In: P. Hammer, M. Alirezaie, C. Strannegård (eds.), Proc. Artificial General Intelligence 2023, 296-306

Addressing the Unsustainability of Deep Neural Networks With Next-Gen AI

Amanda Vallentin[1,2] & Kristinn R. Thórisson[1,3] & Hugo Latapie[4]

[1] Reykjavik University, Center for Analysis and Design of Intelligent Agents, Menntavegur 1, Reykjavík, Iceland thorisson@ru.is
[2] IT University of Copenhagen
Rued Langgaards Vej 7, Copenhagen, Denmark amgv@itu.dk
[3] Icelandic Institute for Intelligent Machines, Reykjavík, Iceland
[4] Cisco Systems, Emerging Technologies & Incubation, San Jose, California, USA
hlatapie@cisco.com

Abstract. Humanity is currently facing one of its biggest challenges to date: The climate crisis. As a result, most industry sectors are reassessing their ways of working to be better equipped to address their share of the situation. The digital sector often gets set aside in such considerations in talk about the green transition because a significant amount of its work consists of optimizing processes that can save resources. Deep neural networks (DNNs) have gained great attraction and have shown good results regarding process automation. We argue that there are well-known and lesser known negative side-effects to automation frameworks based on DNNs (and related technologies) in terms of energy consumption, pollution, and social equality, that must be questioned. We analyze the operating principles and deployment methods of DNNs, the new era of automation efforts this has launched, and argue on this basis that their continued use is both unsustainable and indefensible. Using three examples of ongoing research, we explain how alternative approaches to develop more general machine intelligence are well-poised to power the next phase of AI-based automation.

Keywords: Artificial Intelligence · Methodology · Deep Neural Networks · General Machine Intelligence · Empirical Reasoning · Automation · Energy Consumption · Pollution · Social Equality · Innovation

1 Introduction

The IT-sector is one of the most innovative and fastest-growing industries worldwide[5]. The bleeding edge lies arguably in automation technologies, in no small part because of the obvious incentive that reduced cost and increased speed

[5] According to Statista, IT-related revenue has a predicted annual growth rate of 6.86% and a predicted market volume of US$1,570.00bn by 2027 (https://www.statista.com/outlook/tmo/it-services/worldwide – *accessed March 1st, 2023)*.

translates directly to increased revenue. Within contemporary[6] applied automation technologies, DNNs are the latest arrivals with significant potential for various applications. Spurred by predictions of its usefulness for a wide range of tasks[7], unbridled optimism has often characterized its coverage in the media. For instance, a Forbes article presents the 13 skills AI already has today (including "smell" and "reading your mind") [19]; the Guardian explains how AI is changing how a number of different industry sectors operate [6]. However, after a period of experimentation it is increasingly clear that DNN deployment is unavoidably hampered by inherent deep limitations [17] and hidden costs [33, 4].

DNNs risk compromising the path towards sustainable development of society[8]. In particular, their runtime and updating methodologies make them unsustainable [33, 16]. Another limitation has to do with how they are developed. Energy consumption during DNN development is incredibly high (PaLM, a language model from Google, consumed about 3.4 GWh in about 2 months [3]). For this reason, and others, very few companies will be able to afford developing them because of their sheer size and compute requirements (the BLOOM model, with 175 billion parameters, cost US$7 million to develop [3]). So this approach is inappropriate for parties with only small data and small funding. DNNs are thus nowhere nearly as appropriate or powerful for being deployed in automation tasks as past and present moves by tech giants might indicate [2, 5].

This paper has two main parts: In sections 2 and 3 we detail what we consider key limitations of DNNs, and in section 4 and 5 we discuss how these could be overcome through research on artificial general intelligence.

2 Deep Limitations of Deep Neural Networks

In recent years both the size and the training data of DNNs have exploded due to an incentive to upscale the models to reach better performance [33, 3]. GPT-3, one of the biggest language models to date, has shown great results in some text generation tasks [5] leading many to think that DNNs can be used to solve any task. We argue that this is neither wise nor possible.

---

[6] Our use of the term 'contemporary AI' refers to a set of methodologies that are currently in active experimentation or use *in industry*, including but not limited to reinforcement learning, ANNs of all kinds, and other well-known methods.

