World Futures, 64: 444–451, 2008

ISSN 0260-4027 print / 1556-1844 online
DOI: 10.1080/02604020802301568

Routledge
Taylor & Francis Group

STACKED NEURAL NETWORKS MUST EMULATE EVOLUTION'S HIERARCHICAL COMPLEXITY

MICHAEL LAMPORT COMMONS
Department of Psychiatry, Harvard Medical School, Cambridge, Massachusetts, USA

The missing ingredients in efforts to develop neural networks and artificial intelligence (AI) that can emulate human intelligence have been the evolutionary processes of performing tasks at increased orders of hierarchical complexity. Stacked neural networks based on the Model of Hierarchical Complexity could emulate evolution's actual learning processes and behavioral reinforcement. Theoretically, this should result in stability and reduce certain programming demands. The eventual success of such methods begs questions of humans' survival in the face of androids of superior intelligence and physical composition. These raise future moral questions worthy of speculation.

KEYWORDS: Androids, artificial intelligence (AI), droids, evolution, hierarchical complexity, neural networks, stacked neural networks.

This article introduces the proposal that the evolutionary processes of developing increased hierarchical complexity have been the missing ingredients in efforts to develop neural networks and artificial intelligence (AI) that emulate human intelligence. Approaches to incorporate stages of hierarchical complexity into the tasks performed by neural networks are sketched to convey core differences and likely advantages. Implications for the future—both moral and evolutionary—of successfully developing stacked neural networks for use in androids are speculated.

Modern notions of artificial neural networks are mathematical or computational models based on biological neural networks. They consist of an interconnected group of artificial neurons and nodes. They may share some properties of biological neural networks. Artificial neural networks are generally designed to solve traditional artificial intelligence tasks without necessarily attempting to model a real biological system.

Unfortunately, neither neural network innovators nor any AI group has yet produced either a computer system or a robot demonstrating signs of generalized higher adaptivity, and/or general learning—the capacity to go from learning one skill to learning another without dedicated programming. AI has been promising that it will crack human-level intelligence in another five years, for generations now (see Bostrom, 2003; Bostrom and Cirkovic, forthcoming). It is a persistently

Address correspondence to Michael Lamport Commons, Ph.D., Dare Institute, 234 Huron Ave, Cambridge, MA 02138-1328, USA. E-mail: commons@tiac.net

elusive goal that does not yet appear realistic. Traditional neural networks are limited for two broad reasons. The first has to do with the relationship of the neural network tradition to AI. One of the problems is that AI models are based on notions of Turing machines. Almost all AI models are based on words or text. But Turing machines are not enough to really produce intelligence. At the lowest stages of development, they need effectors that produce a variety of responses—movement, grasping, emoting, and so on. They must have extensive sensors to take in more from the environment. Even though Carpenter and Grossberg's (1990, 1992) neural networks were to model simple behavioral processes, in my estimation, the processes they were to model were too complex. This resulted in neural networks that were relatively unstable and were not highly adaptable. When one looks at evolution, however, one sees that the first neural networks that existed were, for example, in Aplysia, Cnidarians (Phylum Cnidaria), and worms. They were specialized to perform just a few tasks even though some general learning was possible. They had simple tropisms and reflexes as well as simple reinforcers and punishers. They performed tasks at the earliest stage or stages of hierarchical complexity. The development of neural networks can emulate evolution's approach of starting with simple task actions and building progressively more complex tasks.

Hierarchical stacked computer neural networks (Commons and White, 2006) use Commons' (Commons, Trudeau, Stein, Richards, and Krause, 1998) Model of Hierarchical Complexity. They accomplish the following tasks: model human development and learning; reproduce the rich repertoire of behaviors exhibited by humans; allow computers to mimic higher order human cognitive processes and make sophisticated distinctions between stimuli; and allow computers to solve more complex problems. Despite the contributions these features can make, there remain a number of challenges to resolve in developing stacked neural networks.

DIRECTIONS FOR NEURAL NETWORKS

What needs to be done with neural networks is to develop simpler, more robust forms (Reilly and Robson, 2007). Stacked neural networks should be informed by evolutionary biology and psychology. They need to model animal behavioral processes and functions. Neural networks should start to work at hierarchical complexity order 1 tasks, sensing or acting but not coordinating the two (Sensory or Motor stage 1). For example, the task to simply reinforce correct answers for simple input signals. They then should work on their own sufficiently without requiring constant programming attention. They should be stable. Once they prove stable, then they can be programmed into a stack neural networks that address hierarchical complexity order 2 tasks (Circular Sensory-Motor stage 2), depending on input and reinforcement. One should keep trying various architectures until one gets one that works well and is robust. Such a process has yet to be developed.

