



HAL
open science

Commonsense reasoning, commonsense knowledge, and the SP Theory of Intelligence

J Gerard Gerard Wolff

► **To cite this version:**

J Gerard Gerard Wolff. Commonsense reasoning, commonsense knowledge, and the SP Theory of Intelligence. 2019. hal-01970147v3

HAL Id: hal-01970147

<https://hal.science/hal-01970147v3>

Preprint submitted on 5 Jun 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Commonsense reasoning, commonsense knowledge, and the SP Theory of Intelligence

J Gerard Wolff*

June 4, 2019

Abstract

Commonsense reasoning (CSR) and commonsense knowledge (CSK) (together abbreviated as CSRK) are areas of study concerned with problems which are trivially easy for adults but which are challenging for artificial systems. This paper describes how the *SP System*—meaning the *SP Theory of Intelligence* and its realisation in the *SP Computer Model*—has strengths and potential in several aspects of CSRK. A particular strength of the SP System is that it shows promise as an overarching theory for four areas of relative success with CSRK problems—described by other authors—which have been developed without any integrative theory. How the SP System may help to solve four other kinds of CSRK problem is described: 1) how the strength of evidence for a murder may be influenced by the level of lighting of the murder as it was witnessed; 2) how people may arrive at the commonly-accepted interpretation of phrases like “water bird”; 3) interpretation of the horse’s head scene in “The Godfather” film; and 4) how the SP System may help to resolve the reference of an ambiguous pronoun in sentences in the format of a ‘Winograd schema’. Also described is why a fifth CSRK problem—modelling how a cook may crack an egg into a bowl—is beyond the capabilities of the SP System as it is now and how those deficiencies may be overcome via planned developments of the system.

Keywords: Commonsense reasoning; commonsense knowledge; SP Theory of Intelligence; SP Computer Model; SP System; information compression.

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

1 Introduction

‘Commonsense reasoning’ with ‘commonsense knowledge’ are things that we use everyday to solve commonplace problems such as how to make a cup tea, how to go shopping, and so on. Although they seem trivial for most adults, they present some challenging problems for AI, as described for example by Ernest Davis and Gary Marcus [13]. They say that, for an artificial system to achieve human-level performance with these kinds of problems “basic knowledge of the commonsense world—time, space, physical interactions, people, and so on—will be necessary. Although a few forms of commonsense reasoning, such as taxonomic reasoning and temporal reasoning are well understood, progress has been slow.” [13, p. 92].

The main aim of this paper is to describe how the *SP System*—meaning the *SP Theory of Intelligence* and its realisation in the *SP Computer Model*—may help solve problems in the areas of commonsense reasoning and commonsense knowledge, and to describe some shortcomings of the system in that area and how they may be overcome.

1.1 Presentation

There is an outline description of the SP System in Appendix A, with enough detail to ensure that the rest of the paper makes sense. Readers who are not already familiar with the SP System are urged to read that appendix first.

Section 2 provides a selective review of CSRK-related research, including a summary of what Davis and Marcus [13] say about successes in that area. Then Section 3 describes some apparent strengths and weaknesses of the SP System as it is now as a means of solving CSRK-related problems, and how the weaknesses may be overcome. The main body of the paper—Sections 4 to 8—presents some CSRK examples to illustrate what can and cannot be done with the SP System as it is now, and how its shortcomings may be overcome.

For the sake of brevity in this paper:

- Davis and Marcus, as they write in [13], will be referred to as ‘DM’;
- Problems of commonsense reasoning will normally be referred to as ‘CSR’; how commonsense knowledge may be represented will normally be ‘CSK’; and the two together may be referred to as ‘CSRK’. CSK (and CSRK) will be understood to include the problem of how commonsense knowledge may be learned.
- The name ‘SP’ in the SP System, the SP Theory of Intelligence, and the SP Computer Model is short for *Simplicity* and *Power*, for a reason explained in Appendix A.1, and another reason explained in Appendix A.1.4.

- ‘Information compression’ may be shortened to ‘IC’, and ‘information compression via the matching and unification of patterns’ may be referred to more briefly as ‘ICMUP’.
- The expression ‘human learning, perception, and cognition’ may be shortened to ‘HLPC’.
- And a system which is a version of logic will be referred to as ‘VLS’.

2 Research related to CSRK

This section first describes a selection of relatively recent research related to CSRK. Then, towards the end of the section, four subsections outline where there has been relative success with CSRK, as described by DM, with discussions of the strengths or weaknesses of the SP System in those areas.

An important strand of work is the application of VLSs to CSRK problems. Books in this area include *Representations of Commonsense Knowledge* by Ernest Davis [11] and *Commonsense Reasoning: An Event Calculus Based Approach* by Erik Mueller [37]. A recent survey by Davis of research on this topic is in [12].

‘ATOMIC’ is an atlas of everyday commonsense reasoning developed by a team at the Allen Institute for Artificial Intelligence [45]. Unlike existing systems that concentrate on taxonomic knowledge, ATOMIC focuses on inferential knowledge organized as *if-then* relations. With relevant training, neural models can acquire simple commonsense capabilities and reason about previously unseen events. Favourable experimental results are reported.

The Cyc project, initiated and led for many years by Douglas Lenat (See, for example, [59, 43]) has assembled a very large database of knowledge about basic concepts and ‘rules of thumb’ about how the world works. The intention has been to facilitate the creation of AI applications that may perform human-like reasoning and which can cope with novel situations that were not preconceived.

DM write quite extensively about this project [13, pp. 99–103], with remarks such as “No systematic evaluation of the contents, capacities, and limitations of CYC has been published.” [13, p. 101], and “The [CSRK] field might well benefit if CYC were systematically described and evaluated. If CYC has solved some significant fraction of commonsense reasoning, then it is critical to know that, both as a useful tool, and as a starting point for further research. If CYC has run into difficulties, it would be useful to learn from the mistakes that were made. If CYC is entirely useless, then researchers can at least stop worrying about whether they are reinventing the wheel.” [13, p. 103].

‘ConceptNet’ is a knowledge graph that connects words and phrases of natural language with labeled edges and is designed to represent the general knowledge

involved in understanding language (see, for example, [50]).

‘SenticNet’ is a three-level knowledge representation for sentiment analysis. The project uses recurrent neural networks to infer primitives by lexical substitution and for grounding common and commonsense knowledge by means of multi-dimensional scaling (see, for example, [10]).

Yukun Ma and colleagues (see, for example, [30]) propose an extension of “long short-term memory” (LSTM), termed “Sentic LSTM”. This augments a LSTM network with an hierarchical attention mechanism comprising “target-level attention” and “sentence-level attention”. In the system, commonsense knowledge of sentiment-related concepts is incorporated into the end-to-end training of a deep neural network for sentiment classification. Experiments show that the system can outperform state-of-the-art methods in targeted aspect sentiment tasks.

Joseph Blass and Kenneth Forbus [8] describe an approach called *analogical chaining* to create cognitive systems that can perform commonsense reasoning. In this approach, ‘commonsense units’ are provided to the system via natural language instruction.

Leora Morgenstern and colleagues [36] discuss plans to run a competition modelled on the ‘Winograd Schema Challenge’, a type of challenge in the interpretation of natural language, described by Terry Winograd [58], which is easy for people but, normally, hard for computers.

Shiqi Zhang and Peter Stone [75] discuss aspects of reasoning in intelligent robots, including challenges in modelling the kinds of commonsense reasoning that people find easy. Following a discussion of some alternative frameworks, they introduce the ‘CORPP’ algorithm, with apparent advantages over those alternatives.

Somak Aditya and colleagues [1] explore the use of visual commonsense knowledge and other kinds of knowledge for scene understanding. They combine visual processing with techniques from natural language understanding.

Nicole Maslan and colleagues [32] present a set of challenge problems for the logical formalization of commonsense knowledge which, unlike other such sets, is designed to support the development of logic-based commonsense theories.

