Ontogenesis of Agency in Machines: A Multidisciplinary Review

Byoung-Tak Zhang
Biointelligence Laboratory & Institute for Cognitive Science
School of Computer Science and Engineering & Brain Science and Cognitive Science Programs
btzhang@bi.snu.ac.kr

Abstract

How does agency arise in biocognitive systems? Can the emergence of agency be explained by physical laws? We review these and related questions in philosophy, psychology, biology, and physics. Based on the review we ask the questions: i) Can machines have agency? and, if so, ii) How can we build machines with agency? We examine existing work in artificial intelligence and machine learning on self-motivated, self-teaching, and self-developing systems with respect to the ontogenesis ("coming into being") of agency in computational systems. The impact of these "autogenic" systems or machine agency on science, technology, and humanity will be discussed.

Introduction

During the last decades artificial intelligence has achieved enormous successes both in technology and industrial applications. Breakthroughs in speech recognition, image classification, and face recognition, for example, have been obtained by advanced machine learning technology (e.g., Graves et al., 2013). Though significant, much of these successes are based on the brute-force memory and computing power combined with big data. To really achieve human-level artificial intelligence it is necessary for the machines to have the capability of continually searching and collecting the right examples for themselves. Ultimately a fully-autonomous learning system that has a purpose and intention, i.e. agency, to improve itself will be desirable.

Here we review the discourse on the emergence of agency, i.e. acting on its own behalf, in biological and cognitive systems from a multiplicity of disciplines. In philosophy, Aristotle identified the end-directedness of final cause or agency. Later, however, Descartes chopped away the agency from the domain of natural science, separating psychology from physics. Kant reintroduced agency in terms of teleology, i.e. goal-directedness, as a distinguishing property of biological organisms. In the 20th century the dominating reductionism again excluded agency from natural science and it became almost a taboo to talk about teleology in biology.

However, in the 1970's and 80's, scientists began to readdress teleology. Maturana and Varela's work (1980) on autopoietic (i.e. self-forming) organization was influential in bringing about this change. This work was also instrumental to the modern turn of cognitive science into embodied cognition. Evolutionary developmental systems theorists (evo-devo) also contributed to the emergentist movement on life and mind. Contrary to the reductionistic and mechanistic approaches, these holistic and systems science approaches redefined the causality as cyclic loops. In particular, developmentalists emphasize self-organization as well as natural selection as the driving force for evolution and emergence of meaning, purposiveness, function, value, and thus agency (Salthe, 1993; Deacon, 2012).

A big question is whether or not biological systems, and thus agency, can be explained by the physical laws as we know them today. Schroedinger (1944) discussed how the living process and the second law of thermodynamics are related through entropy production. Prigogine discovered that biological systems work far from equilibrium and thus living is "becoming", rather than "being" (Prigogine, 1980). Kauffman emphasized the role of self-organization, in addition to selection, as an important source of order in complex adaptive systems such as biological organisms (Kauffman, 2000). The origin-of-life and complex systems research has shed light on our understanding of how life, and thus mind and agency, emerge from matter.

Mathematical biologists and systems scientists have also contributed to the understanding of autonomy in organisms. Rosen (1985), for example, presents a mathematical model of anticipation for the emergence of function, meaning, and purpose. Ecological psychologists have studied the perception-action cycle and its feedback mechanism for the goal-directed behavior in organism-environment interactions (Turvey, 2013).

The rest of the paper is organized as follows. We review in more detail the problem of agency in philosophy, psychology, biology, chemistry, and physics. Our discussion is focused on whether life and mind can emerge from matter, i.e. is there a need for additional physical laws to explain the living and mental systems? Much of the review discusses the possibility and the potential processes for the emergence of biological from physical and of psychological from biological systems. In the final section, we ask the following questions: i) Can machines have agency? and, if so, ii) How can we endow machines with agency? We review some existing work in this direction of research especially from lifelong machine learning, and hint on the future of human-level artificial intelligence.

Agency in Philosophy

All biological organisms "act on their own behalf." Agency and purposiveness distinguish living systems from the non-living and, thus, cannot be ignored in characterizing life and mind. Avoiding purposiveness from natural science is one thing and avoiding the phenomenon of purposiveness is another.

In ancient philosophy, Aristotle distinguished four distinct modes of causality of the change of objects: material, formal, efficient, and final causes. Material cause is what determines the structural stability of the raw material. Formal cause is the form of the object being changed. Efficient cause is the modification of the material to create this form. Final cause is the aim of the process. The final cause refers to the purpose and goal, and thus agency, of the system. However, at the beginning of the Renaissance, several thinkers, like Francis Bacon and Rene Descartes, progressively downplayed the role of final causality. Thus, the Cartesian and Newtonian world of mechanism mainly dealt with the first three causes while ignoring the fourth, final cause.

