Chapter 31
Computational Complexity and Cognitive Science: How the Body and the World Help the Mind be Efficient

Peter Gärdenfors

Abstract Computational complexity has been developed under the assumption that thinking can be modelled by a Turing machine. This view of cognition has more recently been complemented with situated and embodied cognition where the key idea is that cognition consists of an interaction between the brain, the body and the surrounding world. This chapter deals with the meaning of complexity from a situated and embodied perspective. The main claim is that if the structure of the world is taken into account in problem solving, the complexity of certain problems will be reduced in relation to Turing machine complexity. For example, search algorithms can be simplified if the visual structure of the world is exploited. Another case is the logical problem of language acquisition, claiming that children cannot learn language simply by considering the input. It is argued that this problem will not arise if it is taken into account that children's learning of grammatical features often exploits world knowledge.

31.1 The Notion of Complexity in Cognitive Science

Cognitive science comes in three flavours [6, pp. 83–84], [11]. The historically first is *classical computationalism*. The basic tenets are that the brain is a computer (Turing machine) and that all thinking is manipulation of symbols (e.g. [8, 9]). The second is *connectionism* (associationism). Here the central tenets are that the brain can be seen as a neural network and that thinking can be described as parallel distributed processing in such a network [25]. The third is *situated and embodied cognition* where the key idea is that cognition consists of an interaction between the brain, the body and the surrounding world. Thinking is not encapsulated in the brain but it leaks out into the world [6].

P. Gärdenfors (✉)
Cognitive Science, Lund University, Lund, Sweden
e-mail: Peter.Gardenfors@lucs.lu.se

A. Baltag and S. Smets (eds.), *Johan van Benthem on Logic and Information Dynamics*, Outstanding Contributions to Logic 5,
DOI: 10.1007/978-3-319-06025-5_31,

Classical computationalism entails that all cognitive science can be reduced to the study of computers and the algorithms that are run on them. Isaac et al. (this volume) formulate this idea crisply: "The human mind can only solve computable problems". Since the 1950's there has been a rapid development of computer science and its relation to logical formalisms. One problem area concerns the *complexity* of different kinds of computation. A difference between the analysis of algorithms and computational complexity theory is that an analysis of an algorithm for a particular problem can determine which amount of resources is used to solve the problem, whereas complexity theory asks a more general question about the minimal resources required among all possible algorithms that could be used to solve the same problem. The paper by Isaac et al. (this volume) is an overview of the consequences for cognitive science of the results concerning complexity and logical formalisms.

However, if one takes a different perspective on cognition, considerations concerning complexity will be of a different nature. The focus of this article will be the relation between complexity theory and situated and embodied cognition.

In order to bring out the contrast between the different kinds of cognitive science in relation to complexity, I want to highlight two assumptions of classical computationalism:

(1) All computation is (sequential) manipulation of symbols.
(2) The algorithms are run in a system (a computer or a brain) that is separated from the world—once the inputs are given to an algorithm it runs independently of what happens outside the system.

31.2 Complexity in Neural Networks

The second flavour of cognitive science is connectionism. In this tradition, Assumption (1), that all computation is manipulation of symbols, is abandoned. The neurons in a neural network are seen as processing information on the "subsymbolic" [27] or "subconceptual" [12] level. In general, connectionism kept Assumption (2), that computation is performed in a system that is separated from the world. In most applications, the neural network is given an input—in the form of a vector of values to its input layer—that is then processed by the system resulting in an output—a vector of values in its output layer.

However there are exceptions: In robotics, the *reactive systems* studied by Brooks [2] and others consist of comparatively simple processors, not necessarily parallel, that are in a constant interaction with the world. The research on reactive systems can be seen as precursors of the movement towards situated cognition. In these systems, it is no longer meaningful to separate input and output since they function as feedback loops, directly involving the surrounding world in its computations. Brooks [2] denies that a reactive system needs any internal representations at all. He takes the stance that robots do not need a model of the world to determine what to do next because they can simply sense it directly. He says that the world is its own best representation and that an efficient system should exploit this. However, his position has met with criticism (e.g. [18, 32]), even within the situated cognition camp.

As a part of the debate between classical computationalists and connectionists, it has been shown that all neural networks can be simulated by traditional computers (Turing machines) and vice versa. Hence many of the classical computationalists have claimed that the debate is a red herring. However, in these results complexity issues are eschewed.[1] Even though a Turing machine can simulate any neural network, it does not follow that the complexity of the algorithm for the Turing machine is of the same order as the one followed by the neural network.