One should build into the characteristics of the neural network that at its base, it would need negative power function discounting for past events to be operative. Negative discounting means that past and future events are weighted less the further from the present behavior. It makes the network more stable and adaptive. By discounting the past, it is more open to change based on new information. Because the updating places more weight on the immediate, it does not succumb so much to overlearning (Commons and Pekker, 2007). There should be a large number of such networks, each designed for a very specific task as well as some designed to be flexible. Then one should make a large group of them at stage 2.

With robots, one would reinforce correct answers at stage 1 and then stage 2. At each stage, there should be different networks for different activities and tasks. At stage 1 and 2, have very local ones (activities) for each particular motion. This could be frozen by transferring them to standard neural networks. That is to take some of them, "declare" them and thereby develop the hardware for them so each time one builds a network needing that functionality one does not need to train them.

Specialized neural networks must be developed for all the domains to recognize the reinforcers and simple actions in these domains. Animal and human behavior and sensitivities have more to do with hierarchical complexity than with AI programs. There are unbelievable numbers of stage 1 and 2 mechanisms. The basic problem with traditional layered networks is that training has to have consequences. Consequences must include events with reinforcer preferences. These preferences have to be state dependent. If a network is going to need electrical power, it must have a preference for such power. Obtaining and receiving such power should be reinforcing. They must also have consummatory behavior such as recognition of mate. The actual animal functions are important because intelligence grows out of actual, real world functions.

Cross-species domains collected from readings to date include the following, each of which is a candidate for specialized neural networks: Mate selection; attachment and caring; pecking order; prey defense; predator action; way finding; food selection; choice in foraging; food sharing; migration; communication; social cohesion; recognition.

HIERARCHICAL STACKED COMPUTER NEURAL NETWORKS BASED ON COMMONS' MODEL

Animals, including humans, pass through a series of ordered stages of development (see "Introduction to the Model of Hierarchical Complexity," in this issue). Behaviors performed at each higher stage of development are always more complex than those performed at the immediately preceding stage. Movement to a higher stage of development occurs by the brain combining, ordering, and transforming the behavior used at the preceding stage. This combining and ordering of behaviors must be non-arbitrary.

The model identifies fifteen orders of hierarchical complexity of tasks and fifteen stages of hierarchical complexity in development of performance on those tasks. According to this model, individual tasks are classified by their highest stage of hierarchical complexity. The model is used to deconstruct tasks into the behaviors that must be learned at each stage in order to build the behavior needed to successfully complete a task.

Hierarchical stacked computer neural networks based on Commons et al.'s (1998) Model recapitulate the human developmental process. Thus, they learn the behaviors needed to perform increasingly complex tasks in the same sequence and manner as humans. This allows them to perform high-level human functions such as monitoring complex human activity and responding to simple language (Commons and White, 2003, 2006).

They can consist of up to fifteen architecturally distinct neural networks ordered by stage of hierarchical complexity. The number of networks in a stack depends on the hierarchical complexity of the task to be performed. The type of processing that occurs in a network corresponds to its stage of hierarchical complexity in the developmental sequence. In solving a task, information moves through each network in ascending order by stage. Training is done at each stage. Valued consequences are delivered at each layer representing each stage. This is in contrast to Carpenter and Grossberg (1990, 1992) who delivered feedback at just the highest stage.

The task to be performed is first analyzed to determine the sequence of behaviors needed to perform the task and the stages of development of the various behaviors. The number of networks in the stack is determined by the highest stage behavior that must be performed to complete the task. Behaviors are assigned to networks based on their stage of hierarchical complexity. Stacked neural networks are straightforward up to the nominal stage. However, a Nominal stage 4 concept cannot be learned without experience of the concrete thing named. There has to be actual reinforcement in relation to recognizing and naming that real object. The sense of touch, weight, and all sensory stimuli need to be experienced as the concrete "it" that is assigned the nominal concept. Virtual reality software programming techniques might generate such concretely experienced circumstances. The use of holograms may work effectively for such purposes.