Kant states that "an organized natural product is one in which every part is reciprocally both end and means" (Deacon, 2012). Kant argued that "An organized being is then not a mere machine, for that has merely motive power, but it possesses in itself formative power of a self-propagating kind which it communicates to its materials though they have it not of themselves". Implicit in Kant's abstract characterization of "formative power" is the fact that organisms are organized so as to resist dissolution by replacing and repairing their degraded components and structural characteristics, and eventually replacing themselves altogether by reproduction. More importantly, he emphasizes that this is a reciprocal process. No component process is prior to any other. Kant's characterization is prescient. He puzzles over the question of whether there is something like intrinsic teleology in organisms. Kant concludes that this formative reciprocity constitutes what he calls "intrinsic finality".

A closely related topic to agency is intention. Intention is a specific purpose in performing an action. It is the end or goal that is aimed at. Intentional behavior can also be just thoughtful and deliberate goal-directedness. Psychologists have long studied conscious experience. Like intention, conscious experience is intrinsically consequence-oriented and normative. Dennett distinguishes three levels of abstraction (Dennett, 1985). The most concrete, physical level is concerned with mass, energy, velocity, and chemical composition. The more abstract, biological level is concerned with purpose, function and design. The most abstract is the intentional stance, which is the domain of software and minds, and concerned with belief, thinking, and purpose.

Developmental Systems

What are the fundamental properties that distinguish the living from non-living systems? What mechanisms allow for cognition to emerge from biology? Can the biology be explained by the first principles of physical laws? Here we review these questions in natural sciences. These are all involved with self-organization and developmental processes of organisms.

Autopoietic Systems

Maturana and Varela (1980) introduced the notion of "autopoiesis" (i.e. self-production or self-formation) to explain the nature of living systems. Autopoietic organization is a system capable of reproducing and maintaining itself. A standard example of an autopoietic system is the biological cell, consisting of various biochemical components and organized into bounded structures such as the cell membrane and cytoskeleton. These structures, based on an external flow of molecules and energy, produce the components which, in turn, continue to maintain the organized bounded structure that gives rise to these components. In contrast to self-organization, an autopoietic system is autonomous and operationally closed, in the sense that there are sufficient processes within it to maintain it as a whole. Autopoietic systems are "structurally coupled" with their medium, embedded in a dynamic of changes that can be recalled as sensory-motor coupling. This continuous dynamic is considered as a rudimentary form of knowledge or cognition and can be observed throughout life-forms.

Developmental Systems Theory

Evolutionary developmental systems or developmental systems theory is a theoretical perspective on biological development, heredity, and evolution (e.g., Callebaut et al., 2005; Salthe, 1993; Tommasi et al., 2009; Vijver et al., 1998). It is fundamentally opposed to reductionism and emphasizes the shared contributions of genes, environment, and epigenetic factors on developmental processes. The developmental systems theory argues the inadequacy of modern evolutionary views on the roles of genes and natural selection as the principle explanation of living structures.

Developmental systems theory focuses on the observation that all biological processes operate by continually assembling new structures. Each such structure transcends the structures from which it arose and has its own systematic characteristics, information, functions and laws. Conversely, each such structure is ultimately irreducible to any lower (or higher) level of structure, and can be described and explained only on its own terms. Furthermore, the major processes through which life as a whole operates, including evolution, heredity and the development of particular organisms, can only be accounted for by incorporating many more layers of structure and process than the conventional concepts of 'gene' and 'environment' normally allow for. Not only are major elements of the environment built and inherited materially as any gene but active modifications to the environment by the organism (for example, a termite mound or a beaver's dam) demonstrably become major environmental factors to which future adaptation is addressed. This inheritance may take many forms and operate on many scales, with a multiplicity of systems of inheritance complementing the genes. According to the developmental theory, 'information' depends wholly on the content and context out of which that information arises, within which it is translated and to which it is applied.

Self-organizing Molecular Systems

How did life arise from non-living material? Much of the principles of biological systems and their origin have been studied on chemical dynamical systems. Autocatalysis and self-assembly are self-organizing molecular processes that are capable of occurring spontaneously in a wide variety of conditions (Kauffman, 2000; Deacon, 2012).

Autocatalysis and Self-assembly

Autocatalysis is a catalytic chemical reaction in which a small set of catalysis each enhances the production of another member of the set, so that ultimately all members of the set are produced. It is a self-amplifying chemical process occurring in non-equilibrium. Because of the widespread presence of chemical circles in living systems, autocatalysis offers a starting point for the origins of life, and it has been promoted by many, as the chemical analogue of self-replication (Eigen, 1979; Kauffman, 2000). The complex network analysis shows that complex reciprocally catalytic networks of molecular interactions characterize the metabolic processes of living cells.