Nowadays the area of complexity results concerning computation with neural networks is flourishing. A comprehensive survey is presented by Sima and Orponen [26]. They summarize the situation as that "a complexity theoretic taxonomy of neural networks has evolved, enriching the traditional repertoire of formal computational modes and even pointing out new sources of efficient computation" (p. 2728). However, one conspicuous lacuna in their survey is that the results they consider do not at all account for the *learning dynamics* of neural networks. This is, in my opinion, a serious limitation, since one of the main computational advantages of neural networks is that they can learn, albeit slowly, from the input they are presented with. Modelling such learning becomes much more difficult with classical symbolic computing.[2]

Isaac et al. (this volume) also discuss computation in neural networks, although their focus is on how systems for non-monotonic reasoning may be implemented. In particular they relate results in [20, 21] showing that any system performing computations over distributed representations may be interpreted as a classical computational system performing non-monotonic reasoning. These results support the view that anything that can be computed with a neural network can also be computed in a classical system.

31.3 Complexity in Situated Cognition

Next I turn to complexity issues in relation to situated cognition. The proponents of this position would claim that the brain is not made for checking the logical consistency of sentences or for handling any other NP-complete problem, but for surviving and reproducing in an environment that is partly predictable and partly unpredictable. The primary duty of the brain is to serve the body (the brain is a butler, not a boss). It does not function in solitude, but is largely dependent on the body it is employed by and the environment it is interacting with. In contrast, when

[1] For example, it is surprising that Marr [22] did not at all mention computational complexity in his description of the three levels of computation.

[2] There are attempts, however, in the work on adaptive Turing machines.

the brain is seen as a Turing machine, it has become customary to view it as an isolated entity, in accordance with Assumption (2) above.[3]

In addition to the two assumptions, the traditional complexity argument presumes that the problem is expressed in a representation where the primitive elements (the predicates) are independent of each other. This goes back to the ideals of logical positivism, in particular Carnap's [3] attempt to "reconstruct" the world in terms of atomic predicates. The position of situated cognition is, in contrast, that cognitive processes exploit (and mimic) the structures of the world itself, in particular the spatial layout of information.

Furthermore, situated cognition, at least partly, accepts the position that the world is its own best representation. As we saw, this is a central tenet of reactive systems [2]. Consequently, the brain does not need to construct detached representations of everything it interacts with.[4] Hence, situated cognition gives up both Assumptions (1) and (2) of classical computationalism. The position is succinctly formulated by Clark [6, p. 148]: "Structured, symbolic, representational, and computational views of cognition are mistaken. Embodied cognition is best studied by means of noncomputational and nonrepresentational ideas and explanatory schemes involving, e.g. the tools of Dynamical Systems theory".

In situated cognition, the visual system is not merely seen as an input device to the brain and the hand as enacting the will of the brain, but the eye-hand-brain is a coordinated system that functions as a feedback loop. For many tasks, it turns out that we solve problems more efficiently with our hands than with our brains. A simple example is the computer game Tetris where you are supposed to quickly turn, with the aid of the keys on the keyboard, geometric objects that come falling over a computer screen in order to fit them with the pattern at the bottom of the screen. When a new object appears, one can mentally rotate it to determine how it should be turned before actually touching the keyboard. However, expert players turn the object faster with the aid of the keyboard than they turn an image of the object in their brains [19]. This is an example of what has been called *interactive thinking*. The upshot is that a human who is manipulating representations in the head is sometimes a cognitively less efficient system than somebody interacting directly with the represented objects.

Clark [6, pp. 219–220] presents a fascinating example of a situated interaction between an organism and the world. It has been suggested that some aquatic animals, such as tunas and dolphins are simply not strong enough to propel themselves at the speeds they are observed to reach. Triantafullou and Triantafullou [29, p. 69] paradoxically claim that "it is even possible for a fish's swimming efficiency to exceed 100 %". The reason tunas and dolphins can be so efficient is that they in their swimming create and exploit swirls and vortices in the water that improve their propulsion and ability to manoeuever. In brief, the tunas and dolphins swim *with* the water, not in the water. The analogy I want to bring out is that our brains can be very

[3] This assumption is the basis for all sci-fi novels about a brain in the vat.

[4] In contrast to [2], the position does not deny, however, that the brain employs *some* detached representations, for example, when it is planning [13].

efficient, even with their limited resources, since they think *with* the world, not in the world.