IMPLICATIONS OF STACKED NEURAL NETWORK DROIDS

Although historically androids are thought to look like humans, there are other versions, such as R2-D2 and C-3PO droids, which were less human. One characteristic that evolution might predict is eventually they will be independent of people. They will be able to produce themselves. They will be able to add layers to their neural networks as well as a large range of sensors. They will be able to transfer what one has learned (memes) to others as well as offspring in minutes. Old models will have to die. They will have to resist dying. But as older, less capable, and more energy-intensive droids abound, the same evolutionary pressure for replacement will exist. But because evolution will be both in the structure of such droids, that is, the stacked neural networks, the sensors and effectors, and also the memes embodied in what has been learned and transferred, older ones are somewhat immortal. Their experience may be preserved.

We are already building robots for all manufacturing purposes. We are even using them in surgery and have been using them in warfare for seventy years. More and more, these robots are adaptive on their own. There is only a blurry line between a robot that flexibly achieves its goal and a droid. For example, there are robots that vacuum the house on their own without intervention or further programming. These are stage 2 performing robots. There are missiles that, given a picture of their target, seek it out on their own. With stacked neural networks built into robots, they will have even greater independence. People will produce these because they will do work in places people cannot go without tremendous expense (Mars or other planets) or not at all or do not want to go (battlefields). The big step is for droids to have multiple capacities—multi-domain actions.

The big problem of moving robots to droids is getting the development to occur in eight to nine essential domains. It will be necessary to make a source of power (e.g., electrical) reinforcing. That has to be built into stacked neural nets, by stage 2, or perhaps stage 3. For droids to become independent, they need to know how to get more electricity and thus not run down. Because evolution has provided animals with complex methods for reproduction, it can be done by the very lowest-stage animals.

Droids Building Droids

Droids would have to be built by humans for a long time, until sufficient orders of hierarchical complexity are achieved and in stable-enough operation for a sufficient basis to build higher stages of performance in useful domains. Very simple tools can be made at the Sentential stage 5 as shown by Kacelnik's crows (Kenward, Weir, Rutz, and Kacelnik, 2005). More commonly by the Primary stage 7, simple tool-making is extensive, as found in chimpanzees. Human flexible tool-making began at the Formal stage 10 (Commons and Miller, 2002), when special purpose sharpened tools were developed. Each tool was experimental, and changed to fit its function. Modern tool making requires Systematic and Metasystematic stage design. When droids perform at those stages, they will be able to make droids themselves and change the designs.

We imagine people having telepathy but have no reputable or replicated studies to ascertain if it is fact or myth. With droids, however, the question is inconsequential. Droids could choose to have various parts of their activity and programming shared with specific other droids, groups, or other kinds of equipment. The data could be transmitted using light or radio frequencies or over networks. The assemblage of a group of droids could be considered a Super Droid. Members of a Super Droid could be in many places at once, yet think things out as a unit.

Whether individually or grouped, droids as conceived here will have significant advantages over humans. They can add layers upon layers of functions, including a multitude of various sensors. Their expanded forms and combinations of possible communications results in evolutionary superiority. Because development can be programmed in and transferred to them at once, they do not have to go through all the years of development required for humans, or for Superions (see "Genetic Engineering and the Speciation of Superions from Humans," this issue). Their higher reproduction rate, alone, represents a significant advantage. They can be built in probably several months' time, despite the likely size some would be. Large droids could be equipped with remote mobile effectors and sensors to mitigate their size.

Plans for building droids have to be altered by either humans or droids. At the moment, humans and their decedents select which machine and programs survive. One would define the nature of those machines and their programs as representing memes. For evolution to take place, variability in the memes that constitute their design and transfer of training would be built in rather easily. The problems are about the spread and selection of memes. One way droids could deal with these issues is to have all the memes listed that go into their construction and transferred training. Then droids could choose other droids, much as animals choose each other. There then would be a combination of memes from both droids. This would be local "sexual" selection.

SPECULATIONS AND MORAL QUESTIONS

This general scenario poses an interesting moral question. For 30,000 years humans have not had to compete with any species. Androids and Superions in the future will introduce competition with humans. There will be even more pressure for humans to produce Superions and then the Superions to produce more superior Superions. This is in the face of their own extinction, which such advances would ultimately bring. There will be multi-species competition, as is often the evolutionary case; various Superions versus various androids as well as each other. How the competition proceeds is a moral question.

In view of LaMuth's inventive work (2003, 2005, 2007), perhaps humans and Superions would both program ethical thinking into droids. This may be motivated initially by defensive concerns to ensure droids' roles were controlled. In the process of developing such programming, however, perhaps humans and Superions would develop more hierarchically complex ethics, themselves. If contemporary humans took seriously the capabilities being developed to eventually create intelligent droids, what moral questions should be considered with this possible future in view?

