Generative Artificial Intelligence

Tijn van der Zant, Matthijs Kouw, and Lambert Schomaker

Abstract. The closed systems of contemporary Artificial Intelligence do not seem to lead to intelligent machines in the near future. What is needed are open-ended systems with non-linear properties in order to create interesting properties for the scaffolding of an artificial mind. Using post-structuralistic theories of possibility spaces combined with neo-cybernetic mechanisms such as feedback allows to actively manipulate the phase space of possibilities. This is the field of Generative Artificial Intelligence and it is implementing mechanisms and setting up experiments with the goal of the creation of open-ended systems. It sidesteps the traditional argumentation of top-down versus bottom-up by using both mechanisms. Bottom-up procedures are used to generate possibility spaces and top-down methods sort out the structures that are functioning the worst. Top-down mechanisms can be the environment, but also humans who steer the development processes.

1 Introduction

The field of Artificial Intelligence has not yet seen an unifying theory that captures the fundamentals for the creation of intelligent machines. Since its conception at the Dartmouth conference in 1956 at least three paradigms have permeated its existence. The top-down paradigm was supposed to be for the creation of models of the mind and eventually led to different types of logic and reasoning and later also

Tijn van der Zant
AI dept., University of Groningen, The Netherlands
e-mail: tijn@ieee.org

Matthijs Kouw
Dept. of Technology, Science and Society, Maastricht University, The Netherlands
e-mail: matthijs.kouw@gmail.com

Lambert Schomaker
AI dept., University of Groningen, The Netherlands
e-mail: schomaker@ai.rug.nl

V.C. Müller (Ed.): Philosophy and Theory of Artificial Intelligence, SAPERE 5, pp. 107–120.
springerlink.com

sub-symbolic processing. The second paradigm focussed on the creation of intelligent machines (robots). Many scientist used introspection of their own mind as a tools for the creation of these machines which was challenged in 1969 [28] using theories of biosemiotics, which is the interpretation of (sensory) signals in biological systems. This led to the bottom-up, behavior-based robotics [6]. The third, but mostly forgotten, paradigm is the field of cybernetics, , which was already being investigated when the Dartmouth conference was being held. Before the conference, in 1950, an article in the Scientific American showed two robots which consisted out of a few vacuum tubes, control loops and feedback mechanisms [16]. The field of cybernetics could be defined as: 'The theoretical study of communication and control processes in biological, mechanical, and electronic systems, especially the comparison of these processes in biological and artificial systems.'[1] The division in the field of Artificial Intelligence cannot be accepted as the answer to the question of how to build intelligent machines. An integrated perspective is necessary. The field of neo-cybernetics tries to bridge the gaps in AI with an extra addition: The property of emergence. It is unknown whether neo-cybernetics is also sufficient. What is it to study Artificial Intelligence if there is not even a common denominator within the field? The entry chosen in this article toward the creation of intelligent machines is a post-structuralist approach based on the dynamical aspects of non-linear systems. It seems that neo-cybernetics and post-structuralism meet each other in non-linear dynamical systems theory and can assist each other. Cybernetics requires the deeper underpinnings of post-structuralism, and post-structuralism can proof itself using intelligent machines based on neo-cybernetic mechanisms. This symbiosis is dubbed: Generative Artificial Intelligence [31]. In Generative AI (GAI) the possibility spaces of post-structuralism are actively being manipulated using neo-cybernetic mechanisms in order to scaffold the minds of intelligent machines.

2 Virtual-Actual

In Deleuzes actual-virtual distinction, the virtual is not so much a possible but rather fully real, waiting to be actualized. The actual is not the point of departure of change and difference, but that which has been effected from potentiality, or, the virtual [12]. This notion of the virtual allows Deleuze to describe the modal relation of potentiality against the actuality of complex systems. Thus, the virtual allows Deleuze to talk about phase spaces of systems and the patterns and thresholds that characterize their behavior. To do so, Deleuze refers to multiplicities, a term he uses to treat the multiple in itself as a substantive, rather than an attribute of substance. The realm of the virtual, also described as the plane of consistency [13] is populated by multiplicities, which provide the virtual pattern or structure of morphogenetic processes that actualize bodies, assemblages, and strata. DeLanda [11] uses Deleuzes actual-virtual distinction to propose a new agenda for science and philosophy. DeLanda wishes to provide scientific explanations of emergence: processes where novel properties and capacities emerge from a causal interaction [11]. Whereas science was previously

