

The *SP Challenge*: that the *SP System* is more promising as a foundation for the development of human-level broad AI than any alternative

J Gerard Wolff*

December 2, 2021

Abstract

The *SP Challenge* is the deliberately provocative theme of this paper: that the SP System (SPS), meaning the *SP Theory of Intelligence* and its realisation in the *SP Computer Model*, is more promising as a foundation for the development of human-level broad AI, aka ‘artificial general intelligence’ (AGI), than any alternative. In that connection, the main strengths of the SPS are: 1) The adoption of a top-down, breadth-first research strategy with wide scope; 2) Recognition of the importance of information compression (IC) in human learning, perception, and cognition—and, correspondingly, a central role for IC in the SPS; 3) The working hypothesis that all kinds of IC may be understood in terms of the matching and unification of patterns (ICMUP); 4) A resolution of the apparent paradox that IC may achieve decompression as well as compression. 5) The powerful concept of SP-multiple-alignment, a generalisation of six other variants of ICMUP; 6) the clear potential of the SPS to solve 19 problems in AI research; 7) Strengths and potential of the SPS in modelling several aspects of intelligence, including several kinds of probabilistic reasoning, versatility in the representation and processing of AI-related knowledge, and the seamless integration of diverse aspects of intelligence, and diverse kinds of knowledge, in any combination; 8) Several other potential benefits and applications of the SPS; 9) In ‘SP-Neural’, abstract concepts in the SPS may be mapped into putative structures expressed in terms of neurons and their interconnections and intercommunications; 10) The concept of ICMUP provides an entirely novel perspective on the foundations of mathematics; 11) How

*Dr Gerry Wolff, BA (Cantab) PhD (Wales) CEng MBCS MIEEE, ORCID ID 0000-0002-4624-8904. CognitionResearch.org, Menai Bridge, UK; jgw@cognitionresearch.org; +44 (0) 1248 712962; +44 (0) 7746 290775; *Twitter*: @gerrywolff65; *Web*: www.cognitionresearch.org.

to make generalisations from data, including the correction of over- and under-generalisations, and how to reduce or eliminate errors in data. There is discussion of how the SPS compares with some other potential candidates for the SP-Challenge. And there is an outline of possible future directions for the research.

Keywords: SP Theory of Intelligence, SP Computer Model, artificial general intelligence; information compression, SP-multiple-alignment, SP-Neural, foundations of mathematics, generalisation.

1 Introduction

The deliberately provocative theme of this paper (the ‘SP-Challenge’) is the suggestion that: the *SP System* (SPS), meaning the *SP Theory of Intelligence* and its realisation in the *SP Computer Model* (SPCM), is more promising as a foundation for the development of human-level broad AI, aka ‘artificial general intelligence’ (AGI), than any alternative.

It will be good if other researchers pick up the gauntlet and make a case for alternative systems. This may help to highlight what may be needed to reach AGI.

1.1 Origin and development

In accordance with the top-down breadth-first research strategy with wide scope, described in Section 2, the SP research programme began with a search for a framework with the breadth and scope to provide explanatory and descriptive concepts in AI and in human learning, perception, and cognition (HLPC). It was only later that mainstream computing and mathematics were added to the scope of the search.

Early work began developing a computer model for finding full matches and good partial matches between sequences of symbols. This was because of a hunch that that kind of full and partial matching lay behind many aspects of intelligence, and also because there is a clear connection with information compression (IC), and there are reasons to believe that IC is important in the workings of brains and nervous systems (Section 3).

Finding good full and partial matches between two sequences led to the idea that finding good full and partial matches between two *or more* sequences might be important. That concept known as ‘multiple sequence alignment’ in bioinformatics is illustrated in Figure 1. It began to become clear how that concept might be adapted to model different aspects of intelligence.

Although that insight only took a few hours to take shape, a much longer period of development—about 17 years—has been needed to create the concept of

SP-multiple-alignment SPMA within the SPCM, to bring it to its relatively mature state now, to explore how the SPMA, within the SPCM, might exhibit different aspects of intelligence (Section 8), and to write a book about the research: *Unifying Computing and Cognition* [30].

Later again, after a break of nearly seven years, there was more work exploring potential benefits and applications of the SPS, including several outside the realm of AI (Section 9). Later again, research has shown the clear potential of the SPS to solve 19 problems in AI research, 17 of them described by influential experts in AI (Section 7).

Without me wishing to claim the status of ‘genius’, it seems appropriate here to quote Thomas Edison’s saying that “Genius is one percent inspiration and ninety-nine percent perspiration.”

Why was there so much ‘perspiration’ in the SP research? Most of the work was developing *hundreds* of different versions of the SPCM, each one developing a possible idea for how the SPMA or other part of the SPCM might work—and applying relevant tests, taking notes of the results, and exploring ideas for what might be tried next.

1.2 Sources of information

The SPS is described most fully in the book *Unifying Computing and Cognition* [30], quite fully but more briefly in a shortened version of the book [32], and in outline in Appendix A.

Any of those sources, including Appendix A, should provide readers with sufficient information to understand the rest of the paper.

1.3 The name ‘SP’

Since people often ask, the origin of the name ‘SP’ is explained in Appendix A.2. It is intended that ‘SP’ should be treated as a name like ‘BBC’ or ‘IBM’, not an abbreviation.

1.4 Presentation

To be clear, this paper is about theory, it is not reporting experiments with sections like ‘Method’, ‘Results’, and ‘Discussion’. The paper is about the merits of the SP System as a foundation for the development of AGI.

Because the achievement of AGI is far into the future, most of the arguments are probabilistic. Nevertheless, the paper is unambiguously scientific, in the same way that probabilistic thinking was needed to nail down such copper-bottomed

theories as the theory of evolution by natural selection, the theories of special and general relativity, and more.

The strengths and potential of the SPS are described in the following sections:

- The benefits of a top-down, breadth-first research strategy with wide scope (Section 2).
- The central importance of information compression in the workings of brains and nervous systems, and in the SPS (Section 3).
- ‘Information Compression via the Matching and Unification of Patterns’ (ICMUP) (Section 4).
- A resolution of the apparent paradox that IC may achieve decompression as well as compression (Section 5).
- A major breakthrough: the discovery and development of the powerful concept of *SP-multiple-alignment* (Section 6).
- A major result: the potential of the SPS to solve 19 problems in AI research (Section 7).
- The strengths and potential of the SPS in modelling aspects of intelligence (Section 8). Here, the main sections are:
 - Several kinds of intelligent behaviour (Section 8.1);
 - Several kinds of reasoning (Section 8.2).
 - The representation and processing of several kinds of AI-related knowledge (Section 8.3).
 - The seamless integration of diverse aspects of intelligence, and diverse kinds of knowledge, in any combination, Section 8.4.
- Other potential benefits and applications of the SPS (Section 9).
- SP-Neural, a ‘neural’ version of the SPS (Section 10).
- How to make generalisations without either over-generalisations (under-fitting) or under-generalisations (over-fitting); and how to minimise corrupted or ‘dirty’ data (Section 11).
- The concept of ICMUP provides an entirely novel perspective on the foundations of mathematics (Section 12).

- Advantages of the SPS compared with other potential candidates for the SP-Challenge (Section 13).
- Future directions for the SP research (Section 14).

Abbreviations used in this paper are listed in Appendix B.

2 The benefits of a top-down, breadth-first research strategy with wide scope

Unlike most other research in AI or cognitive science, the SPS has been developed via a top-down, breadth-first research strategy of wide scope.

This should help to meet the concerns of Gary Marcus and Ernest Davis when they wrote: “What’s missing from AI today—and likely to stay missing, until and unless the field takes a fresh approach—is broad (or “general”) intelligence.” [14, p. 15] (see also writings of Allen Newell and Pamela McCorduck in Section 13).

In keeping with what Marcus and Davis say, the scope of the SP research strategy is unusually wide, *aiming for a simplification and integration of observations and concepts across AI, mainstream computing, mathematics, and ‘human learning, perception, and cognition’*.

This wide scope in research strategy is already yielding breadth in the strengths and potential of the SPS, as described in Section 7 (the clear potential of the SPS to solve 19 problems in AI research), Section 8 (the strengths and potential of the SPS in several aspects of intelligence), and Section 9 (the clear potential of the SPS in several areas of application apart from AI).

3 The central importance of information compression in the workings of brains and nervous systems, and in the SPS

A potent idea, pioneered by Fred Attneave [1, 2], Horace Barlow [3, 4], and others, is that much of the workings of brains and nervous systems may be understood as information compression (IC). This strand of research has been pursued by various researchers up to the present. Evidence for the importance of IC in HLPC, is reviewed in [41].

The importance of IC in HLPC became central in a programme of research developing computer models of the learning of a first language by children [28], and it is bedrock in the SPS.