André Freitas and colleagues [18] describe a selective graph-navigation mechanism based on a distributional-relational semantic model which can be applied to querying and reasoning with a variety of knowledge bases, and they discuss how it may be applied.

A paper by Gary Marcus and Ernest Davis [31], a little earlier than the already-referenced [13], discusses an issue related to commonsense reasoning: whether or not the mind should be viewed as a near-optimal or rational engine of probabilistic inference. They argue that this view is markedly less promising than is widely believed. They also argue that the commonly-supposed equivalence between probabilistic inference, on the one hand, and rational or optimal inference, on the other,

is not justified.

And another paper by the same two authors [14] examines whether and how the kinds of ‘simulation’ that serve in such things as flight simulators for the training of pilots may be seen as systems for automated reasoning, including commonsense reasoning. They argue that, for the modelling of automated reasoning, “in many cases, simulation can play at most a limited role.” [31, Abstract].

In addition to the research just outlined, much other research related to CSRK, published a little earlier, is described in [56].

2.1 The four areas of relative success with CSRK, described in by DM

The four subsections that follow describe four areas where, as suggested DM, researchers have achieved relative success with CSRK [13, pp. 94–97]. There is discussion of these four areas of relative success, in relation to the SP System, in each of the subsections below, and in Section 2.1.5.

2.1.1 Taxonomic reasoning

The first area of success in CSRK-related processing discussed in [13, pp. 94–97] is taxonomic reasoning, with its close association with inheritance of attributes. They say:

“Simple taxonomic structures such as those illustrated here are often used in AI programs” [13, p. 95], “Many specialized taxonomies have been developed in domains such as medicine and genomics” [13, p. 96], and “A number of sophisticated extensions of the basic inheritance architecture described here have also been developed. Perhaps the most powerful and widely used of these is description logic. Description logics provide tractable constructs for describing concepts and the relations between concepts, grounded in a well-defined logical formalism. They have been applied extensively in practice, most notably in the Semantic Web ontology language OWL.¹” [13, p. 96].

Of course, class hierarchies with inheritance of attributes were introduced as useful constructs for simulation in the Simula computer language [7] and have subsequently been widely adopted in many ordinary programming languages.

The representation of taxonomic knowledge and inheritance of attributes are clear strengths of the SP System, as outlined in Appendices B.1, B.2, and B.4. Those strengths with taxonomic knowledge and taxonomic reasoning are not coded

¹The OWL language is described on the web page www.w3.org/OWL/, retrieved 2018-11-06.

into the system by hand but arise as a by-product of IC, ICMUP, and SP-multiple-alignment which are bedrock in the SP System, giving rise to its versatility across AI and related areas.

2.1.2 Temporal reasoning

DM write:

“Representing knowledge and automating reasoning about times, durations, and time intervals is a largely solved problem [16]. For instance, if one knows that Mozart was born earlier and died younger than Beethoven, one can infer that Mozart died earlier than Beethoven.” But “Integrating such reasoning with specific applications, such as natural language interpretation, has been much more problematic. Natural language expressions for time are complex and their interpretation is context dependent. ... However, many important temporal relations are not explicitly stated in texts, they are inferred; and the process of inference can be difficult. Basic tasks like assigning timestamps to events in news stories cannot be currently done with any high degree of accuracy [52].” [13, p. 96].

As noted in Section 3.5.4, the representation and processing of magnitudes like time are shortcomings in the SP Computer Model where more work is needed.

2.1.3 Action, events, and change

DM write:

“Another area of commonsense reasoning that is well understood is the theory of action, events, and change. In particular, there are very well established representational and reasoning techniques for domains that satisfy [constraints such as ‘events are atomic’, ‘every change in the world is the result of an event’, and so on] [44]” [13, p. 97] and

“For domains that satisfy these constraints, the problem of representation and important forms of reasoning, such as prediction and planning, are largely understood. Moreover, a great deal is known about extensions to these domains [such as ‘continuous domains’, ‘simultaneous events’, and more] ... The primary successful applications of these kinds of theories has been to high-level planning [44], and to some extent to robotic planning, for example, [15].” [13, p. 97].

There has been no attempt yet to examine this area from the perspective of the SP programme of research. It is anticipated that any such perspective would express concepts like ‘action’, ‘event’, and ‘change’ in terms of information and the compression of information within the framework of SP-multiple-alignment.

In view of the versatility of the SP System in diverse aspects of intelligence (Section 3.1.1), including diverse forms of reasoning (Section 3.1.2), and the representation of diverse forms of knowledge (Section 3.1.3), there seems to be a reasonable chance that, at some stage, concepts like action, events, and change, may be modelled in the SP System.

2.1.4 Qualitative reasoning

DM write:

“One type of commonsense reasoning that has been analyzed with particular success is known as qualitative reasoning. ... If the price of an object goes up then (usually, other things being equal) the number sold will go down. If the temperature of gas in a closed container goes up, then the pressure will go up. ... For problems within the scope of the representation, the reasoning mechanism works well. However, there are many problems in physical reasoning, particularly those involving substantial geometric reasoning, that cannot be represented in this way, and therefore lie outside the scope of this reasoning mechanism. For example, you want to be able to reason a basketball will roll smoothly in any direction, whereas a football can roll smoothly if its long axis is horizontal but cannot roll smoothly end-over-end.” [13, p. 97].

From the description given above, it seems that, despite the name “qualitative reasoning”, the main focus in this area is quantitative reasoning (“the price of an object”, “the number sold”, “the temperature of a gas”, “the pressure [of the gas]”), not qualitative reasoning. As noted in Section 3.5.4, quantitative reasoning is an aspect of the SP System where more work is needed.

2.1.5 Discussion

With regard to the four areas of relative success with CSRK outlined in the foregoing subsections, there is a problem: *they have apparently been developed quite independently of each other*, the kind of fragmentation in AI that has been lamented by Pamela McCorduck [33, p. 417] and others, and that the SP System has aimed to overcome (Appendix A.1.1). Consequently, there is little or no integration amongst them, and they have little or nothing to say to each other. Since they all deal with

aspects of CSRK, it would be good if there was some kind of overarching theory for them.

For now, the SP System cannot claim to provide that overarching theory. But it has clear strengths and potential with taxonomic reasoning, and it has at least some potential in each of the other three areas.

3 CSRK-related strengths and weaknesses of the SP System

Most of this section describes features of the SP System that appear to be favourable for CSRK. Some apparent weaknesses of the SP System in relation to CSRK, and how they may be overcome, are described in Section 3.5.

3.1 A favourable ratio of Simplicity to descriptive or explanatory Power

As noted in Appendix A.1.1, the SP System is the product of a unique attempt to simplify and integrate observations and concepts across several fields—to develop a system that combines conceptual Simplicity with descriptive or explanatory Power.

As suggested in Appendix A.8.1, this endeavour has been largely successful: it appears that the SP System, with SP-multiple-alignment at centre stage, provides a more favourable combination of ‘Simplicity’ and ‘Power’ than any of the alternatives—with versatility in diverse aspects of intelligence including diverse forms of reasoning, with versatility in the representation of diverse forms of knowledge, and with the seamless integration of diverse aspects of intelligence and diverse forms of knowledge, in any combination. These kinds of versatility and integration are illustrated schematically in Figure 1.

To be more specific:

- A long-term aim in the development of the SP System is to represent *all* kinds of knowledge with SP-patterns, which are simple arrays of atomic symbols in one or two dimension (Appendix A.2).
- Despite the simplicity of SP-patterns, they come to life within the framework of SP-multiple-alignments, demonstrating versatility in the representation of knowledge (Section 3.1.3).
- Because the SP-multiple-alignment construct is a generalised version of six more-specific versions of ICMUP [73, Section 2.1.7], there are reasons to think that, with the aid of SP-multiple-alignments, SP-patterns, including




Figure 1: A schematic representation of versatility and integration in the SP System, with SP-multiple-alignment centre stage. Reproduced with permission from Figure 6 in [73].