[1] From: http://www.answers.com/topic/cybernetics

preoccupied with simple laws acting as self-evident truths (axioms) from which all causal effects could be deduced as theorems Today a scientific explanation is identified not with some logical operation, but with the more creative endeavor of elucidating the mechanisms that produce a given effect. [11] To describe emergence, DeLanda deploys a conceptual apparatus that that consists of emergent properties, capacities, and tendencies. The sharpness of a knife is an example of an emergent property. The shape of the cross-section of the knife makes up its sharpness, which requires the knifes metallic atoms to be arranged in such a manner that they form a triangular shape. Sharpness features emergence since individual metallic atoms cannot produce the required triangular shape. What is more, sharpness provides the knife with the capacity to cut things. However, this capacity remains potential without a relational event, in this case an encounter with something that has the capacity to be cut by the knife. Similarly, the metallic atoms of the knife must have the capacity to be arranged in such a manner that sharpness emerges. Finally, the knifes blade may have the tendency to liquefy if certain conditions change, for instance in case its environment exceeds a particular temperature. Like capacities, tendencies are closely related to relational events (e.g. rising temperatures), but also to emergent properties since the metallic atoms of the knife need to interact in such a manner that the blade melts, something individual atoms cannot do. Whereas tendencies can be enumerated (e.g. the states in which a particular material can be, such as solid, liquid, or gaseous), capacities are not necessarily finite due to their dependence on being affected and / or affecting innumerable other entities. In such events, DeLanda argues in Deleuzian fashion, capacities and tendencies become actual, but neither tendencies nor capacities must be actual in order to be real. [11] Here DeLanda draws upon Deleuzes actual-virtual distinction, which allows him to ascribe reality to the virtual rather than brushing it off as a mere possible that lacks reality.

2.1 Flat Ontologies and the Machinic Phylum

A wide variety of systems can be described in terms of virtual potentialities and actualizations thereof. DeLanda [11] describes a wide variety of systems ranging from meteorological phenomena and insect intelligence to early human civilizations and stone age economics in terms of their emergent properties, capacities, and tendencies, which constitute a structure of the space of possibilities [11] that can be explored by means of computer simulations. Explanations of these different systems may builds upon explanations of lower hierarchies in a process called bootstrapping: a realist ontology may be lifted by its own bootstraps by assuming a minimum of objective knowledge to get the process going and then accounting for the rest. [10] The structures of spaces of possibilities have an objective existence [11] that can be investigated mathematically by the imposition of an arrangement through formalization or parametrizing. [11] Computer simulations enable exploration by allowing experimenters to stage interactions between different entities and investigate the emergent wholes that are the result of these interactions, thereby gaining an understanding of mechanisms of emergence. Philosophy can fulfil the role of synthesizing simulation-enabled insights into an emergent materialist world view that finally does justice to the creative powers of matter and energy. [11] In the aforementioned process of bootstrapping, DeLanda wishes to avoid the postulation of general entities (ideal types, eternal laws), since for a realist whose goal is to create a mind-independent ontology, the starting point must be those areas of the world that may be thought of as having existed prior to the emergence of humanity on this planet. (DeLanda 2009, 28) Here DeLanda aligns himself with contemporary critiques of correlationism the idea according to which we only ever have access to the correlation between thinking and being, and never to either term considered apart from the other. [20] By focusing on mechanisms of emergence, science now has the ability to describe [w]holes the identity of which is determined historically by the processes that initiated and sustain the interactions between their parts. [11] Concepts that do not elucidate sequences of events that produce emergent effects are considered irrelevant for scientific analyses. Philosophy emerges renewed, banished of reified generalities like Life, Mind, and Deity. (Ibid.) This desire to rid scientific explanations of reified generalities relates closely to the refutation of typological thinking advanced by Deleuze and Guattari [13]. Typological thinking implies that individuals are defined in terms of the species they belong to. Deleuze and Guattari argue that the members of species are not so much defined by essential traits, but by similarities in morphogenetic processes. The individual is the condition for the emergence of species, rather than vice versa. One cannot identify a species without referring to the individuals that constitute it, and the changes these individuals go through cannot be explained through the limitations put on them by the species they are said to belong to. Such imposed limits are merely restrictions of what the processes of becoming that characterize individuals, which forces them into neatly fitted categories. Deleuze and Guattari describe interacting parts (machinic elements, and emergent wholes (nomadic spaces drawn up by interacting machinic elements). These wholes may deliver assemblages that exist in a different spatio-temporal time scale when compared to their constituent parts (i.e. organisms, families, governments, nations, etc.), but they do not have a different ontological status compared to their elements [9] Similarly, researchers working in the field of complexity science explain how systems attain higher levels of complexity without relying on external organizing agents.DeLanda defines ontologies committed to the quirks and whims of individuals and their processes of becoming as flat ontologies, which can be related to Deleuze and Guattaris machinic philosophy. Such flat ontologies cannot be overcoded in dimensions supplementary to their own. Deleuze and Guattari [13] speak of a machinic phylum as a set of self-ordering material processes inherent in material, which enables emergent effects. There are in fact several phyla that tap into the self-ordering forces of material. These phyla are effectuated by assemblages, which are actualizations of the virtual (Ibid.). Machinic phyla may be explored by what Deleuze and Guattari identify as artisans, who follow the traits of materials and thereby actualize new assemblages [13]. Artisanal production relies on natural processes and the activities of the aforementioned artisans, which makes the machinic phylym as much artificial as natural: it is like the unity of human beings and