Self-assembly is a form-generating process at the molecular level. It is responsible for much of the microstructure of eukaryotic cells. Self-assembly of macromolecular structures is a special case of crystallization. Like the crystal lattice formation, the growth of multi-unit macromolecular structures is an expression of the intrinsic geometry of component molecules, the collective symmetries these offer in aggregate, and the lower energy state of the crystallized forms. Examples of self-assembling macromolecular structures in cells include laminar and tubular structures, lipid bilayers, and protein matrices.

Autocatakinetic Systems

Autocatakinetic systems are spontaneously ordered systems or flow structures that by pulling resources into themselves maintain their identities through the flux, or motion of their components (Swenson 1998). These ideas are related with the open systems of von Bertalanffy (1950), the dissipative structures of Prigogine (1978), and the autopoietic systems of Maturana and Varela (1980). More formally, an autocatakinetic system is defined as one that maintains its "self" as an entity constituted by, and empirically traceable to, a set of non-linear (circularly causal) relations through the dissipative or breakdown of field (environmental) potentials (or resources) in the continuous coordinated motion of its components.

Physical and Mathematical Theories

Can the purposiveness emerge spontaneously by natural laws from physical and molecular processes? Significant insights concerning the nature of life came from the studies of thermodynamic processes that are far from equilibrium.

End-Directedness and Thermodynamics

The second law of thermodynamics says that the world evolves toward maximum entropy. However, this end-directedness of the second law appears to run directly opposite to the active, end-directedness of living things and the intentional dynamics. Von Bertalanffy (1950) partially solved this dilemma by showing that spontaneous order can appear in open systems (systems with energy flows running through them) by their ability to bring their order by dissipating potentials in the environments. Schroedinger (1944) points out that living things exist far from equilibrium by feeding on negative entropy in their environments. These ideas were further developed by Prigogine (1980). Schroedinger's point was that as long as living things produce entropy (or minimize potentials) at a sufficient rate to compensate for their own internal ordering (their ordered persistence away from equilibrium) then the balance equation of the second law would not be violated. Order can arise spontaneously, and living things are permitted to exist as long as they "pay their entropy debt" (Swenson, 1998).

Kauffman (2000) proposes a tentative fourth law of thermodynamics, in which the workspace of the biosphere expands, on average, as fast as it can in this coconstructing biosphere. Based on the concept of the autokatakinetic systems, Swanson proposed the law of maximum entropy production (LMEP) as a basis for end-directed ordering determined by meaning (Swenson, 1998). It says that potentials are minimized at the fastest rate given the constraints, and, following the balance equation of the second law of thermodynamics, autocatakinetic or spontaneously ordered systems work spontaneously to increase the rate. As a consequence, the world acts opportunistically to produce as much order as it can. The epistemic dimension, the urgency towards existence characterizing the intentional dynamics of living things is thus not only commensurable with universal first principles, but a direct manifestation of them. Thus, the law is also known as the fourth law of thermodynamics (Morel and Fleck, 2006). This view provides a principled basis for uniting living things and their environments, knower and known, or self and other as a reciprocal part of an active world acting back on itself in its own becoming.

The classical statement of the second law says that entropy will be maximized, or potentials minimized, but it does not ask or answer the question of which out of available paths a system will take to accomplish this end. The answer to the question is that the system will select the path or assembly of paths out of otherwise available paths that minimize the potential or maximize the entropy at the fastest rate given the constraints. The statement of the law of maximum entropy production is the physical selection principle that provides the nomological explanation for why the world is in the order production business.

Anticipatory Systems and Relational Biology

Mathematical models have been developed for end-directedness of biological systems. Recognizing the inadequacy of reductionism to explain and describe the behavior of biological systems, Rosen (1985), for example, proposes the notion of anticipatory systems. An anticipatory system is a natural system that contains an internal predictive model of itself and of its environment, which allows it to change state at an instant in accord with the model's predictions pertaining to a later instant.

Rosen proposes the (M, R)-system as a formal model of biological organism, where M stands for metabolism and R for repair. Building on Rashevsky's work on n-ary relations in organismic sets, Rosen further develops relational biology and elaborates on the concept of organization. Organizations have more than purely structural or material aspects. Rosen's abstract relational biology approach focuses on a definition of living organisms in terms of their internal organization as open systems that cannot be reduced to their interacting components because of the multiple relations between metabolic, replication and repair components that govern the organism's complex biodynamics.

Emergent Teleodynamics

For the emergence of agency, or mind in general, three levels of dynamics can be distinguished: homeodynamics, morphodyanmics, and teleodynamics (Deacon, 2012). Homeodynamics is the spontaneous inorganic processes that generate macroscopic order. Morphodyanmics is the chemical self-organizing processes that create and maintain structural and dynamical regularity. Teleodynamics is the organizing causal processes that are end-directed and consequence-organized. Deacon argues that life requires the additional emergent transition of teleodynamics which is superveninent on morphodynamics and therefore also on thermodynamics.