[7] For instance, the annual prediction that "full self-driving cars will be available next year" has been updated at a rate of one year per year by Tesla's CEO (*"Watch Elon Musk Promise Self-Driving Cars 'Next Year' Every Year Since 2014,"* `https://futurism.com/video-elon-musk-promising-self-driving-cars` — *accessed March 1st, 2023.*

[8] The UN defines 'sustainable development' as harmony between economic growth, social inclusion, and environmental protection. `https://www.un.org/sustainabledevelopment/development-agenda/` — *accessed April 4th, 2023.*

2.1 DNNs: Expensive to develop and use

A search with Google's new chatbot Bard can cost the company 10 times more than a traditional key word search [21]. However, the total cost of the models are already high before they leave the lab (cf. the BLOOM model required $7 million worth of computing time during its development). To provide necessary computing power for the training phase, the developers of DNNs also need access to expensive specialized hardware [33]. Attempts have been made to measure the environmental footprint of large language models (LLM)[33, 16]. Luccioni et al. 2022 uses a life cycle analysis approach to estimate a more realistic environmental footprint for a LMM called BLOOM. When they add emissions from all training activities and experiments (not just the final training run) as well as emissions from the infrastructure that maintains the hardware and emissions from manufacturing the hardware, the total footprint of BLOOM is ∼124 tons CO_2eq [16]. However, the carbon intensity of the grid used to train BLOOM was only $57gCO_2eq/kWh$(trained in France) compared to GPT-3 where the carbon intensity of the grid was $429gCO_2eq/kWh$(trained in the US). Unfortunately, we only know the power consumption of the final training phase of GPT-3. Comparing the estimated carbon emissions from the two models' final training phase, BLOOM was ∼25 tons CO_2eq and GPT-3 was ∼502 tons CO_2eq, which is a significant difference since the models have about the same amount of parameters [16]. The environmental footprints of DNNs are strongly influenced by carbon intensity of the energy grid. Considering that US and China are the biggest players in the AI market[20], and they use about 81% and 83% fossil fuels[27], the estimated footprint of only a small part of GPT-3's life cycle is worrying.

Product	Description	CO_2eq (kg)	Scale
One laptop [1]	Entire life cycle incl. power use (avg.)	423	1
One automobile [33]	Entire life cycle incl. fuel use (avg.)	57,153	135
One GPT-3 [16]	Final training phase	502,000	1187

Table 1. Comparison of the CO_2 emission of different products. The **Scale** column shows the emissions multiplier matching a laptop computer's complete lifecycle.

Another aspect that can raise the economic and environmental price is when the model needs to be corrected after it is deployed. In this case, the model needs to be taken down and retrained since DNNs cannot be taught anything new once they have left the lab. What often happens is that once the models are released, they act in unexpected ways and the developers need to spend more resources on making them behave. In 2016 when Microsoft created a twitter account for the chatbot Tay, it went from tweeting innocent tweets like "I love feminism now" to "Hitler was right I hate jews" in a single day, despite being trained on safe data as Microsoft claims [9], resulting in the bot having to be taken down. Seven years later, Microsoft ran into a similar problem when a journalist at the New York Times had a conversation with their new chatbot that ended with the bot confessing its love to him and telling him to leave his wife [28]. After the incident, attempts were made to prevent the chatbot from answering personal questions, but even with countless reboots and alterations, it could not be guaranteed that it behaved according to plan. Some have even made this into a sport (called 'JailBreaking') where they share and test ways of getting around the "lobotomized" chatbots and make them say racist, misogynistic, etc. statements. Considering the enormous resources spent on controlling the DNNs' behaviors after they leave the lab, it seems that proper kinds of control mechanisms are missing. This is, however, hardly surprising, since DNNs are primarily based on statistical methods and have no obvious ways of being steered through explicit goals or hierarchical rules.

2.2 The limited "learning" of statistics-based systems

All animals learn cumulatively because the world does not reveal itself to anyone all-at-once. The "learning" that contemporary AI systems practise is a very special case of what is normally called 'learning,' and it greatly limits which kind of tasks they can be "trained" to solve well. Research by Eberding et al. [7] compared several different types of DNN-based learners (they also tested the AGI-aspiring NARS — we discuss this in a later section) on the well-known cart-pole balancing task, which consists of learning to balance a stick standing on a cart by issuing right and left commands ('R' and 'L'). Once the various AI learners had achieved this task, the researchers reversed the directional commands. The performance of the various learning algorithms to adjust their prior training to this new condition is recorded. In a final scenario, the researchers switch back to the original control method. The performance of all tested DNN-based learners dropped significantly in the reversed phase, and it takes them many more iterations to reach the same performance, once the controls are switched back to the original settings. The change in the task had to be "unlearned" through enough new interactions before the performance could return to what it was before the controls were switched. None of them returned to their original performance.