2D SP-patterns, have potential to represent all kinds of knowledge, and that would include all kinds of 'commonsense' knowledge.

- In addition to providing a means of realising the expressiveness of SP-patterns, the SP-multiple-alignment construct lies at the heart of the SP System's versatility in aspects of intelligence (Section 3.1.1) including versatility in reasoning (Section 3.1.2), and for the seamless integration of diverse aspects of intelligence and diverse forms of knowledge, in any combination (Section 3.1.4).

- These things suggests that the SP System has potential in CSRK.

3.1.1 Versatility in aspects of intelligence

Largely because of the mechanisms of SP-multiple-alignment, the SP System demonstrates strengths and potential in several aspects of intelligence ([62, Sections 5 to 12], [61, Chapters 5 to 9]) including: unsupervised learning; natural language processing; fuzzy pattern recognition; recognition at multiple levels of abstraction; best-match and semantic forms of information retrieval; several kinds of reasoning (more in Section 3.1.2, below); planning; and problem solving.

How the SP Computer Model may demonstrate some of these aspects of intelligence is shown in Appendix B. The reason for showing these examples in an appendix rather than the main text is to avoid disrupting the main presentation.

Although the aspects of intelligence mentioned in this section do not have the word ‘reasoning’ in their names, it is likely that one or more of them will prove relevant to CSR because it is often difficult to distinguish reasoning from other aspects of intelligence. For example, when Sherlock Holmes infers that Watson has returned to medical practice because he is “smelling of iodoform”, has “a black mark of nitrate of silver upon his right forefinger”, and has “a bulge on the right side of his top-hat to show where he has secreted his stethoscope”,² this may be seen as pattern recognition just as well as it may be seen as reasoning.

3.1.2 Versatility in kinds of reasoning

Although reasoning is an aspect of intelligence (Section 3.1.1), it has been given a subsection to itself because of the versatility of the SP System in this area and because all of the kinds of kinds of reasoning that can be demonstrated with the SP Computer Model are likely to prove useful in CSR.

In reasoning, strengths of the SP System, described quite fully in [62, Section 10] and [61, Chapter 7], are in brief: 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’ (as discussed by Judea Pearl in [42, Sections 1.2.2 and 2.2.4]); causal reasoning; and reasoning that is not supported by evidence.

As described in Appendix B.4, the SP System also supports inference via inheritance of attributes. It appears that there is also potential for spatial reasoning [64, Section IV-F.1], and for what-if reasoning [64, Section IV-F.2].

It would not be either feasible or appropriate to reproduce everything in this area that has been published before. But to see some of the versatility of the SP

²A *Scandal in Bohemia*, Sir Arthur Conan Doyle, 1891.

System in modelling different kinds of reasoning—and corresponding potential in CSR—readers may consult the afore-mentioned sources, [62, Section 10] and [61, Chapter 7], and, in particular:

- How the SP System may model nonmonotonic reasoning and reasoning with default values ([62, Section 10.1], [61, Section 7.7]).
- How the SP System may model Bayesian networks with ‘explaining away’ ([62, Section 10.2], [61, Section 7.8]).

3.1.3 Versatility in the representation of diverse forms of knowledge

Apart from its versatility in aspects of intelligence including forms of reasoning, the SP System combines simplicity in its organisation with versatility in the representation of diverse forms of knowledge.

That last feature is relevant to CSRK because all of the kinds of kinds of knowledge that can be represented with the SP Computer Model are likely to prove useful in CSK. There is potential in the SP System for the representation of other kinds of knowledge, but that potential has not yet been explored.

Despite their simplicity, SP-patterns, within the SP-multiple-alignment framework, have strengths and potential in representing several forms of knowledge, any of which may serve in CSK. These include: the syntax of natural language (see Figure 13); class hierarchies (see Figure 14), class heterarchies (meaning class hierarchies with cross-classification); part-whole hierarchies (see Figure 14); discrimination networks and trees; entity-relationship structures; relational knowledge; rules for use in reasoning; and procedural knowledge.

With SP-patterns in two dimensions, there is potential for representing: diagrams and images; structures in three dimensions [63, Sections 6.1 and 6.2]; and parallel streams of information [64, Appendix C].

There is more detail throughout [62] and [61], and there are references to further sources of information in [65, Section III-B]. Some examples are shown in Appendix B.

3.1.4 Seamless integration of diverse aspects of intelligence and the representation of diverse forms of knowledge, in any combination

Three features of the SP System suggests that it should facilitate the seamless integration of diverse aspects of intelligence (including several forms of reasoning) and seamless integration of diverse kinds of knowledge, in any combination:

- The adoption of one simple format—SP-patterns—for all kinds of knowledge.

- That one relatively simple framework—SP-multiple-alignment—is central in all kinds of processing.
- That the relatively simple format for knowledge, and the SP-multiple-alignment framework for processing, provide for several aspects of intelligence (Section 3.1.1) including several kinds of reasoning (Section 3.1.2), and for the representation of several different kinds of knowledge (Section 3.1.3).

Seamless integration of structures and processes may be seen in concepts relating to Figure 14 (Appendix B.1):

- Class-inclusion relations and part-whole relations work together in the representation of knowledge without awkward incompatibilities (Appendix B.2).
- Pattern recognition (an aspect of general intelligence), class-inclusion relations and part-whole relations (aspects of the representation of knowledge), and inheritance of attributes (a form of reasoning), are intimately related in what is, in effect, one type of operation (Appendices B.1, B.2 and B.4).

For the understanding of natural language and the production of language from meanings, it is likely to be helpful if there is seamless integration of syntax and semantics, and it seems likely that this will be facilitated by representing both of them with SP-patterns, and by processing both of them together via the building and manipulation of SP-multiple-alignments.

Some preliminary examples from the SP Computer Model show how this kind of integration may be achieved with both the understanding and production of natural language [61, Section 5.7]. There is clear potential with the SP System for the comprehensive integration of the syntax and semantics of natural language.

Seamless integration of diverse aspects of intelligence and diverse kinds of knowledge in any combination is likely to be critically important in the modelling of CSRK since, as a matter of ordinary experience, CSR means a willingness to use any and all relevant forms of knowledge, and a willingness to be flexible in one's thinking, using diverse forms of reasoning with diverse kinds of intelligence where appropriate. More generally, that kind of seamless integration appears to be essential in any artificial system that aspires to the fluidity, versatility, and adaptability of human intelligence.

3.2 Turing equivalence, plus aspects of human intelligence

The SP System has clear potential to be Turing-equivalent in the sense that, with some more development, the SP Computer Model would probably be able

to perform any computation within the scope of a universal Turing machine [61, Chapter 4].

But, unlike a ‘raw’ Turing machine (without any programming), or a ‘raw’ conventional computer, the SP System has strengths and potential in AI (Appendix A.8, Section 3.1), something which, as Turing recognised ([54, 55, 57]), is missing from the original concept of a universal Turing machine.

Thus the SP System has the kind of generality needed for CSRK, and its strengths and potential in AI give it a head start in modelling CSRK.

3.3 Information compression and CSRK

As described in Appendix A.2, a central part of the SP System is ICMUP and, more specifically, ‘information compression via SP-multiple-alignment’.

For CSRK, IC in the SP System has a three-fold significance:

- *IC and concepts of inference and probability.* As noted in Appendix A.1.5, the intimate connection that is known to exist between IC and concepts of inference and probability provides for the strengths of the SP System in the making of inferences and in the calculation of their probabilities.
- *Generality in the representation of knowledge.* The generality of IC suggests that, in principle, *any* kind of knowledge may be represented in a succinct form in the SP System.
- *The DONSVIC principle.* Unsupervised learning in the SP Computer Model as it has been developed to date, conforms to the DONSVIC principle: the *Discovery Of Natural Structures Via Information Compression* ([62, Section 5.2]). It seems likely that DONSVIC forms of knowledge are those that are most relevant to CSRK.