In the SP programme of research, a working hypothesis is that all kinds of IC may be understood as the Matching and Unification of Patterns (ICMUP). There is more about this in Section 4.

In these connections, it is of some interest that, as far back as 1969, Barlow wrote with some prescience:

“... the operations needed to find a less redundant code have a rather fascinating similarity to the task of answering an intelligence test, finding an appropriate scientific concept, or other exercises in the use of inductive reasoning. Thus, redundancy reduction may lead one towards understanding something about the organization of memory and intelligence, as well as pattern recognition and discrimination.” [4, p. 210].

where “find[ing] a less redundant code” is at the heart of IC.

The importance of IC for the SP-Challenge is as follows:

- In view of the evidence for IC as a unifying theme in HLPC [41], it seems clear that IC should be central in the workings of any system that aspires to AGI, as it is in the SPS;
- The central role for IC in the SPCM—mediated by the concept of SP-multiple-alignment (Appendix A.4 and Section 6)—is largely responsible for: the potential of the SPS to solve several problems in AI research (Section 7); the strength and potential of the SPS in modelling aspects of intelligence (Section 8); and the potential of the SPS in several other areas of application (Section 9).
- And IC also has important roles: as a unifying principle in mathematics (Section 12); in the formation of generalisations without over-fitting or under-fitting, and in the weeding out of ‘dirty data’ (Section 11); and in the analysis and production of data (Section 5).
- In both natural and artificial systems:
 - For a given body of information, **I**, to be stored, IC means that a smaller store is needed. Or for a store of a given capacity, IC facilitates the storage of a larger **I** [41, Section 4];
 - For a given body of information, **I**, to be transmitted along a given channel, IC means an increase in the speed of transmission. Or for the transmission of **I** at a given speed, IC means a reduction in the bandwidth which is needed [41, Section 4].

- Because of the intimate relation between IC and concepts of inference and probability (Appendix A.6), and because of the central role of IC in the SPS, the SPS is intrinsically probabilistic. Correspondingly, it is relatively straightforward to calculate absolute and relative probabilities for all aspects of intelligence exhibited by the SPS, including several kinds of reasoning ([30, Section 3.7], [32, Section 4.4]), in keeping with the probabilistic nature of human inferences and reasoning.

4 Information compression via the matching and unification of patterns

A working hypothesis in the SP research is that all kinds of IC may be understood as ‘Information Compression via the Matching and Unification of Patterns’ (ICMUP).

In this research, seven main variants of ICMUP are recognised [42, Sections 5.1 to 5.7]:

- *Basic ICMUP.* Two or more instances of any pattern may be merged or ‘unified’ to make one instance [42, Section 5.1].
- *Chunking-with-codes.* Any pattern produced by the unification of two or more instances is termed a ‘chunk’. A ‘code’ is a relatively short identifier for a unified chunk which may be used to represent the unified pattern in each of the locations of the original patterns [42, Section 5.2].
- *Schema-plus-correction.* A ‘schema’ is a chunk that contains one or more ‘corrections’ to the schema. For example, a menu in a restaurant may be seen as a schema that may be ‘corrected’ by a choice of starter, a choice of main course, and a choice of pudding [42, Section 5.3].
- *Run-length coding.* In run-length coding, a pattern that repeats two or more times in a sequence may be reduced to a single instance with some indication that it repeats, or perhaps with some indication of when it stops, or even more precisely, with the number of times that it repeats [42, Section 5.4].
- *Class-inclusion hierarchies.* Each class in a hierarchy of classes represents a group of entities that have the same attributes. Each level in the hierarchy *inherits* all the attributes from all the classes, if any, that are above it [42, Section 5.5].
- *Part-whole hierarchies.* A part-whole hierarchy is similar to a class-inclusion hierarchy but it is a hierarchy of part-whole groupings [42, Section 5.6].

- *SP-multiple-alignment*. The SPMA concept [42, Section 5.7] is described in Appendix A.4 and Section 6.

This list probably does not exhaust the possible variants of ICMUP, but they are the ones that have received most attention so far in the SP programme of research.

5 A resolution of the apparent paradox that IC may achieve decompression as well as compression

It is sometimes said that IC as a central feature of the SPS conflicts with the undoubted fact that people can and do produce information as well as compress it, both in ordinary speech or writing and also in creative areas like creative writing, painting, the composition of music, and so on.

In that connection, an interesting feature of the SPCM is that processes for the analysis of New information are *exactly* the same as may be used for the production of information. For example, with natural language, processes for the production of a sentence are, without any qualification, the same as may be used for the analysis of the same sentence.

Since the SPCM works by compressing information, this feature of the SPCM looks, paradoxically, like “decompression of information by compression of information”.

How the whole system works, and how this paradox may be resolved, is explained in [32, Section 4.5] and [30, Section 3.8].

Although this aspect of the SPS has not yet been explored, there is potential in the SPCM for the creation of entirely new structures which may be seen as novel or creative, but not necessarily artistic.

6 A major breakthrough: the discovery and development of the powerful concept of SP-multiple-alignment

As indicated in Section 1.1, a major discovery in the SP programme of research is the powerful concept of *SP-multiple-alignment* (SPMA) (Appendix A.4). The SPMA concept is the last of seven variants of ICMUP described in Section 4 and in [42, Section 5].

The paper [47] shows in detail how the SPMA concept may be seen to be a generalisation of the other six variants of ICMUP. That generality provides an explanation for the strengths and potential of the SPS as described in Sections 7, 8, and 9.

Bearing in mind that it would be just as bad to downplay any feature of the SPS as would over-selling any aspect of the system, *the SPMA concept promises to be as significant for our understanding of intelligence as is DNA for many aspects of biology. The SPMA concept may prove to be the ‘double helix’ of intelligence.*

6.0.1 Discovery

As indicated in Section 1.1, the main inspiration for the SPMA concept was the concept of ‘multiple sequence alignment’ in bioinformatics. This is an arrangement of two or more symbolic representations of DNA sequences (as illustrated in Figure 1), or two or more symbolic representations of sequences of amino-acid residues.

In Figure 1, there are five DNA sequences which have been arranged alongside each other, and then, by judicious ‘stretching’ of one or more of the sequences in a computer, symbols that match each other across two or more sequences have been brought into line.

A ‘good’ multiple sequence alignment is one in which a relatively large number of symbols are brought into line.

```

      G G A      G      C A G G G A G G A      T G      G      G G A
      | | |      |      | | | | | | | | | |      | |      | | | |
      G G | G      G C C C A G G G A G G A      | G G C G      G G A
      | | |      | | | | | | | | | | | |      | |      | | | |
A | G A C T G C C C A G G G | G G | G C T G      G A | G A
      | | |      | | | | | | | | | |      | |      | | | |
      G G A A      | A G G G A G G A      | A G      G      G G A
      | | |      | | | | | | | | | |      | |      | | | |
      G G C A      C A G G G A G G      C      G      G      G G A

```

Figure 1: A ‘good’ multiple sequence alignment amongst five DNA sequences.

Because the abstract space of possible multiple sequence alignments is so large, it is necessary to use heuristic techniques to search that abstract space in stages, discarding all but the best partial alignments at the end of each stage.

This kind of technique trades accuracy for speed, but in many cases, acceptably good alignments, like the one shown in Figure 1, may be found within a reasonable time.

6.0.2 Development

The bioinformatics concept of ‘multiple sequence alignment’ (Section 6.0.1) became the powerful concept of *SP-multiple-alignment* (SPMA), described briefly in Appendix A.4 and more fully in Section 6.0.3.

As noted in Section 1.1, a lengthy period has been needed to create the SPMA construct within the SPCM, and to explore what the SPCM can do as described in Sections 7, 8, and 9.

6.0.3 Example

An example of the SPMA concept is shown in Figure 2.




Figure 2: The best SPMA created by the SPCM with a store of Old SP-patterns like those in rows 1 to 8 (representing grammatical structures, including words) and a New SP-pattern, '(t w o k i t t e n s p l a y)', shown in row 0 (representing a sentence to be parsed). Adapted from Figure 1 in [31], with permission.

In the figure, row 0 shows an *SP-pattern* (a sequence of *SP-symbols*) that is designated ‘New’ because we imagine that it has just been received from its environment via the system’s ‘senses’, as illustrated in Figure 7. Clearly, in Figure 2, it corresponds with the sentence “Two kittens play”.

Each of rows 1 to 8 in the SPMA shows a single SP-pattern representing a grammatical structure which is designated ‘Old’ because it comes from the system’s store of Old SP-patterns in the system’s ‘head’ (Figure 7).

As with multiple sequence alignments like the one shown in Figure 1, it is necessary in the building of SPMA’s to use heuristic search in several stages, weeding out the bad partial structures at the end of each stage and retaining the good structures.

The structure in Figure 2 is considered ‘good’ because it enables the New SP-pattern to be encoded economically in terms of the Old SP-patterns in the alignment. How that encoding is done is described in [32, Section 4.1] and [30, Section 3.5].