This research exposes how DNNs have a static and simplistic representation of the world. They are not capable of inferring simple relations (in Eberding's [7] experiment, that the controls were switched around) which makes them unusable for many tasks where reasoning is of importance, like math. The best DNN score on the MATH data set is 50% [3]. The developers managed to reach this score by training it only on mathematics-related texts and up scaling the size of the model to astonishing 540 *billion* parameters. With this strategy they were hoping that the model would evolve to be able to perform reasoning through pattern-recognition alone [3]. There have been attempts of creating reasoning abilities in DNNs, for example using chain-of-thought prompting. The technique improves models' scores on certain data sets, but in bigger models [40]. This method does not, in fact cannot, turn ANN-based systems into reliable reasoners.

2.3 DNN autonomous learning after it leaves the lab: 'Undefined'

Another of DNNs limitations is that once they are trained, their knowledge is fixed and they cannot be easily applied to another task. When faced with something that was not part of their training data, performance decreases or they do something that is unpredictable. This is likely one of the reasons why self-driving cars have not met their makers' expectations; there are countless scenarios an artificial driver must be able to navigate before it is safe to let it out on the roads. The upshot is, when it comes to complex tasks, DNNs cannot be trusted, due to the countless road scenarios that may occur. There are attempts to overcome this problem, for example a one-shot learning model can classify images it has not seen in its data set. However, the models are more computationally heavy to run and they only work if the image is similar to the ones in the training set [15].

2.4 DNNs and social inequality

When looking at LLMs, the data size requirements have exploded in recent years. BERT was trained on 16 GB data in 2019 and GPT-3 was trained on 570 GB in 2020 [4]. Firstly, it is difficult to get a hold of this much data and secondly, it is nearly impossible to ensure that the data has the right quality. In LLMs this manifests itself in a bias against minorities because most of their data has been scraped of sources like Reddit, Twitter, and Wikipedia where the majority of writers are white males [4]. In medical AI, we also see discrimination of patients because it is difficult to acquire data sets that are representative for all genders, ages, and races [25]. Due to the data requirements and cost of DNNs, their increased use will risk worsening inequality, as not everyone has equal access [33] or is equally represented. Healthcare models only work for groups represented in the data sets.

Additionally, the price of using the DNNs will limit which users have access them. For instance, ChatGPT has recently made headlines about being able to pass several advanced exams at universities[9]. If not all students have equal access to DNN aids, we risk increasing social inequality[10].

2.5 DNNs' domination of the AI narrative

Despite the known limitations of DNNs, development of alternative approaches to making machines smarter suffers from their media dominance. The private sector has great influence on AI research and they tend to favor data-hungry and computationally heavy DNNs [11].

Due to inordinate emphasis on a single technology, young researchers may be lead to believe that deep learning methods (a) are the end-all, be-all, (b)

[9] ChatGPT has passed the Wharton Exam, US medical licensing exam, law school exam, and others. (`https://www.businessinsider.com/list-here-are-the-exams-chatgpt-has-passed-so-far-2023-1?r=US&IR=T#wharton-mba-exam-1` — *accessed March 4th, 2023).*

[10] As of April 2023, the price is $20 a month for reliable and fast access to ChatGPT, although a free version with slower response is still available. (`https://openai.com/blog/chatgpt-plus` — *accessed April 4th, 2023).*

will overcome all the challenges we accounted for, and (c) will continue to be a key technology in our society [11]. It is no surprise that we see this development because many of the key DNN researchers still seem to believe that the technology will overcome all these challenges with more data and more efficient hardware. Kaplan & McCandlish [10] argue that there exists a scaling law for neural language models, suggesting that there is more to gain if we continue with enlarging the DNNs. Altman predicts an AI revolution because of the incredible wealth that will be created as DNNs replace the majority of our workforce [2].

However, there is ample evidence that DNNs are not living up to such expectations—and it probably never will [37]. The optimism echoes claims made of the Cyc project in the 80s and 90s [14]. Looking at DNNs' abilities regarding common sense, Marcus and Davis [18] recently challenged ChatGPT-3's presumed theory of mind, arguing that the results do not show an ability of common sense but rather that, due to being trained on data about thought-experiments and logic tests, it can predict the answers on purely linguistic principles. When the phrasing of questions changes slightly or the questions are asked in another language, GPT-3 shows no sign of having a theory of mind. In a study by Stojnic et al. [32], they compared DNNs common sense ability to infants and the study revealed that the DNNs failed and did not appear to have common sense.