Nature. [13] The process of stratification by which assemblages are actualized from the machinic phylum can be found in areas as different as geology, biology, metallurgy, and social strata. Thus, the flat ontologies and machinic phylum of Deleuze and Guattari enable the study of processes of actualization in a variety of domains.

2.2 Minor Science and Royal Science

For DeLanda, science need not neutralize the intensive or differentiating properties of the virtual, much like his mentors Deleuze and Guattari argued. In this sense, he has much to offer constructivist debates since his work attempts to provide both an ontological and epistemological alternative to philosophies of science based on axiomatic systems, deductive logic, and essentialist typologies, one that is grounded in creative experiment rather than theory, in the multiplication of models rather than the formulation of universal laws. [3] However, unlike his mentors, DeLanda grants a particularly authoritative role to science in enabling a rigorous ontology of the virtual. A sense of ontological completion takes root in DeLandas work over the course of his various publications: from a more speculative alternative history produced by a robot historian [8], via the erudite exploration of the ability of science to engage intensities [9], to his latest book that exerts a confidence in the exploratory potential of computer simulations [11]. However, the rigorous approaches to the virtual enabled by the flat ontologies and machinic phylum of Deleuze and Guattari should not be approached in teleological terms, or a way to provide more robust criteria to evaluate scientific progress. Deleuze and Guattari emphasize the importance of what they call minor science [13], which is the kind of science deployed in artisanal production, as outlined above. Minor science works by pushing systems to their intensive states in order to follow traits (indications of 'forces', that is, singularities or self-ordering capacities) in material to reveal their virtual structures or multiplicities. [4] The difference between minor science and Royal science,

> Refers only to a differential in the rhythm and scope of the actual-virtual system From our own historically specific point of view, some terms and concepts will necessarily appear more adequate to the task than others. Science does not describe an objective state of affairs so much as inscribe a more or less mobile point of view within things themselves, causing a plurality of worlds to emerge from the virtual state of flux. [15]

Science produces more and less robust explanations, whose objectivity concerns a coalescence of relations at a particular point in time. However, the virtual always exceeds the scientific gaze and will continue to haunt the scientific observer: science thus makes a leap into ontology simply by bringing its own laws and principles into contact with the problem the chaos that haunts it thereby facilitating and allowing itself to be swept away by the movement of becoming. [15] What is more, scientific explanations intervene in the movement of becoming of the virtual on the basis of the socio-technical conditions of the scientific enterprise. A more thorough emphasis on data-driven methods will need to continuously tap into the force of the virtual as described by Deleuze and Guattari. In the phase-space of virtual exist abstract machines that are so powerful that they form the base of many of the living structures we see around on. This article tries to describe (a part of) perhaps the most important one, the one that creates thinking matter. The basics of this abstract machine consist of the following: There are self-maintaining generative mechanisms that create structures, these structures interact with an environment that selects on this process. Usually there are many (as in millions or billions) of the generative mechanisms with variations between them. Sometime these mechanisms form meshworks of interacting components, and some of these meshworks become generators themselves where (other) with other selection mechanisms, ad infinitum. This describes many processes in living organisms. Some examples are the following. Often animals and plants spawn a wealth of offspring (tens of thousands or even millions), the environment deletes the worst ones and the unlucky ones and the few that remain can become the next that generate offspring. Some plants/animals form interacting meshworks which can be in many forms such as feeding on each other, symbiotic relations, collaborations, sacrifice for the genes, ... Another example is the neurogenesis of the hominid brain. Around birth half of the neurons destruct themselves. First many neurons are created, and then the worst ones are selected against. During the first three years of the human infant many neurons have an axon battle, where the amount of axons is reduced from approximately 6 to exactly 1. Again, this is a generative mechanism (create many axons) followed by a selection mechanism (destroy the ones with the least amount of working connections). The same happens around puberty with the synaptogenesis of the dendrites, where many connections are formed in the beginning, only to be followed by almost a decade of the pruning of the synapses. In neurology these processes are called progressive and regressive processes [14]. It is the fundamental nature of these two processes, not their implementation, that Generative AI is discussing. Actual implementations will most likely not resemble the biological mechanisms created by the known process of biological evolution. It is the way in which the abstract machines operate and are implemented that bootstraps the emergence of an intelligent sub system and determines how well it operates in its environment. There is ample proof that this abstract and generative machine, if reasonably well implemented, can lead to rather flexible implementations that can operate in many different environment and handle themselves in many different situations, as exemplified by the emergence of humans during the course of biological evolution.