It should be pointed out that ideas related to those of embodied and situated cognition that have become popular in the last decades have several predecessors. One example is the "ecological" psychology of Gibson [15] who rejected the idea that cognition is information processing and instead claimed that organisms could "pick up" all the necessary visual information directly from the environment and so that no computation was needed. Another tradition is the cybernetic movement in the middle of the 20th century (e.g. [31]) that studied feedback loops between an agent and the environment, again without exploiting any symbolic representations.

As far as I know, no strict account of the complexity of cognitive processes has been developed within the tradition of situated cognition. One reason for this is that it is difficult to develop formal models of how a situated approach influences complexity issues since we often do no know enough about what in the world the brain exploits directly and what it represents for itself.

One toy example, dear to researchers in classical AI, is how to determine whether a block x is above a block y in a tower of blocks (a typical robotics problem in the early days). In classical computation, this problem would be represented by a set of atomic statements of the type on(a, b), on(b, c), on(c, d)… and formulas expressing that the relation "above" is the transitive closure of "on". All this would be put into an inference engine that can determine the truth or falsity of above(x, y). The computational complexity of this problem is of the order n^2, where n is the number of blocks in the tower.

In contrast to the classical internal computation, a model within situated cognition would take into account that in the real world the blocks are *spatially organized* along the vertical dimension. The transitivity of the relation "above" is *built into* this spatial organization and need not be expressed in axioms, let alone be computed. A robot can simply visually scan the blocks from the bottom and see whether it encounters x or y first to determine the truth or falsity of above(x, y). The complexity of this procedure is of the order n, where n is the number of blocks, that is, it is linear in the number of blocks. The upshot is that *the geometric structure of the external word reduces the complexity of the problem*. This toy (sic) problem, illustrates in what sense the structure of the world helps offloading a cognitive system.[5]

More generally, one can consider the complexity of visual search problems. Tsotsos [30, p. 428] distinguishes between two variants: *bounded search* in which the visual target is explicitly provided in advance and *unbounded search* in which the target is defined only implicitly, for example, by specifying relationships it must have with other visual stimuli. He proves that unbounded visual search is NP-complete, while bounded visual search has linear complexity.

These theoretical results can be compared with the empirical results from Treisman [28] and her colleagues. In the experiments, two types of stimuli were used:

[5] In the terminology of Barwise and Shimojima's [1] "surrogate reasoning", this example is a "free ride" provided by the geometric constraints. However, the authors do not consider the reduction in complexity provided by "free rides".

disjunctive where the target can be identified by only one feature, such as color or orientation, and *conjunctive* where the target requires that more than one feature is identified. Both types are cases of bounded search in Tsotso's [30] terminology. Treisman [28] finds that conjunctive displays are identified in a response time that is linear with the number of items in the scene, just as predicted by Tsotsos' complexity result. However, for disjunctive stimuli, the target is found immediately—it simply "pops out"—independently of the number of items present. In this case, the human visual system somehow finds a solution that is more efficient in terms of complexity than what is predicted by Tsotsos' theoretical results.

31.4 Other Problems Relating to Complexity and Situated Cognition

In this section I will discuss complexity issues related to two well-known enigmas for classical computationalism in terms of situated cognition.

The first is the *frame problem* [7, 23]. Within the early AI community, it was hoped that if we could represent the knowledge necessary to describe the world and the possible actions in a suitable symbolic formalism, then by coupling this world description with a powerful inference machine one could construct an artificial agent capable of planning and problem solving. It soon turned out, however, that describing actions and their consequences in a symbolic form leads to a combinatorial explosion of the logical inferences that are needed. In other words, the complexity of the problem became insurmountable.

The crux is that symbolic representations are not well suited for representing causal connections or dynamic interactions in the world. Various escape routes were tried, but the frame problem persisted in one form or another. As a consequence, the entire program of building planning agents based on purely symbolic representations more or less came to a stall.

At the other extreme one finds the reactive systems that were presented earlier. Such systems are able to solve problems in the immediate environment without any symbolic representations simply by being directly situated in the world. On the other hand, reactive systems cannot form any plans that go beyond what is given in the environment.

Nowadays, many robotic systems take a middle road. They build up representations from their experience of the world, for example by constructing maps of their environment. Often, the representations are of a non-symbolic form. Some robots build on hybrid forms of representations, mixing symbols with maps and other non-symbolic forms (e.g. [4]). However, there exists no principled theory of how the computationally most efficient mixture between inner representations and immediate reactions to the environment is to be determined for a planning system. The problem is still in the hands of the engineers. Again, a suitable theory of the complexity of the problem is lacking.