Along with over-promising in the field of DNNs, there is a lack of innovation that misleads newcomers, governments, and institutions who continue to support research on the topic. By ignoring other strategies, society is not only wasting precious resources but also risking the field of AI as a whole to lose trust.

3 Summary of Limitations

Based on the foregoing, there can be little doubt that contemporary AI methodologies, in all their variations, come with significant limitations. DNNs are monolithic technologies with limited scope. They only work well when they are built for a well-defined limited task with extensive amounts of data of a certain quality. If any changes are necessary due to unwanted behavior or a slightly different task, the models must be rebuilt, repeating their resource-demanding training cycles. Combined with the potential decrease in social equality, we have a technology that both compromises social inclusion and environmental protection. This is unsustainable. To summarize the limitations of DNNs discussed so far:

- are exceedingly expensive to develop and use
- have a large environmental footprint due to energy consumption
- are difficult to control
- only work well for certain types of tasks
- are difficult and expensive to adapt to new tasks
- can increase inequality in the world
- take away focus and resources from other approaches in AI

It is neither good for the field of AI nor for society at large that the inordinate amount of funding and effort poured into DNNs and related technologies continues. How can we move forward to more sustainable AI?

4 Breaking the Stalemate Through Innovation

Examples of similar situations can be found in recent history of innovation, where a single framework had become too entrenched too early. One example is the global windmill industry in the 1970s. Due to the energy crisis at the time, there was a push towards finding cheaper energy sources and many countries tried to develop megawatt windmills [22]. In Denmark another approach was taken, where smaller companies developed smaller and more experimental windmills and met up at annual windmills conferences and shared their results [22]. The companies had incentive to do so because many private individuals were interested in buying their own local windmill, since the government would pay 30 percent of such investment [22]. Due to this approach, the development of a new type of windmill was undertaken, one in which risk was lowered due to the willingness of the Danish population to buy smaller windmills. As a result the windmill industry was born in Denmark, which produced the best windmills.

In other countries a more conservative approach was chosen by attempting to upscale the best existing windmills at the time, with the aim of turning them into megawatt windmills. All of those approaches failed as they could not compete with the Danish models, which were cheaper yet more robust [22]. The current development of contemporary AI where researchers and companies upscale their frameworks (more data and bigger models), believing that "bigger is better," resembles what we saw in the windmill industry in the 70's. In our view, a wholly new methodological paradigm is called for to develop more autonomous AI that is more capable and whose behavior is easier to manage and predict.

5 Sustainable Automation Via AGI

The main limitations of DNNs can be grouped into three sets based on their source: (a) opaqueness, (b) learning style, and (c) representation. All of three have made a regular occurrence throughout much of AI research [37], or certainly since the start of the annual AGI conference series in 2008. Here we present an overview of selected recent work focusing on these areas.

From opaque to transparent knowledge. A powerful way to represent knowledge[11], that makes it directly inspectable by human or machine, is to make its structure explicitly hierarchical. Representing knowledge explicitly was of course common in the expert systems of the 1970s, and some research in AI has continued this tradition. The approach comes with known limitations, which can be overcome by taking specific steps. For instance, the Non-Axiomatic Reasoning System (NARS) represents knowledge as defeasible [26] statements that nevertheless support reasoning; indeed, NARS-based systems learn through reasoning processes that mix (non-axiomatic) deduction, abduction, and induction (cf. [39, 13, 8]). Other systems take a compatible approach but use a different

[11] By 'knowledge' we mean a form of 'actionable information'—that is, information that can be used for making plans and getting things done in a particular environment.

knowledge representation scheme, e.g. the Autocatalytic Endogenous Reflective Architecture (AERA [24]). The results demonstrated by prototypes developed by Latapie et al. [13] show that systems relying on explicit knowledge representation have come a long way, yet their funding is in no way proportional to the results achieved. These systems work on vastly smaller data than DNN-based systems, and thus use much less energy.[12]

Besides non-axiomatism, another way to overcome the limitations of approaches based on logic statements is to step up to second-order representation, allowing the system to inspect and operate on its own knowledge [34]. Such reflective systems have unfortunately not been given sufficient attention in the AI literature. The results of Nivel et al.'s [23] research on teaching an AERA-based agent to learn by observation how to conduct TV-style interview on the topic of recycling in under 21 hours, including learning the syntax and semantics of a 100-word vocabulary, how to take turns in dialog, manipulation of objects, deictic gestures of various forms, and more – from scratch – should suffice to convince anyone that this very iconoclastic approach to machine learning should be pursued more vigorously by the AI community.