3 Closed and Open Systems in Artificial Intelligence

Systems in the field of Artificial Intelligence tend to be closed. As far as the authors know, all systems in AI are closed systems. These closed systems do not allow new properties to emerge. If there is flexibility at all, it only leads to a solution that the creator wanted the machine to find. This implies that for every problem a *human* has to create a new solution. This way of working will probably not lead to intelligent machines on a human-level time scale since for every little problem someone has to create a solution in the form of software. Only open-ended systems

systems [25] display interesting properties such as self-organizing and emergence [17], [32], which are required for the scaffolding of the mind [7]. Clark states:

> ...the old puzzle, the mind-body problem, really involves a hidden third party. It is the mind-body scaffolding problem. It is the problem of understanding how human thought and reason is born out of looping interactions between material brain, material bodies, and complex cultural and technological environments.

This scaffolding of the mind is what we need for the creation of intelligent machines. Without automated procedures and mechanisms that can grow, diversify and transform, humans will still be required for the creation of AI. Generative Artificial Intelligence defines itself as the field of science which studies the (fully) *automated* construction of intelligence. This is in contrast to contemporary AI, which studies the understanding and construction of intelligence *by humans*. The hidden variable is often that is requires many man-hours of work to create even the simplest solutions. What is needed for the creation of intelligent machines are automated generative methods that can be steered by humans, instead of every detail being created by humans. It is not clear what these procedures will be exactly, but the first glimpses has been seen in research that turn the usual methodology up-side-down. AI systems usually try to limit the search space of possible solutions. By doing so they also limit the possibilities of anything new arising. The closed systems from AI suffer from the problem that they all follow the same methodology, namely: Input $\rightarrow$ Process $\rightarrow$ Output (IPO). After the output the system halts, or waits for a new input. Such an IPO system will not get the needed diversity of inputs needed to find singularities in the phase space of solutions. For example, if a system is only using visual input and no tactile information, then these inputs will not increase the possibility of a learning algorithm to find the connection between hitting an object with a manipulator and seeing the object move. If on the other hand tactile information is added, then this extra amount of information flow through the system will create an extra singularity where all this information is combined. So instead of lowering the chance that a machine learning algorithm can find the connection because of the increase of information in the input space as is usually thought, it actually increases the probability of finding a solution due to an extra singularity that solves the problem. In Generative AI is it important to create generative methods that create possible solutions to problems that the machine encounters while interacting with its environment. Figure 1 give a graphical representation of the movements through a phase space. These generative methods can be implemented using software, as will be explained in the next section, but can also be due to the configuration of the machine itself, as in the previous example. The machine has sorting methods to filter out the worst solutions, and generates new solutions continuously using the best ones it has so far. The sorting machines can be manually created by humans, as in the case of Genetic Programming [18], but this would not lead to an open-ended method. Only if the machine has the opportunity to also create sorting mechanisms, partially due to pre-programmed predispositions and partially steered by its interaction with the environment (nature vs. nurture), it will be capable of displaying interesting emergent properties.