3.4 Versions of logic, and the SP System

As noted in Section 2, an important strand of CSRK-related research is in the application of VLSs to problems in CSRK (eg, [11, 37, 12]). There are undoubted benefits from this research but there appears to be greater long-term promise in the SP System.

3.4.1 Benefits of VLSs

CSRK-related research with VLSs has yielded at least five useful results:

- As noted in Section 8, below, three different CSRK-related research projects using VLSs have shown that something as apparently simple as cracking an egg into a bowl is remarkably complicated. As Murray Shanahan [46, p. 142] has argued:
 - It is no longer plausible to argue that it is feasible or reasonable to encode such knowledge by hand.
 - Armchair theorising in this area must be tested in computer models.
- Ernest Davis points out [12, pp. 710–711]:
 - That there are many forms of CSK that, as far as anyone knows, can only be expressed in one or other of the several VLSs. And there are many examples of CSR that, as far as anyone knows, can only be characterised via the use of a VLS.
 - That CSRK-related research with VLSs serves as a useful reminder of the relative lack of precision and transparency in the workings of deep learning networks, both in the way that they represent knowledge, and in the way that they draw inferences from their knowledge.

3.4.2 Apparent shortcomings of VLSs compared with the SP System

Notwithstanding the strengths of VLSs in CSRK-related research, as outlined above, it appears that the SP System holds out more promise as a foundation for an understanding of CSRK, either natural or artificial.

In general, on the strength of evidence summarised in Section 3.1, it appears that the SP System has *a more favourable ratio of Simplicity to Power than any other AI-related system, including VLSs*. More specifically:

- *Psychological validity*. As a general rule, research with VLSs has paid little or no attention to research on the learning, storage, or use of knowledge by human brains. Thus Erik Mueller writes:

“I do not claim that the methods for commonsense reasoning presented in this book are the methods used by humans. This book presents one way of automating commonsense reasoning. How humans perform commonsense reasoning is an interesting topic, but it is not the topic of this book. There is evidence both for and against the use of logic in human reasoning.” [37, p. xxii].

At least with the benefit of hindsight in developing the SP System, it seems that VLSs, in their development, could have benefitted from closer attention to what is known about human cognitive psychology.

- *Learning.* Each VLS says nothing about how its knowledge might be learned. In keeping with the previous point (‘psychological validity’), a learning perspective may have been of value in the development of VLSs, as it has been in the development of the SP System.
- *Conceptual primitives.* Like any natural or artificial computational system, a VLS requires conceptual primitives from which other concepts may be built. For example, the ‘ALC’ kind of ‘Description Logic’ has primitive concepts like ‘conjunction’, ‘disjunction’, ‘negation’, ‘existential restriction’, and ‘value restriction’ [3, p. 140].

In terms of Simplicity and Power (Section 3.1), these VLS conceptual primitives seem to be simpler than concepts like ‘IC’ and ‘ICMUP’ (Section A.2) and ‘SP-multiple-alignment’ (Section A.3) which form the foundations of the SP System. But:

- IC and perhaps ICMUP have substantial empirical support as important features of HLPC [73].
- The descriptive and explanatory power of the SP System considerably outweighs the ‘cost’ of its conceptual foundations (Section 3.1).
- Although no attempt has been made to examine whether or how Description Logics or other VLSs may be modelled within the SP System, it has been shown that several of the concepts that are familiar in mathematics, logic, and computing, may be seen as emergent properties of the SP System and how it works. These include: ‘function’, ‘variable’, ‘value’, ‘set’, and ‘type definition’ [69, Appendix B.2]; and the union and intersection of sets [72, Section 5.1.3].

3.5 Apparent shortcomings of the SP System with CSRK, and how they may be overcome

The SP System is a work in progress, with shortcomings in areas that have not yet been developed. The following subsections describe the shortcomings which are most relevant to CSRK, with indications of how those shortcomings may be overcome, mainly via planned developments in the SP System, described in [41].

3.5.1 The representation and processing of information in two or more dimensions

At present, the SP Computer Model works only with one-dimensional SP-patterns, although it has been envisaged from the outset that it would be generalised to work with two-dimensional SP-patterns (Appendix A.2). That development will allow the system to work with pictures and diagrams. Less obviously, it will also allow the system to encode information in three dimensions, in the manner of commercial programs that create 3D computer models from photographs, as outlined in [63, Section 6].

As noted in [64, Section IV-F.1], the encoding of 3D structures and spaces in the SP System will open up possibilities for spatial reasoning, such as for example, the ‘mental’ exploration of how furniture may be arranged in a room, much as people sometimes do using cardboard shapes to represent furniture, with a plan of a room.

How SP-patterns in the SP Computer Model may be generalised to two dimensions is described in [41, Section 9].

3.5.2 Recognition of low-level perceptual features in sound and visual images

The SP Computer Model is designed to work with data composed of atomic SP-symbols at the most fine-grained level, where an SP-symbol is a mark that can be matched in a yes/no manner with any other SP-symbol.

This works well with SP-symbols like ‘a’, ‘b’, ‘c’, and so on, or even ‘house’, ‘car’, ‘person’, and so on. And it is clearly applicable to features of speech like ‘formant ratio’ or ‘formant transition’, and visual features such as ‘edge’ or ‘corner’.

But it less clear how the symbolic aspects of sound or vision may be isolated from the fuzziness of raw auditory or visual input, or whether SP principles apply, and if so, how (Appendix A.7). How the SP System may be developed for the recognition of low-level features in sound and vision is described in [41, Section 10]).

3.5.3 Unsupervised learning

As noted in Appendix A.5, the SP Computer Model 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. But there is further work to be done in this area.

With respect to CSRK, unsupervised learning is important for two main reasons:

- A fully developed theory of unsupervised learning in the SP System will be needed for a comprehensive account of how the SP System may be applied in CSRK.
- Although some progress can be made by researchers in defining CSK structures, that process is far too laborious and prone to error to be satisfactory in the long run. It will be essential, with the SP System or any other AI system, to replace that kind of manual coding with robust and effective automatic coding via a well-developed version of unsupervised learning. There are further remarks about this issue in Section 8.

How unsupervised learning may be further developed in the SP Computer Model is described in [41, Section 12].

3.5.4 The representation and processing of numbers and magnitudes

Although the SP System has some potentially useful things to say about the nature of mathematics [61, Chapter 10] and has led to the proposal that much of mathematics, perhaps all of it, may be understood as a set of techniques for the compression of information and their application [72], the SP Computer Model is not yet an effective means of representing numbers and doing such basic things as addition, subtraction, multiplication, and division (Appendix A.7).

Likewise, there is more work to be done in the SP programme of research on the representation and processing of concepts that relate to magnitudes, such as motion and speed, distances, volumes, time, and so on. It is envisaged that the development of concepts in this area may be approached via the further development of unsupervised learning in the SP Computer Model ([41, Sections 9.3, 12, and 14], [63, Sections 5.3 and 6]). There are some notes on a possible way forward in Appendix C.

4 “The meaning of noun phrases”

This main section and the following three main sections describe CSRK problems where the SP System is relatively successful. After that, Section 8 describes the CSRK problem of modelling how a cook may crack an egg into a bowl which, despite its apparent simplicity, is well beyond what the SP system can do as it is now, but where planned future developments may help.

“The meaning of noun phrases” is a CSRK problem contributed by Ernest Davis to the ‘Commonsense Reasoning Problem Page’ (bit.ly/2qjdMBj), described as follows:

“There are many ways in which the meaning of a two word noun phrase can be related to the meanings of the individual nouns, and syntax gives little indication of which applies in any given case. Some such phrases are purely idiomatic and must be individually learned (e.g. ‘tag sale,’ ‘mustard gas’) but in most cases a speaker who has never seen the particular phrase can figure out its meaning from semantic constraints and commonsense knowledge.