Autocatalysis and self-assembly are the primitive morphodynamic processes that can build autonomous systems or autogens (Deacon, 2012). Self-assembly provides the conditions for sustaining autocatalysis, and autocatalysis complements self-assembly by continual production of identical molecules. Through their reciprocal complementarity, these two self-organizing processes can be spontaneously linked. It also creates the potential for self-reconstitution. Though each component process is self-undermining in isolation and co-dependent, together they are reciprocally self-limiting, so that their self-undermining features are reciprocally counteracted. Thus, an autogenic system can establish its capacity to re-form before exhausting substrates. This self-reconstitution property provides a potential mechanism for self-reproduction. That is, a disrupted autogen will be as likely to produce two identical autogens as one.

Once autogens are capable of self-replication, they are also potential progenitors of autogen lineage. An autogen will increase in numbers so long as there are sufficient substrate molecules in the surrounding environment. There will also be competition among many lineages as well as variation. These provide the conditions for natural selection. Their self-reconstituting properties in favorable environments spontaneously bring into being the systemic conditions that are sufficient to initiate a persistent form of natural selection. Autogen lineage competition for resources will tend to lead to the evolution of variant lineages differentially "fitted" to their local environments.

The notion of constraint plays an important role in autogenic systems. A constraint is "that which binds together" and the cause of that, restriction or reduction of variety. The morphodynamic processes maximize entropy flow, i.e. the dissipation of constraints, but as fast as entropy increase diffuses the introduced distributional asymmetry, external perturbation reestablishes it. The result is a constant throughput of energy and a constant rate of entropy generation. In dissipative systems, those processes that dissipate constraints more slowly and less efficiently will tend to be spontaneously supplanted by those that do so more rapidly and more efficiently. The buildup of local constraints creates conditions where a fraction of their potential to do work is diverted into the generation of global constraints that progressively increase global dissipation rates. In this way, continual dynamical perturbation causes local impediments to dissipation to become self-eliminating. Self-organization is a process that works in the service of the second law, i.e. the maximum entropy production (Swenson, 1998; Kauffman, 2000).

Thus, autogens conserve constraints and, by doing that, preserve information about past adaptive organization. The retained foundation of reproduced constraints is effectively the precursor to genetic information. Whether it is embodied in specific information-bearing molecules (as in DNA) or merely in the molecular interaction constraints of a simple autogenic process, information is ultimately constituted by preserved constraints. By transforming the formative dynamics of self-organizing processes into structures that transiently resist degradation, living organisms and autogens provide the foothold for additional linked forms of constraint propagation to develop.

Teleodynamic organization emerges due to reciprocally organized morphodynamic processes, and results in constraint stabilization rather than constraint amplification, and entropy rachet rather than entropy production. The emergence of teleodynamic processes includes more than merely consequence-organized phenomena. Teleodynamic organizations are robust to perturbation. They maintain their integrity. The physically dissociated components retain their systemic identity sufficiently to be able to reassociate into identically organized unit structures. They have functions, i.e. each of the components is present for the sake of the other. They have values. Different features of the surroundings of the organizations are beneficial or harmful.

Autogenic processes have the minimal precursors of function, adaptation, teleology, valuation, and the dim anticipation of information about the environment and a self with respect to which all this matters. The emergence of these attributes can be understood without the need to attribute them to mysterious or weird forms of causality. This demonstrates that real teleological and intentional phenomena can emerge from physical and chemical processes previously devoid of these properties. However, there can be no simple one-to-one mapping of teleodynamic relationships to mechanistic relationships. The link between these should be bridged by an intervening level of morphodynamic processes. The living dynamics and cognitive processes (i.e., mental processes) cannot be directly mapped to simpler physical processes (i.e. mechanistic processes). The supervenient relationship between them is necessary.

Cognitive Neuroscience of Agency

Agency or “free will” in action has attracted much attention in neuroscience and cognitive science (Libet, 1999; Friston, 2009; Goschke, 2013). Here we review cognitive neurosciences of how intentions and volition emerge in goal-oriented actions, based on the latest article of Herwig, Beisert, and Prinz (2013).

An action has two elements. The first is ongoing behavioral activity, i.e. body movement. The second is the orientation of these movements toward prospective goal state, i.e. toward the effect created by these movements. The term goal is used both as a descriptive and as an explanatory term. At a descriptive level, goals characterize the behavior of systems whose actions are apparently directed toward the achievement of certain goals-in-the-world. In this respect, the behavior of the organism can be considered to be goal directed. At an explanatory level, goals do not tell us anything about whether they actually play a functional role in action.