From once-and-for-all learning to cumulative learning. Learning in nature has no choice but to proceed incrementally, because the world does not reveal itself to learners all-at-once. This means that the knowledge representation scheme must be updatable piece-wise [36]. Furthermore, any autonomous system deployed in the physical world will encounter situations that are not identical to something experienced before. Automating the handling and learning from these is imperative for advancing the state of industrial automation. Viable solutions to cumulative learning have already been proposed [8, 36, 39].

Compositional knowledge representation. This topic is closely related to the first point, which is to say that compositional knowledge representation goes hand-in-hand with knowledge transparency and cumulative learning. The ability to construct a goal hierarchy autonomously is a foundational requirement for any AI that is to operate autonomously (or even semi-autonomously); the designers cannot possibly foresee every and all situations that the system may encounter. A goal hierarchy that the system can itself manipulate safely is a necessity. Thórisson [38] presents arguments that general autonomous learning is not possible without the capacity for some form of explanation generation.

While fully-functional AGI systems are still in their early phase of development, some examples are leading the way (cf. [13, 8, 31, 30, 12, 29]). All this points towards next-generation systems having the potential to become a green alternative to DNNs, promising easier reuse, increased generality, significantly less energy consumption, lower data requirements, less compute power, and a wider range of applications.

Acknowledgments. This work was supported in part by Cisco Systems, the Icelandic Institute for Intelligent Machines and Reykjavik University.

[12] The AERA system, for example, learned to do a TV-style interview after learning for only 20 hours on a 6-core office desktop machine [35].

References

1. What is the carbon footprint of a laptop? Circular Computing (2021), https://circularcomputing.com/news/carbon-footprint-laptop/
2. Altman, S.: Moore's law for everything (2021), https://moores.samaltman.com/
3. Ananthaswamy, A.: In AI, is bigger always better? Nature **615**, 202–205 (2023)
4. Bender, E.M., McMillan-Major, A., Gebru, T., Shmitchell, S.: On the dangers of stochastic parrots: Can language models be too big? Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency pp. 610–623 (2021)
5. Brown, T.B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., et al.: Language models are few-shot learners. Advances in Neural Information Processing Systems **33** (2020)
6. Butler, S.: From retail to transport: how ai is changing every corner of the economy. The Guardian (2023), https://www.theguardian.com/technology/2023/feb/18/from-retail-to-transport-how-ai-is-changing-every-corner-of-the-economy
7. Eberding, L.M., Thórisson, K.R., Prabu, A., Jaroria, S., Sheikhlar, A.: Comparison of machine learners on an aba experiment format of the cart-pole task. Proceedings of Machine Learning Research **159**, 49–63 (2021)
8. Hammer, P.: Reasoning-learning systems based on non-axiomatic reasoning system theory **192**, 88–106 (2022)
9. Hunt, E.: Tay, microsoft's AI chatbot, gets a crash course in racism from twitter. The Guardian (2016), https://www.theguardian.com/technology/2016/mar/24/tay-microsofts-ai-chatbot-gets-a-crash-course-in-racism-from-twitter
10. Kaplan, J., McCandlish, S., Henighan, T., Brown, T.B., Chess, B., Child, R., Gray, S., Radford, A., Wu, J., Amodei, D.: Scaling laws for neural language models (2020)
11. Klinger, J., Mateos-Garcia, J.C., Stathoulopoulos, K.: A narrowing of AI research? SSRN Electronic Journal (2020)
12. Kommrusch, S., Minsky, H., Minsky, M., Shaoul, C.: Self-supervised learning for multi-goal grid world: Comparing leela and deep q network. Proceedings of Machine Learning Research **131**, 81–97 (2020)
13. Latapie, H., Gabriel, M., Kompella, R.: Hybrid AI for IoT actionable insights & real-time data-driven networks. Proceedings of Machine Learning Research **159**, 127–131 (2022)
14. Lenat, D.: Cyc: A large-scale investment in knowledge infrastructure. Communications of the ACM **38**, 33–38 (2023)
15. Logunova, I.: A guide to one-shot learning (2022), https://serokell.io/blog/nn-and-one-shot-learning
16. Luccioni, A.S., Viguier, S., Ligozat, A.L.: Estimating the carbon footprint of bloom, a 176b parameter language model (2022), https://arxiv.org/abs/2211.02001
17. Marcus, G.: Deep learning is hitting a wall. Nautilus (2023), https://nautil.us/deep-learning-is-hitting-a-wall-238440/
18. Marcus, G., Davis, E.: How not to test gpt-3 (2023), https://garymarcus.substack.com/p/how-not-to-test-gpt-3
19. Marr, B.: 13 mind-blowing things artificial intelligence can already do today. Forbes (2019), https://www.forbes.com/sites/bernardmarr/2019/11/11/13-mind-blowing-things-artificial-intelligence-can-already-do-today/?sh=39e42add6502
20. Maslej, N., Fattorini, L., Brynjolfsson, E., et al.: Artificial Intelligence Index Report 2023. Technical Report, Stanford University (2023), https://aiindex.stanford.edu/report/