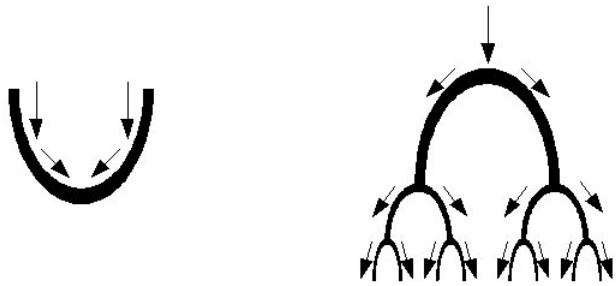

Fig. 1 Classical AI and Generative AI: In Classical AI (left figure) there is often an optimization toward some end-state and (preferably) the outcome is predictable. In both the training and the execution phase this system can be classified as: Input → Process → Output. The 'Process' part is an implemented model (hand-crafted or learned). The left figure is in a stable equilibrium. In Generative AI (right figure), the path followed through the phase space depends on the internal dynamics of the system and the interactions with the environment. The models are created and tested automatically. The creation process can be steered, but the outcome is unpredictable to some extent. After uphill explorations, the system may drop into a lower (better) energy state, with a solution which is qualitatively different from the preceding state (cf. the transition of handwritten copying to book printing). There is no difference between a training phase and an execution phase. The system learns while executing.

4 Experiments in Generative AI

4.1 Learning

Learning constitutes a core aspect of Generative Artificial Intelligence. Traditionally, learning theories were strongly embedded in reasoning, argumentation and overt cognition in general. Learning was assumed to take place in a categorical world, were instances had categorical properties and newly learned insights may be communicated by the learner using a narrative. Although this perspective on the cognitive process of learning is cogent and recognizable from within 'common sense', the paradigm has produced only few examples of convincing machine learning. Mentionable are the version-spaces symbolic learning model [21, 22] and alignment based learning in grammar induction [1]. While symbolic and explicit, such models are brittle and still far from the goal of explaining what they have learned in a natural narrative. Instead of explicit learning, the successful models of contemporary artificial intelligence are implicit, akin to Polanyi's [24] tacit knowledge: neural-network models [2], hidden-Markov models [26], support-vector machines [5] and Bayesian learning systems [27]. Although such models may either be analog or symbolic in their responses, the underlying learning process assumes a continuous parameter adaptation, either directly, as in the error back-propagation mechanism [29] for the multi-layer perceptron, or indirectly, as a consequence of exemplar weighing which takes place in the support-vector machine. Computer vision, speech and handwriting recognition and robotic control systems are trained using 'analog', numerical rather than discrete, symbolic methods. Such learning mechanisms are functional as well as adaptive and may ultimately lead to more complex models of artificial intelligence that do exhibit the potential for a verbal expression of inner states.

4.2 Humans?

However, the largest stumbling block for such a revolution is the fact that current machine-learning systems require a human researcher to provide a micro world with constraints and performance criteria. Current machine-learning methods require sufficiently large data sets of examples of patterns with their corresponding label or target responses to be produced in this micro world. The informed and motivated researcher or engineer is ever present and is steering the experimentation/exploration in great detail. The gain factor in the dissipative process [25] that takes place between the environment and the learning system is determined by an out-of-equilibrium energy state (cf. 'adrenalin') in the researcher him/herself, further motivated by the thrill of public benchmark tests and the probability of obtaining appreciation in the social context of scientific endeavor. This state of affairs is extremely costly. It leads to isolated 'feats' and successes, such as a particular type of robot filling one particular instance of a glass with a particular amount of fluid. However, the total process of wide exploration of the problem space needs to be repeated by a new PhD researcher for each small variation on the task to be learned. The total amount of costly human labor is massive and puts a ceiling on the level of attainable results in artificial intelligence.

4.3 No Humans, Machines!

What is needed are models that make use of a highly active exchange process between learner and the environment, in such a way that the problem space is continuously explored broadly, thanks to an autonomous and widely diverging bifurcation of system states. Ideally, this process unrolls, devoid of human interference but in any case requiring very little steering by humans. If the necessary input/output relations are achieved, such a system should become 'bored', i.e., divert its attention to other corners in the problem space. Similarly, if a solution pathway fails to provide performance improvement for a prolonged period, this should trigger a large jump to another location in the solution space, preferably with qualitatively different solutions than those explored along the unfruitful path. Human labor is then exchanged with another form of energy dissipation, e.g., in the form of the contemporary silicon-based von Neumann/Turing computer or a more advanced form of massively parallel computation.In a GAI engine, all aspects of human heuristic exploration will be replaced by autonomous mechanisms.