Concerning goals at an explanatory level, there are two classical answers to this question. One answer roots goal-directed behavior in procedural routines that have evolved through phylogenetic or ontogenetic learning. It is assumed that these procedural routines do their work without the active involvement of goals. The second answer roots goal-directed behavior in the activity of goal representations, i.e. goals-in-the-mind, arising from motivational processes. This idea is not only deeply embedded in our common sense psychology but is also a part of scientific theories of motivation and volition that consider internal representations of goal states as an essential part of the intention to act. Thus, goals have two faces: actions are driven by and directed at goals. That is, goals are used both for describing and for explaining action.

Actions build on two kinds of antecedent conditions, external and internal. While external circumstances are represented through perceptual functions, internal circumstances are represented through intentional functions. Intentional functions stand for the driving forces that make things happen. While perception is required to select an appropriate action, intention is required to initiate and realize it. The “cool” approaches to action emphasize the role of perception and cognition, while the “hot” approaches to action emphasizes the role of intention and volition.

Hot sources of action have been relatively neglected because of their difficulty in both method and theory. Regarding the methods, hot internal contributions to action are less easy to access than cool internal contributions. Regarding the theories, intentional representations and functions refer, by definition, to states of affairs that are currently not given. Thus, intentional contributions to action are much more elusive than perceptual contributions. Future research needs to develop a single, coherent framework that combines the operation of cool and hot factors.

The ideomotor theory posits two basic principles, one for learning and another for control. Ideomotor learning claims that (representations of) actions become associated with (representations of) action effects or outcomes they lead to. Ideomotor control claims that, once acquired, these associations can also be used in the reverse direction, linking representations of action outcomes with representations of actions leading to them. These reverse links offer an explanation for goal-directed, voluntary action. The ideomotor theory of voluntary action claims that representations of intended action effects, i.e. goals, have the power to elicit actions that are suited to realize them, i.e. of which previous ideomotor learning has shown that they may lead to these effects. While this addresses the mechanics of learning and control, it fails to address the dynamics of volition.

When and how are intentional representations formed and in what way do they differ from nonintentional kinds of action representations? The extended learning claims that action outcome learning includes a dynamic component of outcome evaluation – an automatic assessment of action outcomes in terms of the actor's current needs and desires. Action outcomes turn from neutral to more or less desirable or undesirable events. Regarding control, we may envisage a sequence of events in the reverse order. Here we may conceive intention formation as a process that starts from given dynamic states, i.e. needs and desires, and then determines action goals that are desirable or undesirable in terms of these states. According to this view, the crucial functional feature that distinguishes intentions from other kinds of action-related representations (like percepts, thoughts, or anticipations) drives from their origin in dynamic states. Intentions are, from the outset, charge and energized by motivational drivers like needs and desires. It is these drivers that make the difference. They furnish intentions with the power to make things happen by moving from representation and prediction to execution and production.

One of the most important implications of this is that it construes close functional links between perception and volition. It also says that perceptual and intentional underpinnings of action are both acquired in the same learning processes, and likewise they are both combined in the formation of intentions for control. As a result, the dynamics of volition use the same kinds of representational resources as the mechanics of perception and cognition. While intentions and action goals originate from dynamic sources, they are created and maintained in the same representational domain that takes care of the mechanics of action and event representations. Action, event, and goal representations are all made of the same stuff: they stand for events that happen, have happened, may happen, or ought to happen.

Human-Level Artificial Intelligence

So far we have seen how agency and goal-directedness can emerge in nature, i.e. psychological, biological, chemical, and physical systems. A big question from artificial intelligence point of view is: Can we build a machine with agency? or How do we endow machines with agency? Before we discuss this ambition, a historical context of human-level AI is in order.

In the Zeitgeist of Turing's and many other contemporary's dream of thinking machines, AI started in the 1950's. In the 1960's AI has focused on general problem solving based on heuristic search, logic, and symbol manipulation. Then, in 1970-80's rule-based expert systems and knowledge-based systems have been developed and deployed in the industry. When the Internet and Web age emerged in the 1990's, intelligent agents were widely used for internet services, such as information retrieval, e-commerce, and recommendation. In the current age of big data and smart mobile services, machine learning technologies are the driving force for the innovation of internet companies. Most recently, many global internet companies, including Google, Facebook, Amazon, Microsoft, Samsung, and Baidu, have founded AI and machine learning research centers on/near the university campuses.

However, AI technologies have so far been focused on specialized problems, while general intelligence of human level is still far from being achieved. In 2006 an issue of AI Magazine deals with the human-level AI as a special topic for its 50th anniversary. In this volume, many researchers, including pioneers like John McCarthy, Marvin Minsky, and Nils Nilsson, reflected on the last 50 years and look forward to the next 50 years.

