natural experiment*. Cambridge: MIT Press.
- Rocha, L. M. (1998): Selected self-organization and the semiotics of evolutionary systems. *Evolutionary systems: Biological and epistemological perspectives on selection and self-organization*. S. Salthe, G. Van de Vijver, and M. Delpo (Eds.). Kluwer Academic Publishers; 341-358.
- Shapiro, J. (2011): *Evolution: A view from the 21st century*. FT Press.
- Sheth, R., Marcon, L., Bastida, M. F., Junco, M., Quintana, L., Dahn, R., Kmita, M., Sharpe, J. & Ros, M. (2012): Hox genes regulate digit patterning by controlling the wavelength of a Turing-type mechanism. *Science* 338(6113), 1476-1480.
- Stieg, A., Avizienis, A., Sillin, H., Aguilera, R, Shieh, H.-H., Marin-Olmos, C., Sandouk, E., Aono, M. Gimzewski, J. (2014): Self-organization and emergence of dynamical structures in neuromorphic atomic switch networks. *Memristor Networks*. A. Adamatsky and L. Chua (Eds.) New York and London: Springer; 173-209.
- Thompson, D. W. (1942): *On Growth and Form*. Cambridge: Cambridge University Press.
- Thompson, E., Varela, F. (2001): Radical embodiment: Neural dynamics and consciousness. *Trends in Cognitive Sciences* 5(10), 418-425.

- Turing, A. (1952a): The chemical basis of morphogenesis. *Philosophical Transactions of the Royal Society of London. Series B* 237(641), 37-72.
- Turing, A. (1952b): Letter to Norman Routledge. Retrieved from <http://www.turingarchive.org/viewer/?id=167&title=1a> (accessed July 9, 2019)
- Turing, A., Wardlaw, C. (1952): A diffusion reaction theory of morphogenesis in plants. Retrieved from <http://www.turingarchive.org/browse.php/C/7> (accessed July 9, 2019)
- Uhlhaas, P., Pipa, G., Lima, B., Melloni, L., Neuenschwander, S., Nikolic, D & Singer, W. (2009): Neural synchrony in cortical networks: history, concept and current status. *Frontiers in Integrative Neuroscience* 3(17), 1-19.
- Waddington, C. (1940): *Organisers and Genes*. Cambridge: Cambridge University Press.