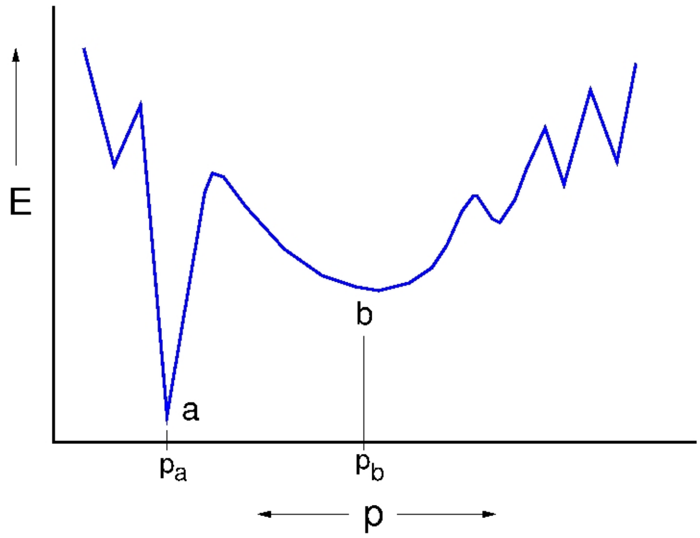

Fig. 2 The best value for parameter p needs to be found by the learner. The solution for p should have a low energy E. Is the global minimum (a) a good solution or is the value for p at point (b) to be preferred? Intelligent learners 'know' that the probability of solution (a) being useful in unseen conditions is fairly low, while the smoothness of the energy bowl at (b) gives high confidence that the value of Pb will not be very wrong in slightly varied problem conditions in the future.

4.4 What Is Needed?

What is needed for generative AI is a broadening of the concept of parameter-value search. For the solution of learning problems, usually a fitness criterion to be maximized or an energy criterion to be minimized is determined in advance. In the exploration of a high-dimensional parameter space, the criterion, say, energy E, will vary. Good solutions have a low energy, bad solutions having high energy. If the problem space is simple and idealized, the energy landscape would consist of a multi-dimensional parabola, with a clear and clean singular minimum point. In practice, however, such energy landscape are highly irregular, with many local minima such that a simplistic Newton-Lagrange method for finding 'the' solution is not feasible. One solution has already been proposed to escape this predicament and it has been widely successful. It consists of the assumption of noisy energy in the learning system, such that the exploration haphazardly jumps out of local minima, thereby increasing the probability that a deeper minimum or trough will be found. When the amount of noise ('temperature') is gradually decreased until the exploration has become deterministic, the search process is more or less guaranteed to find the deepest point. This mechanism is called simulated annealing [23] and its convergence has been demonstrated by theoretical physicist Boltzmann. However, this precursor of generative AI has three important limitations. First, a practical learner does not have the infinite time that is needed to achieve the theoretical minimum, i.e., best solution. Second, it is not always guaranteed that the deepest point in the energy landscape corresponds to the best solution. Its location in parameter space may be the consequence of lack of data. An example would be a needle-shaped pit for which statistically it can be easily demonstrated that its exact position will not be replicated in a slightly changed world. In fact, we see here that the simplistic Newton-Lagrange heuristic: ”zero partial derivatives are good, because horizontality indicates an extremum” is not enough. Not only do we want deep pits, we also prefer solutions that are characterized by a flat smooth bowl rather than a deep and steep energy ravine (Figure 2). The learner needs rich criteria in order to determine that a ’sweet spot’ has been achieved, much the same as a bird would asses a corner of the world to be an appropriate place for nesting, using a number of criteria instead of one zero-crossing of the derivative along one dimension of appropriateness. This means that we would need a much more sophisticated mechanism to evaluate the goodness of local solutions (read: hypotheses) than is currently the case in such annealing systems. A well-known variant of stochastic learning concerns the class of genetic algorithms[18]. Here, the exploration of problem space is slightly more sophisticated in that multiple local solutions are explored in parallel, and blind stochastic exploration is augmented with a ’reuse’ of partial solutions during learning. The third flaw, however, is most important. These laboratory-based learning systems assume that the process is completed once the minimum has been found: It is a training process that is detached from the real environment and its results are exported to the real world to enter the final phase in their life cycle, the operational stage. The feedback process is severed. In no way do current learning models tell us what other portions of the space are to be explored in anticipation of, or in reaction to an ever changing world.