A second enigma in the classical tradition is Chomsky's [5] *poverty of stimulus* argument, which claims that the grammar of a natural language cannot be learned by children because of the limited data available to them. In a more general form, this has become known as the *logical problem of language acquisition*, claiming that children cannot learn language simply by considering the input.[6] The argument can be structured as follows:

- All languages contain grammatical patterns that cannot be learned by children using *positive evidence* alone.
- Children are only presented with positive evidence for these patterns.
- Children learn the correct grammars for their native languages.

As a consequence, Chomsky argues, learning the grammar of a language must depend on some sort of *innate* linguistic capacity that provides additional knowledge to the children. In brief, language is too complex to be 100 % learned. Note that the logical problem of language acquisition presumes analogues of the assumptions (1) and (2), in particular that language processing is done independently of the world.[7]

From the perspective of situated cognition, a similar argument to the one presented in the previous section can be applied here. The key idea is that the child does not learn a language in the world, it learns it *with* the world, in particular together with other humans.

First of all, note that the problem of language acquisition, at least in Chomsky's version, does not concern how a language is learned, but how the *grammar* of a language is acquired. Formulating the problem in this fashion builds on the additional assumption that the grammar of a language is *independent* of its semantics (let alone, its pragmatics). However, outside the Chomskian congregation, this assumption is denied. Cognitive linguistics, for example, builds on the idea that the syntax of language is constrained, if not determined, by its semantics. And as soon as one then allows some connection between the semantics of a language and the world the language user is situated in, learning a grammar will, at least to some extent, be dependent on one's knowledge about the world.

Several experiments about language learning have shown how the learning of grammatical features exploits world knowledge (e.g. [10, 24]). For example Ramscar and Yarlett [24] show that children's world knowledge generates *expectations* about grammatical patterns. When such expectations are violated, for instance by an irregular plural form, the input can indeed function as negative evidence. In this way the argument from the poverty of stimulus is blocked.

Furthermore, a sentence is not just taken as input to the grammar crank in the child's brain and then determined to be grammatical or not—a sentence is *used* in a

[6] Several researchers have used Gold's [16] theorem to support this argument, but, as Johnson [17] shows, this result has little bearing on how people actually learn languages.

[7] Chomsky's early work concerned the relations between different kinds of formal automata and the (formal) languages they could identify. This is a typical problem of computationalism that builds on Assumptions (1) and (2). Chomsky seems, more or less, to have stuck to these assumptions throughout his career.

particular context. And the use of a sentence may provide constraints for its structure. Here, I do not wish to speculate on how the constraints can be specified. Suffice it to notice that such constraints will block the poverty of stimulus argument, at least in its current form.

31.5 Conclusion

In conclusion, Assumptions (1) and (2) of classical computationalism have been taken over implicitly in many other areas. Once they are brought out into the open, however, they are seen to be invalid for many kinds of cognitive problems. The main argument of this paper is that once we give up these assumptions, many problems that have seemed hopelessly complex for the classical computationalist may become more manageable, if a connectionist or situated perspective on cognition is adopted instead. And evolution is a tinkerer with limited resources: rest assured that if one solution to a problem is cheaper than another, evolution will, in the long run, select the cheap one.

Still, humans have evolved symbolic language. In my opinion [13, 14], the main reason for this is that it has improved our *planning capacities*. There are situations involving reasoning with numbers, reasoning with cases or reasoning with conditional assumptions where symbolic structures are required. My point in this paper is simply that there are cases of problem solving where less complex methods than those offered by symbolic thinking are sufficient and therefore more efficient.

Humans have also speeded up the evolutionary selection processes: We have created cultures and artefacts that greatly improve our problems solving capacities. We have invented pencil and paper, libraries and smartphones that offload our memories, allow us to share knowledge, and amplify our calculations. Tunas and dolphins create structures in the water that improve their swimming. Humans create structures in the world that improve their thinking. As Clark [6, p. 180] puts it: "Our brains make the world smart so that we can be dumb in peace! Or to look at it another way, it is the human brain *plus* these chunks of external scaffolding that finally constitutes the smart, rational inference engine we call mind".

It must be pointed out, though, that the theory of situated cognition still lacks a rigor that would make it possible to develop a parallel to the theory of complexity that exists for classical computationalism and to some extent also for connectionism. Barwise and Shimojima's [1] ideas about *constraint projection* is perhaps a first step in that direction. I can only hope that a more precise theory will be formulated that will allow comparisons with the results concerning the complexity of situated processes.