In Figure 2, the whole SPMA may be seen to achieve the effect of parsing the sentence into its parts and subparts, much as is done in linguistics, except that linguists use trees for parsing, not SPMA’s.

The way that the SPMA in Figure 2 imitates the kind of tree used by linguists is just one of many examples of the versatility of SPMA’s in modelling other structures and their workings (Section 8).

7 A major result: the potential of the SPS to solve 19 problems in AI research

Strong support for the SPS has arisen, indirectly, from the book *Architects of Intelligence* by science writer Martin Ford [7]. To prepare for the book, he interviewed several leading experts in AI to hear their views about AI research, including opportunities and problems in the field:

“The purpose of this book is to illuminate the field of artificial intelligence—as well as the opportunities and risks associated with it—by having a series of deep, wide-ranging conversations with some of the world’s most prominent AI research scientists and entrepreneurs.” Martin Ford [7, p. 2].

In the book, Ford reports what the AI experts say, giving them the opportunity to correct errors he may have made so that the text is a reliable description of their thinking.

This source of information has proved to be very useful in defining problems in AI research that influential experts in AI deemed to be in need of solutions. This has been significant from the SP perspective because, with 17 of those problems and two others—19 in all—there is clear potential for the SPS to provide a solution.

Since these are problems with broad significance, not micro-problems with only narrow significance, the potential of the SPS to solve them is a major result from the SP programme of research, demonstrating some of the power of the SPS.

The paper [45] describes those 19 problems and how the SPS may solve them. In the following summary, ‘**P**’ signals a statement of the problem, ‘**SP**’ summarises how the SPCM may solve it, and ‘**DNN**’ or ‘**Other**’ adds relevant comments, if any, about DNNs or other technology:

1. *The symbolic versus sub-symbolic divide.* **P**: The need to bridge the divide between symbolic and sub-symbolic kinds of knowledge and processing.

SP: Although the SPCM has a symbolic flavour, the granularity of the system is not defined, so that an SP-symbol may be as large as a word or as small as a pixel.

DNN: DNNs are generally seen as sub-symbolic. They do not lend themselves well to symbolic styles of processing (see ‘Compositionality’ below).

2. *Errors in recognition.* **P**: The tendency of deep neural networks (DNNs) to make large and unexpected errors in recognition.

SP and **DNN**: The SPCM as it is now does not suffer from this problem, and there are reasons to believe that future versions will not suffer from it either.

3. *Natural languages.* **P**: The need to strengthen the representation and processing of natural languages, including the understanding of natural languages and the production of natural language from meanings.

SP: The SPCM has strengths and potential in both the representation and processing of natural languages. In particular, it provides for the succinct representation of the kinds of structures recognised by linguists. Preliminary investigations show how the SPMA within the SPCM may model the ‘understanding’ of natural language and the production of natural language from ‘meanings’ [30, Section 5.7].

DNN: Although DNNs have been applied successfully to the processing of natural languages, in for example, machine translation, those good results are achieved by a combination of the machine’s ability to recognise patterns and human abilities to divide one language into sections that have meaningful relations to human-identified section in another language. Otherwise, DNNs

do not provide well for the succinct representation of the kinds of structures recognised by linguists. (see ‘Compositionality’ below).

4. *Unsupervised learning.* **P:** Overcoming the challenges of unsupervised learning. Although DNNs can be used in unsupervised mode, they seem to lend themselves best to the supervised learning of tagged examples.

SP: In the SPCM, IC provides a metric for all aspects of intelligence including unsupervised learning. A working hypothesis is that unsupervised learning is the most general form of learning which may be adapted for other kinds of learning such as supervised learning, reinforcement learning, learning by imitation, and so on.

DNN: The minor role for IC in DNNs—there are only a few brief discussions of IC in Schmidhuber’s overview of DNN research [24]—seems to be why it is generally overlooked as an appropriate metric for unsupervised learning.

5. *Generalisation.* **P:** The need for a coherent account of generalisation, under-generalisation (over-fitting), and over-generalisation (under-fitting).

SP: A solution to this issue, and the related one of how to deal with corruption in the source data (‘dirty data’), was developed for models of language learning [28, p. 183] and further refined for the SPCM [32, Section 5.3]. It is described in outline in Section 11.

DNN: The minor role for IC in DNNs, noted above, seems to be why it has been overlooked as an appropriate basis for generalisation and the correction of dirty data.

6. *One-shot learning.* **P:** How to learn usable knowledge from a single exposure or experience.

SP: This is intrinsic to how the SPCM has been designed. No special provision is required.

DNN: Because learning in DNNs means a progressive strengthening of links between layers, and because such learning only becomes useful after several cycles of such strengthening, DNNs are intrinsically incapable of modelling one-shot learning.

7. *Transfer learning.* **P:** How to achieve transfer learning, incorporating old knowledge in new.

SP: Again, this is intrinsic to how the SPCM works, because New information may be combined freely with Old information, whilst preserving both of them accurately. For example, if there is an Old SP-pattern ‘d e l i c i o u s a p p l e p i e’ and a New SP-pattern is ‘u n a p p e t i s

i n g a p p l e p i e', the system will normally create SP-patterns like this: 'A a p p l e p i e #A', 'd e l i c i o u s A #A', and 'u n a p p e t i s i n g A #A'.

Notice how, in this example, the pair of symbols, 'A #A', serves as a 'code' in the chunking-with-codes techniques for IC (Section A.3).

DNN: In any one DNN, new learning wipes out or corrupts old learning (see 'Catastrophic forgetting', below), so it is difficult or impossible to preserve the old learning accurately.

8. *Reducing computational demands.* **P:** How to increase the speed of learning in AI systems, and how to reduce the demands of AI learning for large volumes of data, and for large computational resources.

SP: The paper [43] discusses how the SPCM, relative to DNNs, can speed up learning, reduce the volumes of data needed for learning, and reduce the computational resources needed for learning. In all three areas, the potential appears to be large.

DNN: Major problems for DNNs are slow learning and demands for very large volumes of data and very large computational resources, and it is not clear how these problems may be overcome.

9. *Transparency.* **P:** The need for transparency in the representation and processing of knowledge.

SP: In the SPCM: all knowledge in the system is open to inspection, and is likely to be in forms such as class-inclusion hierarchies which are familiar to people; and the system creates an audit trail for all its processing which is intelligible to people. Those aspects of transparency in the SPCM are described and discussed in [48].

DNN: Major problems with DNNs are: their lack of transparency in the way knowledge is represented; and their lack of transparency in how processing is done.

10. *Probabilistic reasoning.* **P:** How to achieve probabilistic reasoning that integrates with other aspects of intelligence.

SP: The SPCM supports several different kinds of reasoning (Section 8.2), and because of the intimate relation between IC and concepts of inference and probability (Appendix A.6), it is relatively straightforward to calculate absolute and relative probabilities for all aspects of intelligence exhibited by the SPCM (Section 8), including several forms of reasoning (Section 8.2).

DNN: As noted above (‘The symbolic versus sub-symbolic divide’), DNNs do not lend themselves well to symbolic styles of processing (see also ‘Compositionality’, below).

11. *Top-down strategies.* **P:** The need to re-balance research towards top-down strategies.

SP: As noted in Section 2, the SP programme of research is driven by a top-down, breadth-first strategy with wide scope. It appears that this is largely responsible for the strengths and potential of the SPCM across several aspects of intelligence, as summarised in Section 8.

DNN: Successes with DNNs seems to have created a conceptual trap, because adopting a top-down strategy may mean stopping work, at least temporarily, in an area which is providing publishable results.

12. *Self-driving vehicles.* **P:** How to minimise the risk of accidents with self-driving vehicles.

SP: The ‘SP’ solution to problems of generalisation and dirty data, described in Section 11, appears to be more robust and coherent than any of the solutions proposed in the literature on self-driving vehicles. Hence, that SP solution may help to minimise the risk of accidents in self-driving vehicles.

Other: It seems that there is little in the way of a coherent theory in the development of software for self-driving vehicles. The proposed solution to the problems of generalisation and dirty data may help to plug the gap.

13. *Compositionality.* **P:** The need for strong compositionality in the structure of knowledge.

SP: Since unsupervised learning in the SPCM conforms to the ‘DONSVIC’ principle (‘the Discovery Of Natural Structures Via Information Compression’, [32, Section 5.2]), that learning is likely, with most kinds of data, to develop strong compositionality in what it learns (see also item ‘Natural language’, above, and [48]). In general, the SPS supports compositionality, as can be seen in the SPMA in Figure 2.

DNN: In general, DNNs do not provide a robust means of expressing compositional structures such as class-inclusion hierarchies or part-whole hierarchies.

14. *Commonsense reasoning and commonsense knowledge.* **P:** The challenges of commonsense reasoning and commonsense knowledge.