4.5 First Glimpses

In recent work, we have implemented a very large search engine for word search in historical handwritten collections. This system, Monk [30], uses image processing and pattern recognition to identify and rank word candidates from large collections of books spanning several centuries. The diversity of writing styles requires training by human experts. However, it would be vastly expensive if a standard model of experimental machine learning would be used. This would require at least one PhD researcher per collection, with its particular image processing and handwriting style peculiarities. The challenge is to obtain an autonomous engine that accepts word labels of word images from users over internet, but learns independently, in a continuous (’24 hours/7 days’) manner. While users are motivated to correct system errors by providing labels, Monk detects where the friction is largest, either on the basis of human activity in corners of the word space or on the basis of the internal distance and probability measures indicating sub optimal solutions. A problem generator (the abbot) spawns sub tasks (novices) that execute a local word-learning or ranking task. In a cooperation between man and machine, about 300 thousand word

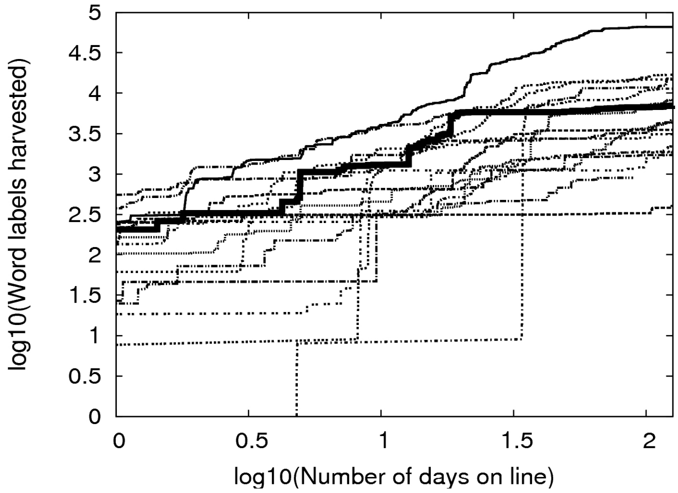

Fig. 3 Temporal evolution of the number of harvested word labels in the Monk system for handwritten word search in a number of books, in a loglog curve. Learning is characterized by growth spurts and flattening periods. The developing number of indexed words for one of the books from the Cabinet of the Queen of the Dutch National Archive is highlighted as the thick curve. The steep points in the curve are determined by the product of human effort and word-recognition performance. The latter is realized thanks to the investment of electrical energy (compute time on a high-performance cluster).

labels could be harvested. This process is ongoing. The continuity and the nature of the problem generator guarantee that both local (down-hill) optimization and diversity (up-hill global exploration) are realized. Figure 3 shows the time course for the number of harvested word labels for a number of historical books. This number is increasing over time, but it is more important to notice the discontinuity of this process. Although there may be underlying random fluctuations in both human and machine effort in training the machine, there is a non-linear speedup as evidenced by the upward jumps in the curves. If the handwriting classifier performs well on a particular word, it becomes very easy for the human volunteers to label large sets of instances as 'correct'. In a similar vein to the development of the guns, from muskets and front-loaded rifles to automatic guns and the development of air planes from the Wright plane up to modern fighter jets, there is, in Monk, a process where energy is spent on friction points in the technology: words not recognized properly elicit human irritation and subsequent efforts to smoothen the world, i.e., to create order from chaos. In our view, the process is a nice example of 'tracking the machinic phylum'. While it is too early to call this learning model in the Monk system a machine implementation of generative artificial intelligence by autonomous bifurcation processes, the results are exciting and indicative of a new way of tackling traditional 'hard' problems such as the recognition of ancient historical scripts.

5 Concluding Remarks

The creation of intelligent machines requires more than the manual tinkering by humans. This article discusses Generative Artificial Intelligence which combines neo-cybernetics and the possibility spaces of post-structuralistic philosophy. By actual experiments we demonstrate how present day machine learning technology can be applied to create generative systems where humans can steer the developmental scaffolding of the machine. Using a profound understanding of non-linear dynamical systems for the creation, and not only for the description, of intelligent systems might lead us not only to a better understanding of how to create intelligent machines. It could lead to machines that can build their own intelligence.