“Characterize the commonsense knowledge used in determining that the correct meaning of the following noun phrases is more plausible than any of the alternative readings:

- water bird (a bird who[sic] lives near the water)³
- marble cake (a cake that looks like marble)
- soda can (a can containing soda)”

(and many more examples, the whole quote retrieved 2018-12-09).

Only the first example will be considered because variations on the proposed solution would be applicable to the other examples in a similar way.

4.1 Learning and interpretation

Although the problem as described above is that “Some ... phrases are purely idiomatic and must be individually learned (e.g. ‘tag sale,’ ‘mustard gas’) but in most cases a speaker who has never seen the particular phrase can figure out its meaning from semantic constraints and commonsense knowledge,” the suggestion here is that the emphasis should be reversed: so that there is more focus on the learning of phrases of all kinds, both idiomatic phrases and those with a more transparent meaning, with less emphasis on interpretation. In accordance with that suggestion, learning and interpretation of phrases may be seen to develop in roughly three stages:

1. *Learning and interpretations with limited knowledge.* With regard to the ‘water bird’ example, a child may learn the names of things like ‘water’ and ‘bird’ but if they have never been to feed ducks on a pond, or done anything similar, they might conclude that ‘water bird’ means a bird that drinks water, or a bird that is made of water, or it might perhaps be the name of a boat. In a similar way, a child with limited experience—perhaps knowing only the meaning of “cake” but not “marble”—might conclude that a ‘marble cake’

³Apart from the instance just shown, “who” has been replaced by “that” in the rest of this paper when the reference is to an animal, in accordance with general usage.

is a cake from a place called ‘marble’, or a cake that is made of ‘marble’, and so on.⁴ And again, a child who has not yet learned the meaning of ‘soda’ might conclude that a ‘soda can’ is a can made of soda, or a can from a place called ‘soda’, and so on.

2. *Learning and interpretations with greater knowledge.* When a child has had the experience of feeding ducks and the like on a pond or lake, or has seen some marble, or has drunk some soda and learned that name, he or she may make guesses at the interpretations of expressions like ‘water bird’, ‘marble cake’, ‘soda can’, which are closer to their meanings for adults.
3. *Developing a mature knowledge of phrases and their meanings.* The new interpretations just described are likely, at first, to be provisional:
 - (a) At first they will be in competition with earlier, naive interpretations such as ‘a bird that drinks water’ as the meaning of ‘water bird’.
 - (b) Further learning will be required for the ‘correct’ interpretations to become firmly established.

In short, the suggestion here is that processes of learning are likely to be prominent in how most people come to understand the meaning of expressions of all kinds, both commonly-used phrases with relatively transparent meanings, as well as more obscure phrases such as “tag sale” and “mustard gas”. Interpretation would also be at work at the same time.

This main section does not attempt to model all that complexity. Instead, the main focus is on showing how the SP System can model processes of interpretation in Stage 3 (a)—where ‘correct’ interpretations of phrases have to compete with naive interpretations. In Section 4.6, there are some observations about the role of learning in all three main stages.

4.2 The effect when frequencies of occurrence of SP-patterns are all equal

With the first example above, “water bird (a bird that lives near the water)”, we would naturally think that the given interpretation was the most plausible. But that is almost certainly because of our extensive knowledge of English, and of birds and their anatomy and how they live.

⁴I can remember as a child being completely baffled by the following words in a hymn: “There is a green hill far away without a city wall”. Not being familiar with the Scottish word “outwith” and its meaning, it seemed nonsensical to say that a green hill did not have a city wall, since green hills that have city walls are vanishingly rare.

For present purposes, we shall put ourselves in the shoes of an ignorant robot (or young child) that is trying to figure out some of the more basic features of the syntax and semantics of English. Let us suppose that the robot has arrived at three possible interpretations of ‘water bird’: a bird that lives near the water; a bird that drinks water; and a bird that is made of water. And, failing any evidence to the contrary, the robot may assume that those three interpretations occur with equal frequency in English.

When the SP Computer Model is run with the New SP-pattern ‘a w a t e r b i r d’ and a collection of relevant Old SP-patterns, all of them with a frequency value of 1, the best three SP-multiple-alignments that it creates are those shown in Figures 2, 3, and 4, *all of them with the same high compression score, which is higher than for any other SP-multiple-alignment produced in this run of the program.* They are described, one in each of the three subsections that follow.

4.2.1 Bird lives near water (all Old SP-patterns have a frequency of 1)

This subsection describes one of the three best SP-multiple-alignments created by the SP Computer Model when all the Old SP-patterns have been assigned a frequency of 1. That SP-multiple-alignment is shown in Figure 2.

In that figure, in the SP-pattern ‘N n5 w a t e r WATER #N’ in column 4, the pair of SP-symbols ‘N ... #N’ marks the structure as a noun, the SP-symbols ‘... w a t e r ...’ represents the surface form of the noun, and ‘... WATER ...’ is merely a label for the meaning of water, without any attempt to show the complexity of water and all of its structure, properties and associations. In a similar way, the robot’s concept of a bird may be represented by the SP-pattern ‘N n8 b i r d BIRD #N’ in column 1.

In general, every SP-pattern for a word contains one or more lower-case letters representing the surface structure of the word, followed by an SP-symbol in upper-case letters that represents the meaning of the word—and likewise for more complex structures.

In the figure, the pair of SP-symbols ‘NP ... #NP’, within the SP-pattern ‘NP np0 D #D N #N N #N npSEM #npSEM #NP’ in column 2, marks the structure as a noun phrase. Within that SP-pattern in column 2:

- The pair of SP-symbols ‘D #D’ in column 2 is a slot for a ‘determiner’ which may be seen to be the word ‘a’ within the SP-pattern ‘D d5 a INDEF #D’ in column 7.
- The first of two pairs of SP-symbols ‘N #N’ in column 2 marks a slot for a noun which may be seen to be the word ‘water’ within the SP-pattern ‘N n5 w a t e r WATER #N’ in column 4.




Figure 2: One of the best three SP-multiple-alignments created by the SP Computer Model with the New SP-pattern 'a w a t e r b i r d' (in column 0) and Old SP-patterns as described in the text, all with a frequency of 1. The other two SP-multiple-alignments are shown in Figures 3 and 4.

- The second of two pairs of SP-symbols ‘N #N’ in column 2 marks another slot for another noun which may be seen to be the word ‘bird’ within the SP-pattern ‘N n8 b i r d BIRD #N’ in column 1.
- Lastly, the pair of symbols ‘npSEM #npSEM’ in column 2 provide a slot for a semantic structure, which is the SP-pattern ‘npSEM npsm1 BRD LIVES_NEAR WTR #npSEM’ in column 3. That semantic structure may be seen to represent the meaning ‘bird that lives near water’.