SP: Preliminary investigation with the SPCM [40], suggests that it has potential in modelling commonsense reasoning and commonsense knowledge, as described in [6].

DNN: DNNs show little promise of providing solutions in this difficult area.

15. *Information compression.* **P:** Establishing the key importance of IC in AI research.

SP: This paper, and others in the SP programme of research (www.cognitionresearch.org/sp.htm) all help to confirm the importance of IC in AI and related research (see also Section 3).

DNN: As noted above, IC is acknowledged as having some role in DNNs but there appears to be no recognition of the central importance of IC in the workings of brains and nervous systems [41], or why, for that reason, IC should be central in any system that aspires to be an AGI.

16. *A biological perspective.* **P:** Establishing the importance of a biological perspective in AI research.

SP: The biological perspective of the SPCM is, for three reasons, connected with the central importance of IC in the SPCM: 1) Evidence for the importance of IC in the workings of brains and nervous systems (Section 3); 2) A recognition that, in all animals including humans, it is almost inevitable that IC would have been favoured by natural selection for reasons related to storage and bandwidth outlined in Section 3; and 3) The intimate connection between IC and concepts of inference and probability, and the importance of those things in natural selection (*ibid.*)

DNN: Some researchers have drawn parallels between the hierarchical organisation of neurons discovered by Hubel and Wiesel [9] and comparable organisation in the workings of ‘convolutional’ DNNs. There is more detail in [45, Section 18].

17. *Distributed versus localist knowledge.* **P:** Establishing whether or not knowledge in the brain is represented in ‘distributed’ or ‘localist’ form.

SP: The way knowledge is organised in the SPCM suggests that in SP-Neural—the ‘neural’ version of the SPCM—knowledge should be stored in localist form [37]. This is in line with evidence for ‘grandmother’ cells or clusters of such cells [22, 5, 23], so that, following a stroke, or with the progressive worsening of brain functions in dementia, a person may indeed lose the ability to recognise a grandmother, or anyone else who is close to them.

DNN: DNNs provide a working model of how distributed knowledge may work. This may be seen as a positive feature for anyone inclined to believe a distributed model of how the brain works. But it may be seen as a negative feature if the weight of evidence favours a localist model.

18. *Adaptation.* **P:** How to bypass the limited scope for adaptation in DNNs.

SP: The SPCM provides for the creation of an unlimited number of SP-patterns, and it builds SP-multiple-alignments with whatever number of rows is needed to model patterns of redundancy in the sensory data.

DNN: By contrast, although DNNs can adapt by changing the strengths of connections between layers, for any one DNN the number of layers is fixed, and the scope for changing the strengths of connections is limited.

19. *Catastrophic forgetting.* **P:** How to eliminate the problem of catastrophic forgetting: for any one DNN, new learning wipes out or corrupts old learning.

SP: In the SPCM, there is no restriction on the number of different SP-patterns that may be formed, so that new learning can be entirely independent of old learning. At the same time, it is entirely possible for new learning to incorporate whatever old learning is appropriate (see ‘Transfer learning’, above).

DNN: Within the DNN framework, there seems to be limited scope for solving this problem, except perhaps for the somewhat inelegant move of providing a new DNN for each of the concepts one may wish to learn.

Judging by the foregoing summary of the potential of the SPS to solve 19 problems in AI research, the evidence favours the SPS compared with DNNs.

8 The strengths and potential of the SPCM in modelling aspects of intelligence

The strengths and potential of the SPS in AI-related functions and structures are summarised in the subsections that follow. Further information may be found in [32, Sections 5 to 12] and [30, Chapters 5 to 9], and in other sources referenced below.

8.1 Several kinds of intelligent behaviour

The SPS has strengths and potential in the following aspects of intelligence: unsupervised learning; the analysis and production of natural language; pattern recognition that is robust in the face of errors in data; pattern recognition at multiple

levels of abstraction; computer vision [33]; best-match and semantic kinds of information retrieval; several kinds of reasoning (next subsection); planning; and problem solving.

8.2 Several kinds of probabilistic reasoning

Kinds of reasoning that may be exhibited by the SPS include: one-step ‘deductive’ reasoning; chains of reasoning; abductive reasoning; reasoning with probabilistic networks and trees; reasoning with ‘rules’; nonmonotonic reasoning and reasoning with default values; Bayesian reasoning with ‘explaining away’; causal reasoning; reasoning that is not supported by evidence; the inheritance of attributes in class hierarchies; and inheritance of contexts in part-whole hierarchies ([32, Section 10], [30, Chapter 7]).

There is also potential in the system for spatial reasoning [34, Section IV-F.1], and for what-if reasoning [34, Section IV-F.2].

The probabilistic nature of the SPS (Appendix A.6) makes it relatively straightforward to calculate absolute or conditional probabilities for results from the system, as for example in its several kinds of reasoning. In that connection, probabilities for inferences may be calculated as described in [32, Section 4.4] and [30, Section 3.7].

8.3 The representation and processing of several kinds of AI-related knowledge

Although SP-patterns are not very expressive in themselves, they come to life in the SPMA framework. Within that framework, they provide relevant knowledge for each aspect of intelligence mentioned in Section 8.1, for each kind of reasoning mentioned in Section 8.2, and more.

More specifically, they may serve in the representation and processing of such things as: the syntax of natural languages; class-inclusion hierarchies (with or without cross classification); part-whole hierarchies; discrimination networks and trees; if-then rules; entity-relationship structures [31, Sections 3 and 4]; relational tuples (*ibid.*, Section 3), and concepts in mathematics, logic, and computing, such as ‘function’, ‘variable’, ‘value’, ‘set’, and ‘type definition’ ([30, Chapter 10], [36, Section 6.6.1], [39, Section 2]).

As previously noted (Appendix A.1), the addition of two-dimensional SP patterns to the SPCM is likely to expand the capabilities of the SPS to the representation and processing of structures in two-dimensions and three-dimensions, and the representation of procedural knowledge with parallel processing.

8.4 The seamless integration of diverse aspects of intelligence, and diverse kinds of knowledge, in any combination

An important additional feature of the SPS, alongside its versatility in aspects of intelligence and diverse forms of reasoning, and its versatility in the representation and processing of diverse kinds of knowledge, is that *there is clear potential for the SPS to provide for the seamless integration of diverse aspects of intelligence and diverse forms of knowledge, in any combination*. This is because those several aspects of intelligence and several kinds of knowledge all flow from a single coherent and relatively simple source: the SPMA framework.

It appears that this kind of seamless integration is *essential* in any artificial system that aspires to human-level broad intelligence.

Figure 3 shows schematically how the SPS, with SPMA at centre stage, exhibits versatility and seamless integration.




Figure 3: A schematic representation of versatility and seamless integration in the SPS, with the SPMA concept centre stage.

9 Other potential benefits and applications of the SPS

Apart from the foregoing distinctive features and advantages of the SPS, it has other potential benefits and applications, several of them not closely related to AI. Relevant publications are outlined here:

- *Overview of potential benefits and applications.* Several potential areas of application of the SPS are described in [36]: the simplification and integration of computing systems; best-match and semantic forms of information

retrieval; software engineering; the representation of knowledge, reasoning, and the semantic web; information compression; bioinformatics; the detection of computer viruses; and data fusion.

- *The development of intelligence in autonomous robots.* The SPS opens up a radically new approach to the development of intelligence in autonomous robots [34].
- *The management of big data.* Somewhat unexpectedly, the SPS has potential to help solve nine significant problems associated with big data [35]: overcoming the problem of variety in big data; the unsupervised learning or discovery of ‘natural’ structures in data; interpretation of big data via pattern recognition, information retrieval, the parsing and production of natural language, and more; the analysis of streaming data; compression of big data; model-based coding for the efficient transmission of big data; potential gains in computational and energy efficiency in the analysis of big data; managing errors and uncertainties in data; and the visualisation of structure in big data.
- *Commonsense reasoning and commonsense knowledge.* Largely because of research by Ernest Davis and Gary Marcus (see, for example, [6]), the challenges in this area of AI research are now better known. Preliminary work shows that the SP System has promise in this area [40].
- *An intelligent database system.* The SPS has potential in the development of an intelligent database system with several advantages compared with traditional database systems [31].
- *Medical diagnosis.* The SPS may serve as a vehicle for medical knowledge and to assist practitioners in medical diagnosis, with potential for the automatic or semi-automatic learning of new knowledge [29].
- *Natural language processing.* The SP System has strengths and potential in the processing of natural language ([32, Section 8], [30, Chapter 5]).
- *Software engineering.* The SP System has potential in several aspects of software engineering [39].
- *Sustainability.* The SP System has clear potential for substantial reductions in the very large demands for energy of standard DNNs, and applications that need to manage huge quantities of data such as those produced by the Square Kilometre Array [43]. Where those demands are met by the burning of fossil fuels, there would be corresponding reductions in the emissions of CO₂.