21. Mok, A.: AI is expensive. a search on Google's chatbot Bard costs the company 10 times more than a regular one, which could amount to several billion dollars. Business Insider (2023), https://www.businessinsider.com/ai-expensive-google-chatbot-bard-may-cost-company-billions-dollars-2023-2
22. Nielsen, K.H.: Vindmøllens historie: Sådan tæmmede danskerne vindens energi. ForskerZonen (2018), https://videnskab.dk/forskerzonen/teknologi-innovation/vindmoellenshistorie-saadan-taemmede-danskerne-vindens-energi
23. Nivel, E., Thórisson, K.R., et al.: Autonomous acquisition of natural language. IADIS International Conference on Intelligent Systems & Agents pp. 58–66 (2014)
24. Nivel, E., Thórisson, K.R., Steunebrink, B., Dindo, H., et al., G.P.: Autocatalytic endogenous reflective architecture. Tech report RUTR-SCS13002, Reykjavik University – School of Computer Science (2013)
25. Panch, T., Mattie, H., Celi, L.A.: The 'inconvenient truth' about AI in healthcare. NPJ Digital Medicine (2019)
26. Pollock, J.L.: Defeasible reasoning and degrees of justification. Argument and Computation **1**(1), 7–22 (2010)
27. Ritchie, H., Roser, M., Rosado, P.: Energy. Our World in Data (2022), https://ourworldindata.org/energy-mix
28. Roose, K.: A conversation with Bing's chatbot left me deeply unsettled. The New York Times (2023), https://www.nytimes.com/2023/02/16/technology/bing-chatbot-microsoft-chatgpt.html
29. Sheikhlar, A., Thórisson, K.R., Eberding, L.M.: Autonomous cumulative transfer learning. In: International Conference on Artificial General Intelligence. pp. 306–316 (2020)
30. Sheikhlar, A., Thórisson, K.R., Thompson, J.: Explicit analogy for autonomous transversal learning **192**, 48–62 (2022)
31. Steunebrink, B., Swan, J., Nivel, E.: The Holon system: Artificial general intelligence as 'work on command'. Proceedings of Machine Learning Research **192**, 120–126 (2022)
32. Stojnic, G., Gandhi, K., Yasuda, S., Lake, B.M., Dillon, M.R.: Commonsense psychology in human infants and machines. Cognition **235** (2023)
33. Strubell, E., Ganesh, A., McCallum, A.: Energy and policy considerations for deep learning in NLP. 57th Annual Meeting of the Association for Computational Linguistics (ACL) (2019)
34. Thórisson, K.R.: A new constructivist AI: From manual construction to self-constructive systems. In Theoretical Foundations of Artificial General Intelligence (Pei Wang and Ben Goertzel, eds.) **4**, 145–171 (2012)
35. Thórisson, K.R.: Seed-programmed autonomous general learning. Proceedings of Machine Learning Research **131**, 32–70 (2020)
36. Thórisson, K.R., Bieger, J., Li, X., Wang, P.: Cumulative learning. In: Proceedings of the 12th International Conference on Artificial General Intelligence. pp. 198–208. Springer, Shenzen, China (2019)
37. Thórisson, K.R., Minsky, H.: The future of AI research: Ten defeasible 'axioms of intelligence'. Proceedings of Machine Learning Research **192**, 5–21 (2022)
38. Thórisson, K.R.: The 'Explanation Hypothesis' in general self-supervised Learning. pp. 5–27. Proceedings of Machine Learning Research, IWSSL-21 (2021)
39. Wang, P.: Rigid flexibility: The logic of intelligence. Springer Science, vol. 34 (2006)
40. Wei, J., Wang, X., Schuurmans, Dale Bosma, M., Ichter, B., Xia, F., Chi, E.H., Le, Q.V., Zhou, D.: Chain-of-thought prompting elicits reasoning in large language models (2023)