Lifelong Machine Learning

To achieve human-level intelligence, human-level machine learning is a prerequisite. Lifelong learning has emerged as a new paradigm for studying human-level intelligence (Carlson, 2010; Zhang, 2008). In contrast to traditional machine learning algorithms that learn passively from given examples, a lifelong learning environment requires several different capabilities of the learning agent. These include active selection of examples from the environments. Active learning deals with the active choice of learning samples by the learner (Freund et al., 1993; Zhang, 1993). When the learner is equipped with sensors and actuators, it is also possible to actively generate novel examples. These creative learning modes, such as self-motivated learning and self-teaching, are especially useful and have the potential for developing "purposive" or "intentional" agents.

Another aspect of lifelong learning is the self-organization capability of the learner's structures or models, such as self-development learning (Zhang, 1993; Zhang et al., 2003). Since the environment can change over time in lifelong learning, the plasticity of the learning model is necessary. This is contrasted to the conventional machine learning paradigms, which focuses on parametric estimation rather than structural learning.

Sequential aspects of lifelong learning is another arising issue. While supervised and unsupervised learning paradigms deal with a static set of given data samples, lifelong learning has to handle a long stream of incoming or sequentially-observed data. The temporal dependency of the data is not to be ignored and proves essential. This setting requires the learning agents to estimate the temporal dependencies of sequential data. Predictive learning, e.g. learning from prediction errors in the perception-action cycle, is a promising paradigm for dealing with this class of problems (Ay et al., 2008; Thelen and Smith, 1994). The agency or intention of the agent will play a great role since a purposive or goal-directed learner can much more efficiently sense the environment than a passive agent.

Some recent work addressed these issues (Zhang, 2013; and references therein). In an interactive learning setting, the learner interacts with the environment between consecutive observations. The problem is to choose a model and an action policy, which are optimal in that they maximize the learner's ability to predict the world, while being minimally complex. The decision function, or action policy, is given by the conditional probability distribution $P(a_t \mid h_t)$. Let the model summarize historical information via the probability map $P(s_t \mid h_t)$. The interactive learning problem is solved by maximizing $I(\{S_t, A_t\}; S_{t+1})$ over $P(s_t \mid h_t)$ and $P(a_t \mid h_t)$, under constraints that select for the simplest possible model and the most efficient policy, respectively, in terms of smallest complexity measured by the coding rate. Less complex models and policies result in less predictive power. This trade-off can be implemented using Lagrange multipliers, λ and μ. Thus, the optimization problem for interactive learning is given by

$$\max_{P(s_t \mid h_t), P(a_t \mid h_t)} \left\{ I(\{S_t, A_t\}; S_{t+1}) - \lambda I(S_t; H_t) - \mu I(A_t; H_t) \right\}$$

where H_t denotes the history up to time *t*.

Another approach is to consider how much influence an agent has on its environment. This is known as empowerment and an information-theoretic generalization of joint controllability and observability of the environment by the agent (Jung et al., 2012). It is defined as the Shannon channel capacity between A_t, the choice of an action sequence, and S_{t+1}, the resulting successor state:

$$\begin{aligned} C(s_t) &= \max\nolimits_{P(a)} I(S_{t+1}, A_t \mid s_t) \\ &= \max\nolimits_{P(a)} \left\{ H(S_{t+1} \mid s_t) - H(S_{t+1} \mid A_t, s_t) \right\} \end{aligned}$$

where $H(a|b)$ is the conditional entropy of *a* given *b*. The empowerment measures to what extent an agent can influence the environment by its actions over time.

Lifelong learning can be naturally formulated in the reinforcement learning framework. The goal of reinforcement learning is to maximize the expected value for the cumulated reward. The reward function is defined as $R(s_{t+1} \mid s_t, a_t)$. This value is obtained by averaging over the transition probabilities $T(s_{t+1} \mid s_t, a_t)$ and the policy $\pi(a_t \mid s_t)$. Given a starting state *s* and a policy π, the value $V^{\pi}(s_t)$ of the state s_t following policy π can be expressed via the recursive Bellman equation, or, alternatively, the value function on state-action pairs:

$$Q^{\pi}(s_t, a_t) = \sum_{s_{t+1} \in S} T(s_{t+1} \mid s_t, a_t) \left[R(s_{t+1} \mid s_t, a_t) + V^{\pi}(s_{t+1}) \right]$$

which is the utility function attained if, in state s_t, the agent execute action a_t, and after that begins to follow π.