- *Transparency in computing.* By contrast with applications with deep neural networks, the SPS provides a very full and detailed audit trail of all its processing, and all its knowledge may be viewed. Also, there are reasons to believe that, when the system is more fully developed, its knowledge will normally be structured in forms that are familiar such as class-inclusion hierarchies, part-whole hierarchies, run-length coding, and more. Strengths and potential of the SP System in these area are described in [48].
- *Vision, both artificial and natural.* The SPS opens up a new approach to the development of computer vision and its integration with other aspects of intelligence, and it throws light on several aspects of natural vision: [33, 44].

10 SP-Neural, a ‘neural’ version of the SPS

Another positive feature of the SPS is that abstract concepts in the SP Theory may be mapped into putative structures expressed in terms of neurons and their interconnections and intercommunications. This preliminary version of the SPS, called ‘SP-Neural’, is outlined in Appendix A.7, and described more fully in [37].

11 How to make generalisations without over-fitting or under-fitting; and how to minimise corrupted or ‘dirty’ data

The central role of IC in the workings of the SPCM provides a solution to two problems with unsupervised learning: how to generalise beyond the data, \mathbf{O} , without either over-generalisations (under-fitting) or under-generalisations (over-fitting); and how to learn correct forms despite the fact that \mathbf{O} normally contains errors of various kinds, otherwise called ‘dirty data’.

The solution proposed within the framework of the SPS ([32, Section 5.3], [30, Section 2.2.12]) appears to be on a firmer foundation than alternative proposals because it flows directly from the SPS and thus shares in the empirical and theoretical justifications for the SPS.

In brief, the solution is:

- Compress \mathbf{O} as much as possible via unsupervised learning in the SPS. The quality of the results depends on the thoroughness of the compression.
- The products of such learning are: an SP-grammar, \mathbf{G} , and an encoding of \mathbf{O} in terms of \mathbf{G} . The encoding may be referred to as \mathbf{E} , and the codes are \mathbf{C} .

- In general but with possible exceptions, **G** contains all the generalisations from the learning process.
- In general but with possible exceptions, **E** contains all the encodings of **O** in terms of **C**. It also contains all the parts of **O** which cannot be encoded, including all the ‘dirty data’—meaning false starts, mispronounced words, coughs, gurgles, and the like.
- Retain **G**, with all its generalisations, and discard **E**, with all the dirty data and encodings which are not normally of much interest.

This has been developed with more detail in [46].

12 The concept of ICMUP provides an entirely novel perspective on the foundations of mathematics

In view of evidence for the importance of IC in HLPC [41], and in view of the fact that mathematics is the product of human brains and has been designed to help human thinking, it should not be surprising to find that IC is central in the structure and workings of mathematics.

In keeping with that line of thinking, the concept of ICMUP provides an entirely novel perspective on the foundations of mathematics, described in the paper [42]. This perspective is a radical alternative to existing ‘isms’ in the foundations of mathematics, with potential connections with structuralism as described in [42, Section 4.4].

13 The SPS compared with other potential candidates for the SP-Challenge

One way another, most of the potential candidates for the SP-Challenge, including the SPS, are likely to be motivated by a need to overcome fragmentation in cognitive science or AI or both:

- In his famous essay “You can’t play 20 questions with nature and win” [16], Allen Newell exhorted researchers in psychology to develop theories where each one is for “a genuine slab of human behaviour” (p. 303)), instead of micro-theories, where each is for a relatively small aspect of human psychology.

This perspective was later expanded in Newell’s book on *Unified Theories of Cognition* [17] which included a description of “Soar”, a first attempt at a unified theory, which has been further developed in later research.

- The fragmentation of AI research has been well described by science writer Pamela McCorduck:

“The goals once articulated with debonair intellectual verve by AI pioneers appeared unreachable ... Subfields broke off—vision, robotics, natural language processing, machine learning, decision theory—to pursue singular goals in solitary splendor, without reference to other kinds of intelligent behaviour.” [15, p. 417]. Later, she writes of “the rough shattering of AI into subfields ... and these with their own sub-subfields—that would hardly have anything to say to each other for years to come.” [15, p. 424]. She adds: “Worse, for a variety of reasons, not all of them scientific, each subfield soon began settling for smaller, more modest, and measurable advances, while the grand vision held by AI’s founding fathers, a general machine intelligence, seemed to contract into a negligible, probably impossible dream.” (*ibid.*).

Apart from the SPS, a selection of possible candidates for the SP Challenge are described briefly in the following subsections. After those subsections, there is a description of the apparent advantages of the SPS compared with other potential candidates like those that have been described (Section 13).

13.1 Soar

As mentioned above, Soar was first described in Newell’s book on *Unified Theories of Cognition*. Now, the most comprehensive description of Soar is in *The Soar Cognitive Architecture* by John Laird [11].

Since this has not been updated since 2012, readers may wish to consult: The Soar home page; or ‘Soar (cognitive architecture)’, *Wikipedia*, retrieved 2021-11-23.

An impression of the organisation of Soar may be gained from Figure 4 which shows its memory structures.

13.2 ACT-R

ACT-R (which stands for “Adaptive Control of Thought—Rational”) is a cognitive architecture developed mainly by John Anderson and Christian Lebiere at Carnegie Mellon University.




Figure 4: Memory structures in Soar. Reproduced from [12, Figure 6], with permission from xxx.

A schematic representation of the ACT-R cognitive architecture is shown in Figure 5.




Figure 5: Schematic of the ACT-R cognitive architecture. Reproduced from [20, Abstract], with permission from xxx.

13.3 Artificial General Intelligence

Research towards the development of ‘Artificial General Intelligence’ (AGI) is a strand of research, referred to with that name mainly since the year 2000, which aims to recapture the interest of the AI pioneers in the breadth and versatility of human intelligence, and avoiding the kinds of fragmentation described at the beginning of Section 13.

Despite the welcome AGI-related goal of researchers in these areas, it appears fair to say that a foundation for the development of human-level broad AI, or AGI, has not yet been achieved, as described in these quotes:

“Despite all the current enthusiasm in AI, the technologies involved still represent no more than advanced versions of classic statistics and machine learning. Behind the scenes, however, many breakthroughs are happening on multiple fronts: in unsupervised language and grammar learning, deep-learning, generative adversarial methods, vision sys-

tems, reinforcement learning, transfer learning, probabilistic programming, blockchain integration, causal networks, and many more.” [10, Preface, Location 51].

And:

“If we want humanlike intelligence, we must figure out a way to construct it from the tools that we have available or we must build new tools. There are some attempts to create general intelligence with current tools, but none of them, so far, has demonstrated any success.” [21, p. 328].

The difficulties in achieving a good path to AGI are also demonstrated in the range of papers in the “Proceedings of the 13th International Conference, AGI 2020”. In general, they are about aspects of intelligence but with little or no synthesis across the field.

13.4 Systems that incorporate ‘Deep Neural Networks’

There is currently much interest in ‘Deep Neural Networks’ (DNNs), largely because of spectacular successes such as:

- The ‘AlphaGo’ system: beating the best human players at the game of Go.¹
- Providing substantial help with the difficult problem of predicting the 3D structure of a protein based on its genetic sequence data.²

And many other applications are being found for DNNs, with or without other programming.

These successes have, not unreasonably, led to suggestions by some researchers that DNNs might hold promise as a basis for the development of AGI. But others are skeptical. For example, Gary Marcus and Ernest Davis write:

“Deep learning still isn’t any kind of universal solvent. Nor does it have much to do with general intelligence, of the sort we need for open systems. In particular, it faces three core problems, each of which affects both deep learning itself and other popular techniques such as deep reinforcement learning that rely heavily on it. [14, p. 55].

¹See, for example: AlphaGo (tinyurl.com/yaxg9vz4); “Mastering the game of Go with deep neural networks and tree search” (tinyurl.com/12ge3oht). Both retrieved on the 18th of November, 2021.

²See, for example, “DeepMind AlphaFold Delivers ‘Unprecedented Progress’ on Protein Folding”.

The three core problems that they identify are:

- *Deep learning is greedy.* “In order to set all the connections in a neural net correctly, deep learning often requires a massive amount of data.” [14, p. 55].
- *Deep learning is opaque.* “Whereas classical expert systems are made up of rules that can be relatively easily understood (like ‘If a person has an elevated white blood cell count, then it is likely that the person has an infection’), neural networks are made up of vast arrays of numbers, virtually none of which make intuitive sense to ordinary human beings.” [14, p. 56].
- *Deep learning is brittle.* For example, a DNN may correctly recognise a picture of a car but may fail to recognise another slightly different picture of a car which, to a person, looks almost identical [27]. It has also been reported that a DNN may assign an image with near certainty to a class of objects such as ‘guitar’ or ‘penguin’, when people judge the given image to be something like white noise on a TV screen or an abstract pattern containing nothing that resembles a guitar or a penguin or any other object [18].