References

1. Adriaans, P., van Zaanen, M.: Computational Grammar Induction for Linguists. Special issue of the Journal "Grammars" with the Theme "Grammar Induction" 7, 57–68 (2004)
2. Arbib, M.A.: The Handbook of Brain Theory and Neural Networks, 2nd edn. MIT Press, Cambridge (2002)
3. Bogard, W.: Book Review: How the Actual Emerges from the Virtual. International Journal of Baudrillard Studies 2(1) (2005)
4. Bonta, M., Protevi, J.: Deleuze and geophilosophy: a guide and glossary. Edinburgh University Press, Edinburgh (2004)
5. Boser, B.E., Guyon, I.M., Vapnik, V.N.: A training algorithm for optimal margin classifiers. In: Haussler, D. (ed.) 5th Annual ACM Workshop on COLT, pp. 144–152. ACM Press, Pittsburgh (1992)
6. Brooks, R.: A robust layered control system for a mobile robot. IEEE Journal of Robotics and Automation (1986)
7. Clark, A.: Natural-Born Cyborgs: Minds, Technologies, and the Future of Human Intelligence. Oxford University Press (2003)
8. DeLanda, M.: War in the Age of Intelligent Machines. Zone Books, New York (1991)
9. DeLanda, M.: Intensive Science and Virtual Philosophy. Continuum, London (2002)
10. DeLanda, M.: Ecology and Realist Ontology. In: Herzogenrath, B. (ed.) Deleuze/Guattari & Ecology, pp. 23–41. Palgrave Macmillan, London (2009)
11. DeLanda, M.: Philosophy and Simulation: The Emergence of Synthetic Reason. Continuum, London (2011)
12. Deleuze, G.: Difference and Repetition. Translated by Paul Patton. Continuum, London (2004)
13. Deleuze, G., Guattari, F.: A thousand plateaus: capitalism and schizophrenia. Translated by Brian Massumi. Continuum, London (2004)
14. Elman, J., et al.: Rethinking Innateness: A connectionist perspective on development, Bradford. MIT Press, Cambridge (1996)
15. Gaffney, P.: The Force of the Virtual. University of Minnesota Press, Minneapolis (2010)
16. Grey, W.: An imitation of life, pp. 42–45. Scientific American (1950)
17. Hendriks-Jansen, H.: Catching Ourselves in the Act: Situated Activity, Interactive Emergence, Evolution, and Human Thought. MIT Press, Cambridge (1996)
18. Koza, J.R.: Hierarchical genetic algorithms operating on populations of computer programs. In: Proceedings of the Eleventh International Joint Conference on Artificial Intelligence IJCAI 1989, vol. 1, pp. 768–774 (1989)

19. Rabiner, L.R.: A tutorial on Hidden Markov Models and selected applications in speech recognition. Proceedings of the IEEE 77(2), 257–286 (1989)
20. Meillassoux, Q.: After Finitude. An Essay on the Necessity of Contingency. Continuum, London (2008)
21. Mitchell, T.M.: Generalization as search. Artificial Intelligence 18(2), 203–226 (1982)
22. Mitchell, T.M.: Machine learning. McGraw-Hill, Boston (1997)
23. Kirkpatrick, S., Gelatt, C.D., Vecchi, M.P.: Optimization by Simulated Annealing. Science 220(4598), 671 (1983)
24. Polanyi, M.: The Tacit Dimension. First published Doubleday & Co. (1966); Reprinted Peter Smith, ch. 1: "Tacit Knowing", Gloucester, Mass (1983)
25. Prigogine, I., Stengers, I.: Order out of Chaos: Man's new dialogue with nature. Flamingo, London (1984) ISBN 0006541151
26. Rabiner, L.R.: A tutorial on hidden Markov models and selected applications in speech recognition. Proc. IEEE 77(2), 257–286 (1989)
27. Shafer, G., Pearl, J. (eds.): Probabilistic Reasoning in Intelligent Systems. Morgan Kaufmann, San Mateo (1988)
28. Simon, H.A.: The Sciences of the Artificial, 1st edn. MIT Press, Cambridge (1969)
29. Werbos, P.J.: Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences. PhD thesis, Harvard University (1974)
30. van der Zant, T., Schomaker, L., Haak, K.: Handwritten-word spotting using biologically inspired features. IEEE Transactions on Pattern Analysis and Machine Intelligence 30(11), 1945–1957 (2008b)
31. van der Zant, T.: Generative AI: a neo-cybernetic analysis. PhD thesis, University of Groningen (2010) ISBN 978-90-367-4470-6
32. Varela, F.J., Thompson, E.T., Rosch, E.: The Embodied Mind: Cognitive Science and Human Experience. MIT Press (1992)