A feature of this SP-multiple-alignment that needs some explanation is the role of the SP-pattern ‘n5 WTR’ in column 5, and of ‘n8 BRD’ in column 6. In broad terms, these two SP-patterns have the effect of transferring semantic information from the SP-patterns ‘N n5 w a t e r WATER #N’ in column 4 and ‘N n8 b i r d BIRD #N’ in column 1 to the ‘semantic’ slot, ‘npSEM #npSEM’, within the SP-pattern ‘NP np0 D #D N #N N #N npSEM #npSEM #NP’ in column 2. More specifically:

- With regard to the SP-pattern ‘n5 WTR’ in column 5:
 - The SP-symbol ‘n5’ matches the same SP-symbol in ‘N n5 w a t e r WATER #N’ (in column 4) and thus helps to select that SP-pattern, including the semantic information ‘WATER’.
 - The SP-symbol ‘WTR’ matches the same SP-symbol in the SP-pattern ‘npSEM npsm1 BRD LIVES_NEAR WTR #npSEM’ (in column 3) and thus helps to select that SP-pattern as the semantic information to fill the ‘semantic’ slot ‘npSEM #npSEM’, within the SP-pattern ‘NP np0 D #D N #N N #N npSEM #npSEM #NP’ (in column 2).
- With regard to the SP-pattern ‘n8 BRD’ in column 6:
 - The SP-symbol ‘n8’ matches the same SP-symbol in ‘N n8 b i r d BIRD #N’ (in column 1) and thus helps to select that SP-pattern, including the semantic information ‘BIRD’.
 - The SP-symbol ‘BRD’ matches the same SP-symbol in the SP-pattern ‘npSEM npsm1 BRD LIVES_NEAR WTR #npSEM’ (in column 3) and thus helps to select that SP-pattern as the semantic information to fill the ‘semantic’ slot ‘npSEM #npSEM’, within the SP-pattern ‘NP np0 D #D N #N N #N npSEM #npSEM #NP’ (in column 2).
- For the sake of clarity and relevance, no attempt has been made with this example to transfer the semantic information ‘INDEF’ within the SP-pattern ‘D

d5 a INDEF #D' (in column 7) to the 'semantic' slot 'npSEM #npSEM', within the SP-pattern 'NP np0 D #D N #N N #N npSEM #npSEM #NP' (in column 2).

As indicated above, the SP-multiple-alignments in Figures 3 and 4 may be interpreted in a similar way. And the same kind of interpretation may be applied to Figure 5 in Section 4.3, and to Figure 6 in Section 4.4.

4.2.2 Bird drinks water (all Old SP-patterns have a frequency of 1)

This subsection describes another of the three best SP-multiple-alignments created by the SP Computer Model when all the Old SP-patterns have been assigned a frequency of 1. That SP-multiple-alignment, which has the same compression score as the SP-multiple-alignments in Figures 2 and 4, is shown in Figure 3.

This SP-multiple-alignment is much the same as the SP-multiple-alignment in Figure 2 but the 'semantic' SP-pattern 'npSEM npsm0 BRD DRINKS WTR #npSEM' (in column 3) has been selected to fill the 'semantic' slot 'npSEM #npSEM' within the SP-pattern 'NP np0 D #D N #N N #N npSEM #npSEM #NP' (in column 2). The whole SP-multiple-alignment may be seen to assign the meaning 'a bird that drinks water' to the surface form 'a w a t e r b i r d'.

4.2.3 Bird made of water (all Old SP-patterns have a frequency of 1)

This subsection describes the third of the three best SP-multiple-alignments created by the SP Computer Model when all the Old SP-patterns have been assigned a frequency of 1. That SP-multiple-alignment, which has the same compression score as the SP-multiple-alignments in Figures 2 and 3, is shown in Figure 4.

This SP-multiple-alignment is similar to the SP-multiple-alignments in the two preceding figures but the 'semantic' SP-pattern 'npSEM npsm2 BRD MADE_OF WTR #npSEM' (in column 3) has been selected to fill the 'semantic' slot 'npSEM #npSEM' within the SP-pattern 'NP np0 D #D N #N N #N npSEM #npSEM #NP' (in column 2). The whole SP-multiple-alignment may be seen to assign the meaning 'a bird that is made of water' to the surface form '[a] w a t e r b i r d'.

What we may learn from these three SP-multiple-alignments is the not-very-surprising result that our inexperienced robot would conclude that its guess about the frequencies of occurrence of SP-patterns representing the three hypothesised interpretations of 'water bird' confirm that, in terms of the robot's elementary knowledge of English and of birds and so on, there is nothing to choose between the three hypothesised interpretations.

Of course, people with a good knowledge of English and of birds are much more likely to favour the 'a bird that lives near water' interpretation than the other two




Figure 3: One of the best three SP-multiple-alignments created by the SP Computer Model with the New SP-pattern 'a w a t e r b i r d' (in column 0) and Old SP-patterns as described in the text, all with a frequency of 1. The other two SP-multiple-alignments are shown in Figures 2 and 4.




Figure 4: One of the best three SP-multiple-alignments created by the SP Computer Model with the New SP-pattern 'a w a t e r b i r d' (in column 0) and Old SP-patterns as described in the text, all with a frequency of 1. The other two SP-multiple-alignments are shown in Figures 2 and 3.

interpretations. The following two sections demonstrate how our robot may come to favour the ‘correct’ interpretation of ‘water bird’.

4.3 The effect of unequal frequencies of occurrence

It is plausible to suppose that ‘a bird that lives near the water’ as the interpretation of ‘water bird’ is much more frequent in the way English is used than ‘a bird that drinks water’ or ‘a bird that is made of water’. To reflect that aspect of English usage with the SP Computer Model, it has been run again with the same New and Old SP-patterns as in Section 4.2 but with frequencies assigned to the three critical SP-patterns as shown in Table 1.

<i>SP-pattern</i>	<i>Frequency</i>
‘npSEM npsm1 BRD LIVES_NEAR WTR #npSEM’	1000
‘npSEM npsm0 BRD DRINKS WTR #npSEM’	15
‘npSEM npsm2 BRD MADE_OF WTR #npSEM’	1

Table 1: Frequencies of occurrence assigned to three SP-patterns amongst the Old SP-patterns used by the SP Computer Model in creating the SP-multiple-alignment shown in Figure 5. The other Old SP-patterns for that run of the program were each assigned a frequency of 1.

No doubt these frequencies are inaccurate for any realistically large sample of English but they should be sufficiently out of balance with each other for present purposes.

In this case, the best SP-multiple-alignment created by the SP Computer Model is the one shown in Figure 5. In this case there is no tie for the first place: all other SP-multiple-alignments created by the program, including those that feature the SP-patterns ‘BRD DRINKS WTR’ and ‘BRD MADE_OF WTR’, have a lower compression score. In other words, a relatively high frequency for the SP-pattern ‘BRD LIVES_NEAR WTR’ means that the best SP-multiple-alignment created by the system, unrivalled by any other, accords with the normal interpretation of ‘water bird’ as ‘a bird that lives near the water’.

It seems reasonable to suppose that, in general, relevant frequencies of occurrence are learned from context and usage. For example, if the expression ‘water bird’ is used predominantly when ducks, geese, moorhens, and so on, are in view or otherwise the focus of attention, then the robot or child would naturally assign a higher frequency to the association of ‘water bird’ with living near water than other associations that are less frequent or vanishingly rare.

In general, this example demonstrates how, in the process of interpreting incoming (‘New’) information, the SP Computer Model is sensitive to the frequencies




Figure 5: The best SP-multiple-alignment created by the SP Computer Model with the New SP-pattern 'a w a t e r b i r d' (in column 0) and Old SP-patterns as in Figures 2, 3, and 4. But in this case, the frequencies of three of the Old SP-patterns are as shown in Table 1, and the rest are 1.

of occurrence of its stored ('Old') items of information. This is broadly consistent with the 'word frequency effect', the long-established psychological phenomenon in which recognition of words by people who are familiar with a given language is, visually or via hearing, and by a variety of measures, more efficient for words that occur frequently in the given language than for those that are rare (see, for example, [9]).

4.4 The effect of disambiguating context

Another thing that may tip the balance towards the normal interpretation of 'water bird' is the linguistic or physical context in which those words are spoken.

Figure 6 shows the best SP-multiple-alignment created by the SP Computer Model with the New SP-pattern 'a w a t e r b i r d f i s h e s' and Old SP-patterns representing relevant syntax and semantics like those that appear in columns 1 to 11 in the figure. By contrast with the example in Figure 5 in Section 4.3, the frequencies of occurrence of all the Old SP-patterns in this example were set to 1.