To deal with multiple optimal policies, Tishby and Polani (2010) and Polani (2011) propose to introduce information cost term in policy learning. If we wish the expected reward $\mathbf{E}[V(S)]$ to be sufficiently large, the information cost for such informationally parsimonious policy will be generally lower. For a given utility level, we can use the Lagrangian formalism to formulate the unconstrained minimization problem

$$\min_{\pi} \left\{ I^{\pi}(S_t; A_t) - \beta E[Q^{\pi}(S_t, A_t)] \right\}$$

where $I^{\pi}(S_t; A_t)$ measures the decision cost incurred by the agent:

$$I^{\pi}(S_t; A_t) = \sum_{s_t} P(s_t) \sum_{a_t} \pi(a_t \mid s_t) \log \frac{\pi(a_t \mid s_t)}{P(a_t)}$$

where $P(a_t) = \sum_{s_{t+1}} \pi(a_t \mid s_{t+1}) P(s_{t+1})$. The term $I^{\pi}(S_t; A_t)$ denotes the information that the action A_t carries about the state S_t under policy π.

The objective function consisting of the value function and the information cost can balance the expected return with minimum cost. However, this lacks any notion of interestingness or curiosity (Ay et al., 2008; Zhang, 2013). The objective function can be extended by the predictive power. Using Lagrange multipliers, we can formulate the lifelong learning as an optimization problem:

$$\arg\max_{q}\left\{I_q^{\pi}\left(\{S_t,A_t\};S_{t+1}\right)+\alpha V_t^{\pi}(q)-\lambda I(S_t;A_t)\right\}$$

where $q(a_t \mid s_t)$ is the action policy to be approximated. The ability to predict improves the performance of a learner across a large variety of specific behaviors. The above objective function embodying the curiosity terms as well as the value and information cost terms can thus be an ideal guideline for a lifelong autonomous learning.

Discussion: Machine Agency and Humanity

So far we have examined and discussed how agency can emerge in biological systems from different viewpoints. The developmental systems theory emphasizes the importance of non-reductionistic, systemic approaches to understanding the biological organisms. It argues that life as a whole can only be accounted for by incorporating many more layers of structures and processes than the conventional concepts of 'gene' and 'environment' normally allow for. Autopoietic systems suggest a model of cognitive functions in biological organization. Autopoietic systems are structurally coupled with their medium, embedded in a dynamic of changes that can be recalled as sensory-motor coupling.

Several proposals provide the possibility of emergence of agency from matter, i.e. by the physical and chemical laws as we know them today. Based on the autocatalysis of complex molecular networks, Kauffman describes how autonomous agents can come into being through self-organization and selection processes. Deacon suggests how purpose and agency may arise through three levels of emergent dynamic transitions, i.e. the thermodynamic (homeodynamic), self-organizational (morphodynamic), and semiotic (teleodynamic) processes.

Mathematical models have been also proposed to explain the nature of biological organisms and their end-directed behavior. The anticipatory system of Rosen, for example, provides an insight of purposive systems. It explains how an internal predictive model of itself and of its environment allows for the future to influence the present behavior. The law of maximum entropy production a.k.a. the fourth law of thermodynamics explains the end-directed behavior of biological organisms can be derived from thermodynamic principles in nature.

Most of these proposals are based on non-reductionistic, self-organizing, and emergent systems, such as, those developed from the communities of complex adaptive systems, general systems science, and developmental systems theories. We have also seen that these "autogenic" system principles and mechanisms can be combined to build human-level AI systems that continually and actively learn in an ever-changing environment over lifetime.

Scientifically, we still have to empirically prove that machines can have agency or "free will" as we observe it in humans. If machine agency becomes a true possibility and reality, it will be a double-edged sword. Technologically, the machines can grow and refine their knowledge without any bound. Therefore, a true human-level machine intelligence or even a superhuman intelligence or hyperintelligence can be a reality. However, superhuman AI might threat humanity and society. It will be necessary for humans to have control over machine agency. It remains to answer the question: Once a machine possesses his/her own agency, how can we control it?

Acknowledgments

This work was supported by the National Research Foundation (NRF), the US Air Force Office of Scientific Research (AFOSR), and Samsung Future Technology Foundation (FTF).

References

AI Magazine. 2006. *Special Issue on Achieving Human-Level AI through Integrated Systems and Research* 27(2), July 2006.

Ay, N., Bertschinger, N., Der, R., Guetter, F., and Olbrich, E. 2008. Predictive Information and Explorative Behavior in Autonomous Robots. *European Physical Journal B* 63:329-339.

Bertalanffy, L. von. 1950. The Theory of Open Systems in Physics and Biology. *Science* 3: 23-9.

Callebaut, C. and Rasskin-Gutman, D. (eds.), 2005. *Modularity: Understanding the Development and Evolution of Natural Complex Systems*. MIT Press.

Carlson, A., Betteridge, J., Kisiel, B., Settles, B., Estevam R. Hruschka Jr., E. R., and Mitchell, T. M. 2010. Toward an Architecture for Never-Ending Language Learning. In *Proceedings of AAAI 2010*. AAAI Press.