13.5 The advantages of the SPS as a foundation for the development of AGI

A direct comparison between the SPS and Soar, ACT-R, and AGI, is difficult because the focus of the SPS is different from the foci of the last three systems. Those latter three systems concentrate on relatively high level structures that may be related to structures in the brain, whereas the SPS has concentrated on developing an abstract model of intelligence that may be mapped on to those kinds of structures (see [37]).

There is a superficial resemblance between the SPS and DNNs because the SPS has SPMAs with rows and DNNs have layers. But the SPS is constantly building new SPMAAs from many different SP-patterns and with many different numbers of rows, whereas each DNN starts with fixed layers and works by strengthening links between them.

Compared with the kinds of systems described above, it seems that the SPS has several advantages as a foundation for the development of AGI:

- *Origin and development* (Section 1.1). The lengthy period of development of the SPS does not in itself mean that the SPS is good. But in that lengthy period of development, hundreds of seemingly plausible ideas have been tested and rejected. The ideas in the system as they are now are ideas that have survived that rigorous process of evaluation.

- *The benefits of a top-down, breadth-first research strategy with wide scope* (Section 2). Unlike most AI-related research, the SPS has been developed via a top-down, breadth-first strategy with a uniquely wide scope, seeking to simplify and integrate observations and concepts across AI, mainstream computing, mathematics, and human learning, perception, and cognition.

Although the SPS is far from being fully developed (as projected in [19]), the descriptive and explanatory range of the current version of the SPS is very wide, as is to be expected from its development via a top-down, breadth-first research strategy with wide scope.

- *Recognising the importance of information compression in human learning, perception, and cognition* (Section 3). There is now considerable evidence for the importance of information compression (IC) in human learning, perception, and cognition (HLPC) (Section 3). For that reason, any AI that aspires to human-level broad AI should have a central role for IC.

It appears that the SPS is the only AI system that qualifies. In research on deep neural networks (DNNs), there is some awareness that IC may be relevant [24], but in the development of DNNs up to now, it has never been given the central role that it has in the SPS.

- *Information compression via the matching and unification of patterns* (Section 4). A working hypothesis in the SP research is that all kinds of IC may be understood as ‘Information Compression via the Matching and Unification of Patterns’ (ICMUP).

Seven main variants of ICMUP are recognised.

This perspective on IC appears to be novel, with clear potential for productive developments.

- *The powerful concept of SP-multiple-alignment* (Section 6). Most of the strengths and potential of the SPS are due to the powerful concept of *SP-multiple-alignment* (SPMA), described in Appendix A.4 and Section 6. This is the last of seven variants of ICMUP described in Appendix A.3.

The SPMA concept has been shown to be a generalisation of the other six variants of ICMUP [47]. This helps to explain why the SPMA is so powerful.

As noted earlier, *the SPMA concept promises to be as significant for our understanding of intelligence as is DNA for many aspects of biology. The SPMA concept may prove to be the ‘double helix’ of intelligence.*

- *The potential of the SPS to solve 19 problems in AI research* (Section 7). The SPS has clear potential to solve 19 problems in AI research. Of those

problems, 17 have been identified by influential experts in AI, as reported by science writer Martin Ford in his book *Architects of Intelligence* [7].

- *Strengths and potential of the SPS in modelling aspects of intelligence* (Section 8). The SPS has strengths and potential in modelling several aspects of intelligence (Section 8): versatility in intelligent behaviour including several kinds of probabilistic reasoning; versatility in the representation and processing of AI-related knowledge; and the seamless integration of diverse aspects of intelligence, and diverse kinds of knowledge, in any combination.
- *Other potential benefits and applications of the SPS* (Section 9). There are several other potential benefits and applications of the SPS, including several described in separate papers. They include: the development of intelligence in autonomous robots; the management of big data; preliminary work suggests that the SPS has potential with commonsense reasoning and commonsense knowledge; the SPS has potential as an intelligent database system; assisting with medical diagnosis; the SPS has strengths and potential in the processing of natural language; the SPS has potential in several aspects of software engineering; sustainability: the SPS has clear potential for substantial reductions in the very large demands for energy of standard DNNs, and of applications that need to manage huge quantities of data such as those that will be produced by the Square Kilometre Array; transparency in computing: the SPS provides an audit trail of all its processing, and all its knowledge may be viewed, probably in forms that are familiar; the SPS opens up a new approach to the development of computer vision, and it throws light on several aspects of natural vision.
- *SP-Neural* (Section 10). Abstract concepts in the SPS may be mapped into putative structures expressed in terms of neurons and their interconnections and intercommunications [37].
- *Mathematics* (Section 12). The concept of ICMUP provides an entirely novel perspective on the foundations of mathematics, described in the paper [42]. This perspective is a radical alternative to existing ‘isms’ in the foundations of mathematics.

In itself, this does not mean that the SPS is a good foundation for the development of AGI, but it suggests a broad significance for the ICMUP concepts, including the SPMA concept, and thus the potential of those concepts as a foundation for the development of AGI.

- *Generalisation and ‘dirty data’* (Section 11). An integral part of unsupervised learning in the SPS is the creation of generalisations from raw

data, apparently without either over-generalisations (under-fitting) or under-generalisations (over-fitting). A related feature is that the system can learn ‘correct’ knowledge and generalisations that reduce or eliminate the distorting influence of errors, aka ‘dirty data’, in the raw data.

14 Future directions for the SP research

It is envisaged that the SPCM will be developed into a first version of an *SP Machine* via the addition of high levels of parallel processing and an improved user interface. The SP Machine will be hosted on a high-end workstation with one or more GPUs.

Figure 6, shows a schematic representation of how the SP Machine may be developed and applied.




Figure 6: Schematic representation of the development and application of the SP Machine. Reproduced from Figure 2 in [32], with permission.

To facilitate future development of the SPS, the software for the SP Machine will be open-source. This will allow researchers anywhere to create clones of the SP Machine. Then individual researchers or teams may investigate any aspect of the system including both strengths and weaknesses of the SPS concept, to develop the SP Machine along the lines described in “A roadmap for the development of the

‘SP Machine’ for artificial intelligence” [19], and to explore potential applications of the SP Machine.

Even if the SPS turns out to be the winner of the SP-Challenge, it would be a mistake for all available research eggs to be put into that one basket. Given the uncertainties attaching to any vision of how AGI might be achieved, it would make more sense to hedge our bets by developing at least one or two of the best alternatives to the SPS, in addition to the SPS itself.

15 Conclusion

The ‘SP-Challenge’ is an invitation to other researchers to refute the claim in this paper that the SP System (SPS), outlined in Appendix A, is more promising as a foundation for the development of human-level broad AI than any alternative.

In that connection, the main strengths and potential of the SPS are as described in Section 13.5.

Acknowledgements

I’m grateful to anonymous referees for constructive comments on earlier versions of this paper. I’m also grateful to many other people, too numerous to list, for their comments and suggestions, which have helped in the development of the SP System.

Appendices

A Outline of the SP Theory of Intelligence and its realisation in the SP Computer Model

The *SP System* (SPS)—meaning the *SP Theory of Intelligence* and its realisation in the *SP Computer Model* (SPCM)—is the product of a lengthy programme of research, from about 1987 to 2021 with a break between early 2006 and late 2012. This programme of research has included the creation and testing of many versions of the SPCM.

A major discovery has been the concept of *SP-multiple-alignment* (SPMA) and, within the SPCM, its versatility in several aspects of intelligence (Section 8 and Appendix A.4).

A.1 High level view of the SPS

In broad terms, the SPS is a brain-like system that takes in *New* information through its senses and stores some or all of it as *Old* information that is compressed, as shown schematically in Figure 7.




Figure 7: Schematic representation of the SPS from an ‘input’ perspective. Reproduced, with permission, from Figure 1 in [32].

The importance of information compression (IC) is described in Section 3.

In the SPS, all kinds of knowledge are represented with *SP-patterns*, where each such SP-pattern is an array of atomic *SP-symbols* in one or two dimensions. An SP-symbol is simply a ‘mark’ that can be matched with any other SP-symbol to determine whether it is the same or different.

Although SP-patterns are very simple, they come alive within the framework of SP-multiple-alignments (Appendix A.4, Section 6).

At present, the SPCM works only with one-dimensional SP-patterns but it is envisaged that it will be generalised to work with two-dimensional SP-patterns as well. This should facilitate:

- The representation and processing of pictures and diagrams;
- For reasons explained in [33, Sections 6.1 and 6.2], the representation and processing of three-dimensional structures as well.

- The representation of two or more procedures in parallel [34, Sections V-G, V-H, V-I, and Appendix C].