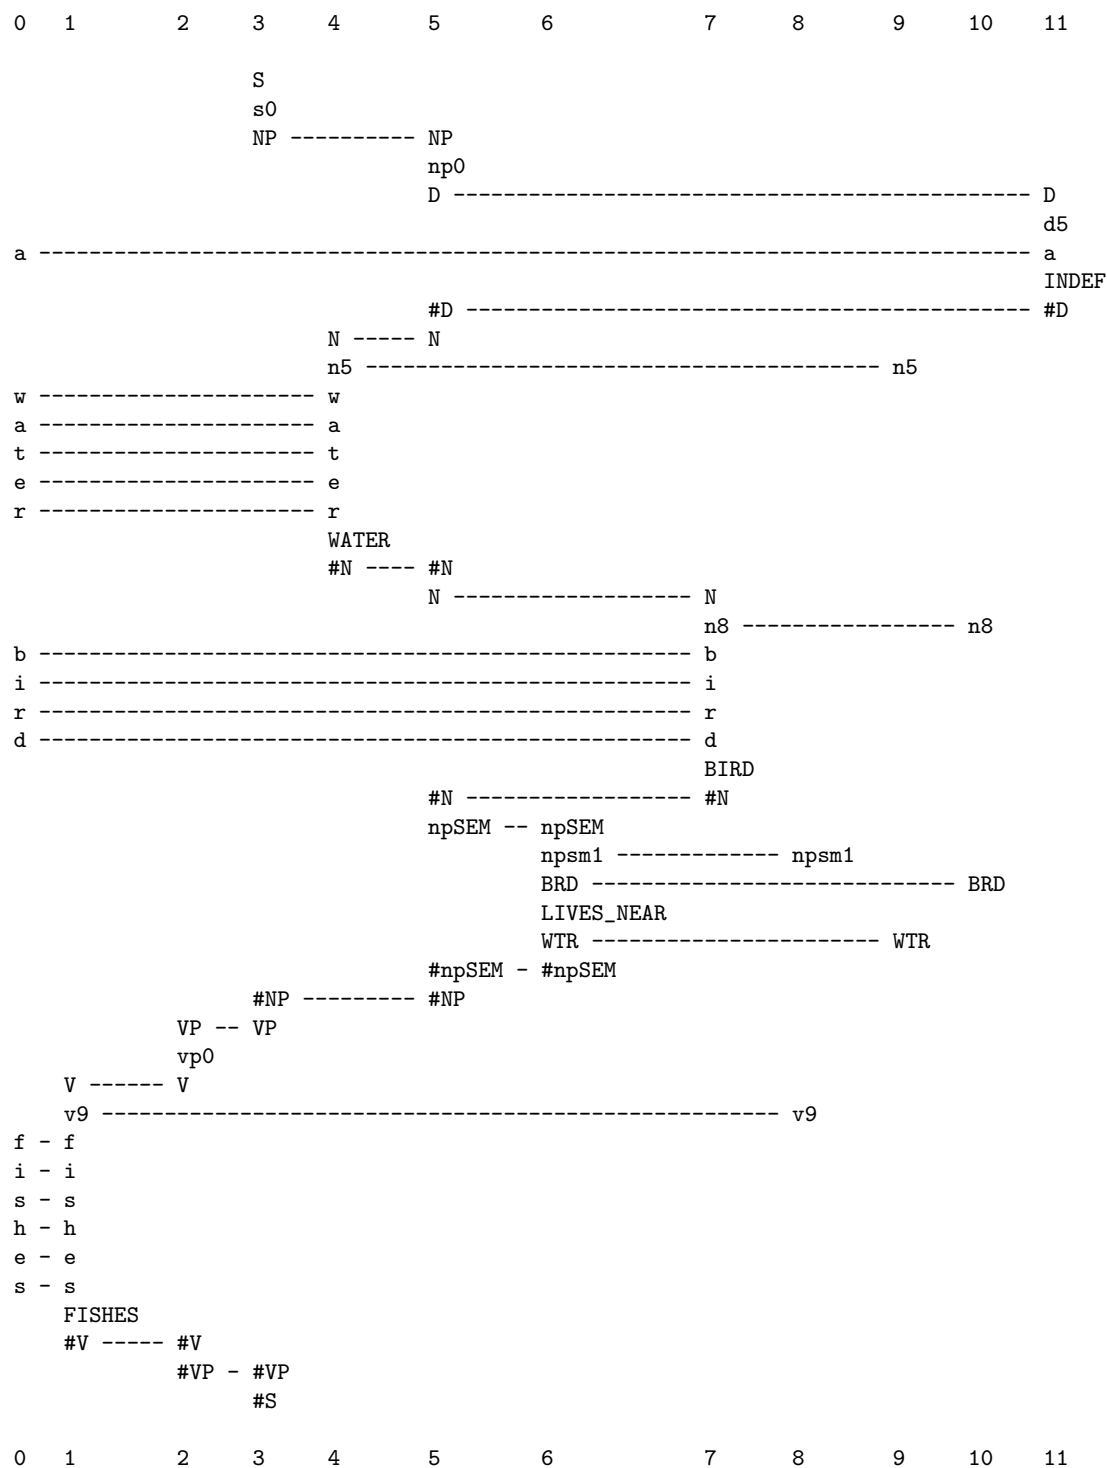


Figure 6: The best SP-multiple-alignment created by the SP Computer Model with the New SP-pattern 'a w a t e r b i r d f i s h e s'

and Old SP-patterns representing relevant syntax and semantics, like those that appear in columns 1 to 11 in the figure.

As with the example discussed in Section 4.3, the best SP-multiple-alignment in this case confirms, via the SP-pattern ‘npSEM npsm1 BRD LIVES_NEAR WTR #npSEM’ (in column 6 in this example), that the most favoured interpretation of “a water bird” is ‘a bird that lives near the water’. In this case, by contrast with the example in Section 4.2, there is no other SP-multiple-alignment that rivals the one shown in the figure.

The reason that the context ‘... f i s h e s’ raises the compression score of the SP-multiple-alignment in the figure is that the SP-pattern, ‘npsm1 v9’, which appears in column 8, shows in effect that there is an association between ‘birds living near water’ and ‘birds catching fish’.

This is because the SP-symbol ‘npsm1’ is, in effect, a label for the SP-pattern ‘npSEM npsm1 BRD LIVES_NEAR WTR #npSEM’ in column 6, and because the SP-symbol ‘v9’ is, in effect, a label for the SP-pattern ‘V v9 f i s h e s FISHES #V’ in column 1. And it is also because the alignment of the SP-symbol ‘npsm1’ in column 8 with the same SP-symbol in column 6, and the alignment of the SP-symbol ‘v9’ in column 8 with the same SP-symbol in column 1, has the effect of raising the compression score for the SP-multiple-alignment, in accordance with the method for calculating compression scores described in [62, Section 4.1] and [61, Section 3.5].

This effect of context in the workings of the SP Computer Model is broadly in accordance with the long-established effect of context in the way that people recognise things (see for example [51]).

4.5 Disambiguating context versus unequal frequencies of occurrence

With the SP Computer Model, a question that arises from the effect of unequal frequencies of occurrence of SP-patterns (Section 4.3) and from a disambiguating context (Section 4.4) is what happens if we combine those two things. As a first step in answering that question, the SP Computer Model was run using the same New and Old SP-patterns as in the example considered in Section 4.4 (the effect of disambiguating context) but with frequencies assigned to the three critical SP-patterns as shown in Table 2. As before, the frequencies of the other Old SP-patterns were set to 1.

In this case, the best SP-multiple-alignment found by the SP Computer Model is exactly the same as the one shown in Figure 6, and there is no other SP-multiple-alignment to rival it. The result shows that, in this case, the context ‘... f i s h e s’ strongly favours ‘a bird that lives near water’ interpretation of ‘water bird’, and completely obscures the effect of frequency.