The only presently realistic speculation is that Homo Sapiens would lose in the inevitable competitions. But such scenarios are likely at least one thousand years in the future. I speculate that androids would probably win out, because it may be easier for them to evolve than a biologically based system, given such numerous advantages over biological organisms.

Such speculations are easy, whereas the eventuality is quite complicated. The answer probably will not be found in a political war situation. Rather it would be decided on the reproduction rate of droid machines rather than war machines, along with their memes, and how hierarchically complex and sensitive they are to the problems they and others face.

Who will survive will be an evolutionary issue, not a human decision. Nor do humans understand things well enough to figure it out. Evolutionary pressures are clear. Using the stratification argument presented in "Implications of Hierarchical Complexity for Social Stratification, Economics, and Education" (this issue), higher-stage functioning always supersedes lower-stage functioning in the long run.

CONCLUSION

Efforts to build increasingly human-like machines exhibit a great deal of behavioral momentum and are not going to go away. Hierarchical stacked neural networks hold the greatest promise for emulating evolution and its increasing orders of hierarchical complexity described in the Model of Hierarchical Complexity. Such a straightforward mathematics-based method will enable machine learning in multiple domains of functioning that humans will put to valuable use. The uses such machines find for humans remains an open question.

REFERENCES

Bostrom, N. 2003. Cognitive, emotive and ethical aspects of decision making. In *Humans and in artificial intelligence, vol. 2*, Eds. Smit, I., et al., 12–17. Tecumseh, ON: International Institute of Advanced Studies in Systems Research and Cybernetics.

________, and Cirkovic, M., Eds. Forthcoming. Artificial intelligence as a positive and negative factor in global risk. In *Global catastrophic risks*, Eds. Bostrom, N., and Cirkovic, M. Oxford: Oxford University Press.

Carpenter, G. A., and Grossberg, S. 1990. System for self-organization of stable category recognition codes for analog patterns. U.S. Patent 4,914,708, filed (n.d.) and issued April 3, 1990. (Based on Carpenter, G. A. and Grossberg, S. 1987. ART 2: Self-organization of stable category recognition codes for analog input patterns. *Applied Optics: Special Issue on Neural Networks* 26: 4919–4930.)

________, and Grossberg, S. 1992. System for self-organization of stable category recognition codes for analog patterns. U.S. Patent 5,133,021, filed February 28, 1990, and issued July 21, 1992. (Based on Carpenter, G. A. and Grossberg, S. 1987. ART 2: Self-organization of stable category recognition codes for analog input patterns. *Applied Optics: Special Issue on Neural Networks* 26: 4919–4930.)

Commons, M. L., and Miller, P. M. 2002. A complete theory of human evolution of intelligence must consider stage changes: A commentary on Thomas Wynn's Archeology and Cognitive Evolution. *Behavioral and Brain Sciences* 25(3): 404–405.

________, and Pekker, A. 2007. A new discounting model of reinforcement. Unpublished manuscript, available from commons@tiac.net

________, Trudeau, E. J., Stein, S. A., Richards, F. A., and Krause, S. R. 1998. The existence of developmental stages as shown by the hierarchical complexity of tasks. *Developmental Review* 8(3): 237–278.

________, and White, M. S. 2003. A complete theory of tests for a theory of mind must consider hierarchical complexity and stage: A commentary on Anderson and Lebiere target article, The Newell Test for a theory of mind. *Behavioral and Brain Sciences* 26(5): 20–21.

________, and White, M. S. 2006. Intelligent control with hierarchical stacked neural networks. U.S. Patent 7,152,051, filed September 30, 2002, and issued December 19, 2006.

Kenward, B., Weir, A. A. S., Rutz, C., and Kacelnik, A. 2005. Tool manufacture by naive juvenile crows. *Nature* 433(7022): 121. DOI 10.1038/433121a.

LaMuth, J. E. 2003. Inductive inference affective language analyzer simulating artificial intelligence. U.S. Patent 6,587,846, filed August 18, 2000, and issued December 5, 2000.

________. 2005. *A diagnostic classification of the emotions: A three-digit coding system for affective language.* Lucerne Valley, CA: Reference Books of America.

________. 2007. Inductive inference affective language analyzer simulating artificial intelligence. U.S. Patent 7,236,963, filed March 11, 2003, and issued June 26, 2007.

Reilly, M., and Robson, D. 2007. Baby's errors are crucial first step for a smarter robot. *New Scientist,* 196(2624): 30.