A.2 The name ‘SP’

Since people often ask, the name ‘SP’ originated like this:

- The SPS aims to simplify and integrate observations and concepts across a broad canvass (Section 2), which means applying IC to those observations and concepts;
- IC is a central feature of the structure and workings of the SPS;
- And IC may be seen as a process that increases the *Simplicity* of a body of information, **I**, whilst retaining as much as possible of the descriptive and explanatory *Power* of **I**.

Nevertheless, *it is intended that ‘SP’ should be treated as a name, without any need to expand the letters in the name, as with such names as ‘IBM’ or ‘BBC’.*

A.3 Information compression in the SPS via the matching and unification of patterns

In the development of the SPS, it has proved useful to understand IC as a process of searching for patterns that match each other and the merging or ‘unifying’ of patterns that are the same.

As noted in Section 4, the expression ‘Information Compression via the Matching and Unification of Patterns’ may be abbreviated as ‘ICMUP’.

A.4 SP-multiple-alignment

A central idea in the SPS, is the simple but powerful concept of *SP-multiple-alignment* (SPMA), borrowed and adapted from the concept of ‘multiple sequence alignment’ in bioinformatics.

Some of its versatility within the SPS is outlined in Section 6, and more in Sections 7, 8, and 9.

The origins, development, and application of this concept in the SP programme of research are described in Section 6.

A.5 Unsupervised learning in the SPS

In the SPS, learning is ‘unsupervised’, deriving structures from incoming sensory information without the need for any kind of ‘teacher’, or anything equivalent (*cf.* [8]).

Unsupervised learning in the SPS is quite different from ‘Hebbian’ learning via the gradual strengthening or weakening of neural connections (Section challenge-of-unsupervised-learning-section), variants of which are the mainstay of learning in DNNs. In the SPS, unsupervised learning incorporates the building of SPMAs but there are other processes as well.

A.5.1 The creation of Old SP-patterns

In brief, the system creates Old SP-patterns from complete New SP-patterns and also from partial matches between New and Old SP-patterns. All learning in the SPS starts with the taking in of information from the environment:

- If that information is the same as one or more Old SP-patterns, then the frequency of the one SP-pattern is increased, or frequencies of the two or more SP-patterns are increased.
- If that information is entirely new, ‘ID’ SP-symbols³ are added at the beginning and end of the pattern so that it becomes an SP-pattern. Then it is added directly to the store of Old SP-patterns.
- If partial matches can be made between the newly-received information and one or more of the stored Old SP-patterns, then each of the parts that match, and each of the parts that don’t match, are made into SP-patterns by the addition of ID SP-symbols at the beginning and end, and the newly-created SP-patterns are added to the store of Old SP-patterns.

A.5.2 The creation of SP-grammars

With a given body of New SP-patterns, the system processes them as just sketched, and then searches for one or two ‘good’ *SP-grammars*, where an SP-grammar is a collection of Old SP-patterns, and it is ‘good’ if it is effective in the economical encoding of the original set of New SP-patterns, where that economical encoding is achieved via SPMA.

As with the building of SPMAAs, the process of creating good grammars is normally too complex to be done by exhaustive search so heuristic methods are

³Appendix A.1

needed. This means that the system builds SP-grammars incrementally and, at each stage, it discards all but the best SP-grammars.

As with the building of SPMA's, the use of heuristic methods helps to ensure that computational complexities in the SPS are within reasonable bounds [30, Sections A.4, 3.10.6 and 9.3.1].

The SPCM has already demonstrated an ability to learn generative grammars from unsegmented samples of English-like artificial languages, including segmental structures, classes of structure, and abstract patterns, and to do this in an 'unsupervised' manner ([32, Section 5], [30, Chapter 9]).

But there are (at least) two shortcomings with unsupervised learning in the SPS, outlined in [32, Section 3.3].

A.6 The probabilistic nature of the SPS

Owing to the intimate relation that is known to exist between IC and concepts of probability (see 'Algorithmic Probability Theory' developed by Solomonoff (see [25, 26], [13]), and owing to the fundamental role of IC in the workings of the SPS, the system is inherently probabilistic ([32, Section 4.4], [30, Section 3.7]).

That said, it appears to be possible to imitate the all-nothing-nature of conventional computing systems via the use of data where all the probabilities yielded by the system are at or close to 0 or 1.

Because of the probabilistic nature of the SPS, it lends itself to the modelling of HLPC because of the probabilistic nature of much of human thinking. Also, the SPS sits comfortably within AI because of the probabilistic nature of most systems in AI, at least in more recent work in that area.

An advantage of the SPS in those areas is that it is relatively straightforward to calculate absolute or conditional probabilities for results obtained in, for example, different kinds of reasoning (Section 8.2).

The very close connection that exists between IC and concepts of probability may suggest that there is nothing to choose between them. But [42, Section 8.2] argues that, in research on aspects of AI and HLPC, there are reasons to regard IC as more fundamental than probability and a better starting point for theorising.

A.7 SP-Neural

A potentially useful feature of the SPS is that it is possible to see how abstract constructs and processes in the system may be realised in terms of neurons and their interconnections. This is the basis for *SP-Neural*, a 'neural' version of the SPS, described in [37].

The concept of an SP-symbol may be realised as a *neural symbol* comprising a single neuron or, more likely, a small cluster of neurons. An SP-pattern maps quite

well on to the concept of a *pattern assembly* comprising a group of inter-connected SP-symbols. And an SPMA may be realised in terms of pattern assemblies and their interconnections, as illustrated in Figure 8.

In this connection, it is relevant to mention that the SPS, in both its abstract, SPCM, and neural forms, is quite different from DNNs [24] and has substantial advantages compared with such systems, as described in [45] and in [38, Section V].




Figure 8: A schematic representation of a partial SPMA in SP-Neural, as discussed in [37, Section 4]. Each broken-line rectangle with rounded corners represents a *pattern assembly*—corresponding to an SP-pattern in the main SP Theory of Intelligence; each character or group of characters enclosed in a solid-line ellipse represents a *neural symbol* corresponding to an SP-symbol in the main SP Theory of Intelligence; the lines between pattern assemblies represent nerve fibres with arrows showing the direction in which impulses travel; neural symbols are mainly symbols from linguistics such as ‘NP’ meaning ‘noun phrase’, ‘D’ meaning a ‘determiner’, ‘#D’ meaning the end of a determiner, ‘#NP’ meaning the end of a noun phrase, and so on.

B Abbreviations

Abbreviations used in this paper are detailed below.

- *Artificial Intelligence*: ‘AI’.
- *Artificial General Intelligence*: ‘AGI’.
- *Deep Neural Network*: ‘DNN’.
- *Human Learning, Perception, and Cognition*: ‘HLPC’.
- *Information Compression*: ‘IC’.
- *Information Compression via the Matching and Unification of Patterns*: ‘ICMUP’.
- *SP Computer Model*: ‘SPCM’.
- *SP-multiple-alignment*: ‘SPMA’.
- *SP System*: ‘SPS’.

As noted in Appendix A.2, it is intended that ‘SP’ should be treated as a name, like ‘IBM’ or ‘BBC’, not an abbreviation.

References

- [1] F. Attneave. Some informational aspects of visual perception. *Psychological Review*, 61:183–193, 1954.
- [2] F. Attneave. *Applications of Information Theory to Psychology*. Holt, Rinehart and Winston, New York, 1959.
- [3] H. B. Barlow. Sensory mechanisms, the reduction of redundancy, and intelligence. In HMSO, editor, *The Mechanisation of Thought Processes*, pages 535–559. Her Majesty’s Stationery Office, London, 1959.
- [4] H. B. Barlow. Trigger features, adaptation and economy of impulses. In K. N. Leibovic, editor, *Information Processes in the Nervous System*, pages 209–230. Springer, New York, 1969.
- [5] A-S. Barwich. The value of failure in science: The story of grandmother cells in neuroscience. *Frontiers in Neuroscience*, 13:1121, 2019.