Deacon, T. 2012. *Incomplete Nature: How Mind Emerged from Matter*. Norton.

Dennett, D. 1987. *The Intentional Stance*. MIT Press.

Eigen, M. and Schuster, P. 1979. *The Hypercycle – A Principle of Natural Self-Organization*. Springer.

Freund, Y., Seung, H. S., Shamir, E., and Tishby, N. 1993. Information, Prediction, and Query by Committee. In S. Hanson et al. (ed.) *Advances in Neural Information Processing Systems 5*, Morgan Kaufmann.

Friston, K. 2009. The Free-energy Principle: A Unified Brain Theory? *Nature Reviews Neuroscience* 11:127-138.

Goschke, T. 2013. Volition in Action: Intentions, Control Dilemmas, and the Dynamic Regulation of Cognitive Control. In W. Prinz, M. Beisert, and A. Herwig (eds.) *Action Science: Foundations of an Emerging Discipline*, Chapter 16, pp. 409-434. MIT Press.

Graves, A., Mohamed, A., and Hinton, G. E. 2013. Speech Recognition with Deep Recurrent Neural Networks. In *IEEE Int. Conf. on Acoustic Speech and Signal Processing* (ICASSP 2013), Vancouver.

Herwig, A., Beisert, M., and Prinz, W. 2013. Action Science Emerging: Introduction and Leitmotifs. In W. Prinz, M. Beisert, and A. Herwig (eds.) *Action Science: Foundations of an Emerging Discipline*, Chapter 1, pp. 1-33, MIT Press.

Jung, T., Polani, D., and Stone, P. 2011. Empowerment for Continuous Agent-Environment Systems. *Adaptive Behavior* 19(1):16-39.

Kauffman, S. 2000. *Investigations*. Oxford University Press.

Libet, B. 1999. Do We Have Free Will? *Journal of Consciousness Studies* 6(8–9):47–57.

Maturana, H. R. and Varela, F. J. 1980. *Autopoiesis and Cognition: The Realization of the Living*. Springer.

Morel, R. E. and Fleck, G. 2006. A Fourth Law of Thermodynamics. *Chemistry* 15(4).

Thelen, E. and Smith, L. B. 1994. *A Dynamic Systems Approach to the Development of Cognition and Action*. MIT Press.

Prigogine, I. 1980. *From Being To Becoming*. Freeman.

Rosen, R. 1985. *Anticipatory Systems: Philosophical, Mathematical and Methodological Foundations*. Pergamon Press.

Salthe, S. N. 1993. *Development and Evolution: Complexity and Change in Biology*. MIT Press.

Schroedinger, E. 1944. *What Is Life?* McMillan.

Still, S. and Precup, D. 2012. An Information-Theoretic Approach to Curiosity-driven Reinforcement learning. *Theory in Biosciences* 131(3):139-148.

Swenson, R. 1998. Spontaneous Order, Evolution, and Autocatakinetics: The Nomological Basis for the Emergence of Meaning. In *Evolutionary Systems*, G. van de Vijver, S. Salthe, and M. Delpos (eds.). Kluwer

Tishby, N. and Polani, D. 2010. Information Theory of Decisions and Actions. In *Perception-Reason-Action Cycle: Models, Algorithm and Systems.* Springer.

Tommasi, L., Peterson, M. A., and Nadel, L. 2009. *Cognitive Biology: Evolutionary and Developmental Perspectives on Mind, Brain and Behavior*. MIT Press.

Turvey, M. 2013. Ecological Perspectives on Perception-Action: What Kind of Science Does It Entail?. In W. Prinz, Beisert, M., and A Herwig (eds.), *Action Science: Foundations of an Emerging Discipline*, MIT Press.

Vijver, G. van de, Salthe, S., and Delpos, M. (eds.) 1998. *Evolutionary Systems: Biological and Epistemological Perspectives on Selection and Self-organization.* Kluwer.

Zhang, B.-T. 1993. Self-development Learning: Constructing Optimal Size Neural Networks via Incremental Data Selection, *Arbeitspapiere der GMD*, No. 768, German National Research Center for Computer Science (GMD), St. Augustin/Bonn.

Zhang, B.-T. 2008. Cognitive Learning and the Multimodal Memory Game: Toward Human-Level Machine Learning, In *IEEE International Joint Conference on Neural Networks (IJCNN 2008)*, pp. 3261-3267, IEEE Press.

Zhang, B.-T. 2013. Information-Theoretic Objective Functions for Lifelong Learning, In *AAAI 2013 Spring Symposium on Lifelong Machine Learning*, pp. 62-69, Stanford University, AAAI Press.

Zhang, B.-T., Yang, J., and Chi, S. W. 2003. Self-organizing Latent Lattice Models for Temporal Gene Expression Profiling. *Machine Learning* 52(1/2):67-89.