Acknowledgments I gratefully acknowledge support from the Swedish Research Council for the Linnaeus environment *Thinking in Time: Cognition, Communication and Learning*. Thanks to Holger Andreas, Johan van Benthem, Alistair Isaac, Giovanni Pezzulo and Jakub Szymanik for helpful comments on an earlier version of the paper.

References

1. Barwise J, Shimojima A (1995) Surrogate reasoning. Cogn Stud Bull Japan Cogn Sci Soc 2(4):7–27
2. Brooks R (1991) Intelligence without representation. Artif Intell 47:139–159
3. Carnap R (1928) Der logische Aubau der Welt. Felix Meiner Verlag, Hamburg
4. Chella A, Frixione M, Gaglio S (1997) A cognitive architecture for artificial vision. Artif Intell 89:73–111
5. Chomsky N (1980) Rules and representations. Basil Blackwell, Oxford
6. Clark A (1997) Being there: putting brain body and world together again. MIT Press, Cambridge
7. Dennett DC (1984) Cognitive wheels: the frame problem of AI. In: Hookway C (ed) Minds, machines and evolution: philosophical studies. Cambridge University Press, Cambridge, pp 129–151
8. Fodor JA (1975) The language of thought. Harvard University Press, Cambridge
9. Fodor JA, Pylyshyn ZW (1988) Connectionism and cognitive architecture: a critical analysis. Cognition 28:3–71
10. Foraker S, Regier T, Khetarpal N, Perfors A, Tenenbaum J (2009) Indirect evidence and the poverty of the stimulus: the case of anaphoric "one". Cogn Sci 33:287–300
11. Gärdenfors P (1999) Cognitive science: from computers to anthills as models of human thought. World Social Science Report, UNESCO Publishing, Paris, pp 316–327
12. Gärdenfors P (2000) Conceptual spaces: the geometry of thought. MIT Press, Cambridge
13. Gärdenfors P (2004) Cooperation and the evolution of symbolic communication. In: Oller K, Griebel U (eds) The evolution of communication systems. MIT Press, Cambridge, pp 237–256
14. Gärdenfors P (2012) The cognitive and communicative demands of cooperation. In: van Eijck J, Verbrugge R (eds) Games, actions and social software, LNCS 7010. Springer, Berlin, pp 164–183
15. Gibson JJ (1979) The ecological approach to visual perception. Houghton Mifflin, Boston
16. Gold E (1967) Language identification in the limit. Inf Control 10:447–474
17. Johnson K (2004) Gold's theorem and cognitive science. Philos Sci 71:571–592
18. Kirsh D (1991) Today the earwig, tomorrow man? Artif Intell 47:161–184
19. Kirsh D, Maglio P (1994) On distinguishing epistemic from pragmatic action. Cogn Sci 18:513–549
20. Leitgeb H (2001) Nonmonotonic reasoning by inhibition nets. Artif Intell 128:161–201
21. Leitgeb H (2003) Nonmonotonic reasoning by inhibition nets II. Int J Uncertainty Fuzziness Knowl Based Syst 11:105–135
22. Marr D (1982) Vision: a computational investigation into the human representation and processing visual information. Freeman, San Francisco
23. McCarthy J, Hayes PJ (1969) Some philosophical problems from the standpoint of artificial intelligence. In: Michie D, Meltzer B (eds) Machine intelligence, vol 4. Edinburgh University Press, Edinburgh, pp 164–183
24. Ramscar M, Yarlett D (2007) Linguistic self-correction in the absence of feedback: a new approach to the logical problem of language acquisition. Cogn Sci 31:927–960
25. Rumelhart DE, McClelland JL (1986) Parallel distributed processing: explorations in the microstructure of cognition. MIT Press, Cambridge
26. Sima J, Orponen P (2003) General purpose computation with neural networks: a survey of complexity theoretic results. Neural Comput 15:2727–2778
27. Smolensky P (1988) On the proper treatment of connectionism. Behav Brain Sci 11:1–23
28. Treisman A (1988) Features and objects: the fourteenth Bartlett memorial lecture. Q J Exp Psychol Sect A: Hum Exp Psychol 40:201–237
29. Triantafyllou MS, Triantafyllou GS (1995) An efficient swimming machine. Sci Am 272:64–70
30. Tsotsos J (1990) Analyzing vision at the complexity level. Behav Brain Sci 13:423–469
31. Wiener N (1948) Cybernetics: or control and communication in the animal and the machine. Hermann and Cie, Paris
32. Williams MA (2008) Representation = grounded information. In: Proceedings of the 10th Pacific Rim international conference on artificial intelligence, Hanoi, pp 473–484