- [6] E. Davis and G. Marcus. Commonsense reasoning and commonsense knowledge in artificial intelligence. *Communications of the ACM*, 58(9):92–103, 2015.
- [7] M. Ford. *Architects of Intelligence: the Truth About AI From the People Building It*. Packt Publishing, Birmingham, UK, Kindle edition, 2018.
- [8] M. Gold. Language identification in the limit. *Information and Control*, 10:447–474, 1967.
- [9] D. H. Hubel. Evolution of ideas on the primary visual cortex, 1955-1975: a biased historical account. *Bioscience Reports*, 2:435–469, 1982.
- [10] M. Iklé, A. Franz, R. Rzepka, and B. Goertzel, editors. *Artificial General Intelligence*, volume 10999 of *Lecture Notes in Computer Science*, Heidelberg, 2018. Springer.
- [11] J. E. Laird. *The Soar Cognitive Architecture*. The MIT Press, Cambridge, Mass., 2012. ISBN-13: 978-0-262-12296-2.
- [12] J. F. Lehman, J. Laird, and P. Rosenbloom. A gentle introduction to soar, an architecture for human cognition: 2006 update. Technical report, University of Michigan, 2006. Work supported by the National Science Foundation under Grant No. 0413013.
- [13] M. Li and P. Vitányi. *An Introduction to Kolmogorov Complexity and Its Applications*. Springer, New York, 4th edition, 2019.
- [14] G. F. Marcus and E. Davis. *Rebooting AI: Building Artificial Intelligence We Can Trust*. Pantheon Books, New York, Kindle edition, 2019.
- [15] P. McCorduck. *Machines Who Think: a Personal Inquiry Into the History and Prospects of Artificial Intelligence*. A. K. Peters Ltd, Natick, MA, second edition, 2004.
- [16] A. Newell. You can't play 20 questions with nature and win: Projective comments on the papers in this symposium. In W. G. Chase, editor, *Visual Information Processing*, pages 283–308. Academic Press, New York, 1973.
- [17] A. Newell, editor. *Unified Theories of Cognition*. Harvard University Press, Cambridge, Mass., 1990.
- [18] A. Nguyen, J. Yosinski, and J. Clune. Deep neural networks are easily fooled: high confidence predictions for unrecognizable images. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR 2015)*, pages 427–436, 2015.

- [19] V. Palade and J. G. Wolff. A roadmap for the development of the ‘SP Machine’ for artificial intelligence. *The Computer Journal*, 62:1584–1604, 2019. <https://doi.org/10.1093/comjnl/bxy126>, arXiv:1707.00614, bit.ly/2tWb88M.
- [20] F. E. Ritter, F. Tehranchi, and J. D. Oury. Act-r: a cognitive architecture for modelling cognition. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2018.
- [21] H. L. Roitblat. *Algorithms Are Not Enough*. The MIT Press, Cambridge, Mass., Kindle edition, 2020.
- [22] A. Roy. An extension of the localist representation theory: grandmother cells are also widely used in the brain. *Frontiers in Psychology*, 4:1–3, 2013.
- [23] A. Roy. Commentary: the value of failure in science: the story of grandmother cells in neuroscience. *Frontiers in Neuroscience*, 14(2):1–2, 2020.
- [24] J. Schmidhuber. Deep learning in neural networks: an overview. *Neural Networks*, 61:85–117, 2015.
- [25] R. J. Solomonoff. A formal theory of inductive inference. Parts I and II. *Information and Control*, 7:1–22 and 224–254, 1964.
- [26] R. J. Solomonoff. The discovery of algorithmic probability. *Journal of Computer and System Sciences*, 55(1):73–88, 1997.
- [27] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. Intriguing properties of neural networks. Technical report, Google Inc. and others, 2014. arXiv:1312.6199v4 [cs.CV] 19 Feb 2014, bit.ly/1elzRGM (PDF).
- [28] J. G. Wolff. Learning syntax and meanings through optimization and distributional analysis. In Y. Levy, I. M. Schlesinger, and M. D. S. Braine, editors, *Categories and Processes in Language Acquisition*, pages 179–215. Lawrence Erlbaum, Hillsdale, NJ, 1988. bit.ly/ZIGjyc.
- [29] J. G. Wolff. Medical diagnosis as pattern recognition in a framework of information compression by multiple alignment, unification and search. *Decision Support Systems*, 42:608–625, 2006. arXiv:1409.8053 [cs.AI], bit.ly/1F366o7.
- [30] J. G. Wolff. *Unifying Computing and Cognition: the SP Theory and Its Applications*. CognitionResearch.org, Menai Bridge, 2006. ISBNs: 0-9550726-0-3 (ebook edition), 0-9550726-1-1 (print edition). Distributors, including Amazon.com, are detailed on bit.ly/WmB1rs. The print version is produced via

print-on-demand from “INGRAM Lightning Source” and, via that technology, is unlikely to go out of print.

- [31] J. G. Wolff. Towards an intelligent database system founded on the SP Theory of Computing and Cognition. *Data & Knowledge Engineering*, 60:596–624, 2007. arXiv:cs/0311031 [cs.DB], bit.ly/1CUldR6.
- [32] J. G. Wolff. The SP Theory of Intelligence: an overview. *Information*, 4(3):283–341, 2013. arXiv:1306.3888 [cs.AI], bit.ly/1NOMJ6l.
- [33] J. G. Wolff. Application of the SP Theory of Intelligence to the understanding of natural vision and the development of computer vision. *SpringerPlus*, 3(1):552–570, 2014. arXiv:1303.2071 [cs.CV], bit.ly/2oIpZB6.
- [34] J. G. Wolff. Autonomous robots and the SP Theory of Intelligence. *IEEE Access*, 2:1629–1651, 2014. arXiv:1409.8027 [cs.AI], bit.ly/18DxU5K.
- [35] J. G. Wolff. Big data and the SP Theory of Intelligence. *IEEE Access*, 2:301–315, 2014. arXiv:1306.3890 [cs.DB], bit.ly/2qfSR3G. This paper, with minor revisions, is reproduced in Fei Hu (Ed.), *Big Data: Storage, Sharing, and Security*, Taylor & Francis LLC, CRC Press, Boca Raton, Florida, 2016, Chapter 6, pp. 143–170.
- [36] J. G. Wolff. The SP Theory of Intelligence: benefits and applications. *Information*, 5(1):1–27, 2014. arXiv:1307.0845 [cs.AI], bit.ly/1FRYwew.
- [37] J. G. Wolff. Information compression, multiple alignment, and the representation and processing of knowledge in the brain. *Frontiers in Psychology*, 7(1584), 2016. arXiv:1604.05535 [cs.AI], bit.ly/2esmYyt.
- [38] J. G. Wolff. The SP Theory of Intelligence: its distinctive features and advantages. *IEEE Access*, 4:216–246, 2016. arXiv:1508.04087 [cs.AI], bit.ly/2qgq5QF.
- [39] J. G. Wolff. Software engineering and the SP Theory of Intelligence. Technical report, CognitionResearch.org, 2017. arXiv:1708.06665 [cs.SE], bit.ly/2w99Wzq.
- [40] J. G. Wolff. Commonsense reasoning, commonsense knowledge, and the SP Theory of Intelligence. Technical report, CognitionResearch.org, 2019. In preparation. viXra:1901.0051v2, hal-01970147 version 3, bit.ly/2RESeut.
- [41] J. G. Wolff. Information compression as a unifying principle in human learning, perception, and cognition. *Complexity*, 2019:38, February 2019. Article ID 1879746. viXra:1707.0161v3, hal-01624595 v2.

- [42] J. G. Wolff. Mathematics as information compression via the matching and unification of patterns. *Complexity*, 2019:25, 2019. Article ID 6427493, Archives: vixra.org/abs/1912.0100 and hal.archives-ouvertes.fr/hal-02395680. This paper is reproduced in *New Ideas Concerning Science and Technology*, Vol. 13, 12 June 2021, Pages 132–169, doi.org/10.9734/bpi/nicst/v13/8449D.
- [43] J. G. Wolff. How the SP System may promote sustainability in energy consumption in IT systems. *Sustainability*, 13(8):4565, 2021.
- [44] J. G. Wolff. The potential of the SP System in machine learning and data analysis for image processing. *Big Data and Cognitive Computing*, 5(1):7, 2021. tinyurl.com/7y9fbe7f.
- [45] J. G. Wolff. Problems in AI research and how the SP System may solve them. Technical report, CognitionResearch.org, 2021. Download: [arXiv:2009.09079v3](https://arxiv.org/abs/2009.09079v3) [cs.CY], [viXra:2009.0012](https://vixra.org/abs/2009.0012), [hal-02928784](https://hal.archives-ouvertes.fr/hal-02928784); tinyurl.com/5buj8swp. Submitted for publication.
- [46] J. G. Wolff. A proposed solution to problems in learning the knowledge needed by self-driving vehicles. 2021. Submitted for publication, tinyurl.com/n6vpxcuf.
- [47] J. G. Wolff. The *sp-multiple-alignment* concept as a generalisation of six other variants of ‘information compression via the matching and unification of patterns’. Technical report, CognitionResearch.org, 2021. tinyurl.com/dukwnkdf.
- [48] J. G. Wolff. Transparency and granularity in the SP Theory of Intelligence and its realisation in the SP Computer Model. In Witold Pedrycz and Shyi-Ming Chen, editors, *Interpretable Artificial Intelligence: A Perspective of Granular Computing*. Springer-Verlag, Heidelberg, 2021. Republished in the book *Interpretable Artificial Intelligence: A Perspective of Granular Computing*, Witold Pedrycz and Shyi-Ming Chen (editors), Springer: Heidelberg, 2021, ISBN 978-3-030-64948-7, DOI: 10.1007/978-3-030-64949-4. Archive of chapter: [arXiv:2009.06370](https://arxiv.org/abs/2009.06370) [cs.AI].