The example in Section 4.3, with the examples in Section 4.4 and this section,

<i>SP-pattern</i>	<i>Frequency</i>
'npSEM npsm1 BRD LIVES_NEAR WTR #npSEM'	1
'npSEM npsm0 BRD DRINKS WTR #npSEM'	15
'npSEM npsm2 BRD MADE_OF WTR #npSEM'	1000

Table 2: Frequencies of occurrence assigned to three SP-patterns amongst the Old SP-patterns used by the SP Computer Model in creating the best SP-multiple-alignment (not shown) which is discussed in this section. The other Old SP-patterns for that run of the program were each assigned a frequency of 1.

are broadly consistent with experimental studies with people which show that both frequency and context have an influence of how we recognise things [40, 6]. But for a more detailed picture, more research would be needed into the relative strengths of the influences of frequency and context in both people and the computer model.

4.6 Learning the meanings of phrases

Section 4.1 suggests that the process of developing an adult knowledge of phrases of all kinds, including their meanings, may be seen to occur in roughly three stages: 1) learning the names of things like water and birds; 2) learning about the kinds of places that ducks and the like often live, and learning the association between 'water bird' and those kinds of bird; and 3) learning the adult meaning of phrases like 'water bird' and discarding immature alternatives.

Although unsupervised learning by the SP Computer Model has been developed to a point where, as mentioned in Appendix A.5, it can learn generative grammars from unsegmented samples of English-like artificial languages, including segmental structures, classes of structure, and abstract patterns, it has deficiencies, also mentioned in that section. These deficiencies mean that it is not yet feasible to apply unsupervised learning in the SP Computer Model as it is now to the kinds of learning described above, although some shortcuts and approximations may be feasible.

5 “Strength of evidence”

'Strength of evidence' is another CSRK problem contributed by Ernest Davis to the 'Commonsense Reasoning Problem Page' (bit.ly/2qjdMBj), described as follows:

“A says that he witnessed B murdering C.

Infer that the evidence that B actually did murder C is stronger:

- If the murder was well lit than if it was in the dark.
- If A already knew B than if he was a stranger.
- If A was sober at the time of the murder than if he was drunk.
- If A is known as a man of good character than if he has previously been convicted of perjury.
- If A has no personal connection to B than if they are enemies.
- If A is testifying under oath than if he is talking casually.”

(retrieved 2018-12-09).

As with the type of problem considered in Section 4, only the first of these pairs of alternative scenarios will be considered—in Sections 5.1 and 5.2—since solutions via the SP System for the other pairs of scenarios would be similar.

5.1 The evidence is strong that B actually did murder C if the murder was well lit

Figure 7 shows the best SP-multiple-alignment created by the SP Computer Model with the New SP-pattern ‘event murder agent B #agent victim C #victim #murder witness A illumination bright #illumination #witness #event’ (which appears in column 0) and a collection of Old SP-patterns describing aspects of the problem (some of which appear in columns 1 to 14, one SP-pattern per column).

As noted in Appendix A.3.1, this figure may be magnified and remain sharp at any convenient size, and this may be done with the copy of the figure in the file ‘spsrk2_figures.pdf’ so that everything else in the main paper need not be magnified. Likewise for Figure 8.

The New SP-pattern may be seen as an expression of the main meanings behind “A says that he witnessed B murdering C”, as described in the next paragraph. The reason for using meanings rather than the surface form of the sentence is to avoid complicating the example by showing how the surface form may be translated into corresponding meanings.

The pair of SP-symbols ‘event ... #event’ at the beginning and end of the New SP-pattern in column 0 mean that the SP-pattern is describing an ‘event’. The SP-symbols ‘murder agent B #agent victim C #victim #murder’ mean that the event is a ‘murder’, performed by an ‘agent’ identified as ‘B’ (marked via the three SP-symbols ‘agent B #agent’), on the ‘victim’ identified as ‘C’ (marked via the three SP-symbols ‘victim C #victim’). And the SP-symbols ‘witness A illumination bright #illumination #witness’ mean that the event was ‘witnessed’ by ‘A’ with ‘illumination’ that was ‘bright’.

```

0         1         2         3         4         5         6         7         8         9         10        11        12        13        14        15        16
event ----- event
murder ----- what ----- what
agent ----- agent ----- murder ----- agent
B ----- person -- person -- person
          pn2
          name ---- name
          B
          #name --- #name
          #mammal
          #mammal
          ...
          #person - #person - #person
#agent ----- #agent ----- #agent
victim ----- victim ----- victim
          person -- person ----- person
          pn3
          name ---- name
          C
          #name --- #name
          #mammal
          #mammal
          ...
          #person - #person ----- #person
#victim ----- #victim ----- #victim
#murder ----- #murder ----- #murder
          #what ---- #what
          time
          #time
          place
          #place
          evidence ----- evidence ----- evidence
          encl
witness ----- witness ----- witness
          person -- person ----- person
          pn1
          name ---- name
          A
          #name --- #name
          #mammal ----- #mammal
          cats
          #cats
          #seats
          #breathes
          #breathes
          sees ----- sees
          ss0
illumination -- illumination
bright ----- bright
#illumination - #illumination
          vision ----- vision ----- vision
          good ----- good ----- good
          #vision ----- #vision ----- #vision
          #sees ----- #sees
          ...
          #mammal ----- #mammal
          ...
          #person - #person ----- #person
#witness ----- #witness ----- #witness
          strength ----- strength ----- strength
          #strength ----- #strength ----- #strength
          #evidence ----- #evidence ----- #evidence
#event ----- #event
          #event
0         1         2         3         4         5         6         7         8         9         10        11        12        13        14        15        16

```

Figure 7: The best SP-multiple-alignment created by the SP Computer Model, with input for “A says that he witnessed B murdering C” when the murder was well lit. Then, the evidence is stronger, as discussed in the text. As noted in Appendix A.3.1, it has been necessary, with this figure, and Figure 8, to shrink the SP-multiple-alignment to ensure that it will fit on the page. But it may be magnified as described in Appendix A.3.1.

As will be explained, the SP-multiple-alignment in Figure 7 confirms that it is reasonable to infer that the evidence that ‘B’ actually did ‘murder’ ‘C’ is (relatively) strong if the murder was well lit. Before getting to that inference, we need to see how the meanings in the New SP-pattern in column 0 are reflected in the SP-multiple-alignment in the figure. Those meanings are expressed by SP-patterns in the SP-multiple-alignment and their interconnections, as shown in Table 3.^{5,6} Incidentally, there is no attempt in this analysis, and the analysis in Section 5.2, to model the possibility that the ‘vision’ of ‘A’ is itself ‘poor’.

From the SP-multiple-alignment shown in Figure 7, the inference that the ‘evidence’ is ‘strong’ that ‘person’ ‘B’ actually did murder ‘person’ ‘C’ may be made like this:

1. The first step is to note that ‘illumination bright #illumination’ in the New SP-pattern is matched with the same three SP-symbols in the SP-pattern ‘sees ss0 illumination bright #illumination vision good #vision #sees’ in column 1.
2. Since ‘illumination bright #illumination’ in column 1 is immediately followed by ‘vision good #vision’ in the same SP-pattern, we may infer directly that, assuming that ‘person’ ‘A’ had normal eyesight, his or her ‘vision’ of the ‘murder’ was ‘good’.
3. We can see that the three symbols ‘vision good #vision’ in column 1 are matched with the same three symbols in the SP-pattern representing ‘evidence’ in column 16: ‘evidence enc1 witness vision good #vision #witness strength strong #strength #evidence’.
4. Since, within that SP-pattern for ‘evidence’, the SP-symbols ‘vision good #vision’ are followed later by the three SP-symbols ‘strength strong #strength’ we may infer that the witnessing of the ‘murder’ by ‘A’ has provided ‘strong’ ‘evidence’ that that the ‘event’ did indeed occur.

⁵A point of interest about Figure 7 and Table 3 is that, although there is only one instance of the SP-pattern ‘person name #name mammal #mammal ... #person’ amongst the Old SP-patterns for this example, it appears three times in the SP-multiple-alignment, in columns 4, 10, and 14. The first appearance connects with the SP-pattern for ‘person’ ‘A’ in column 3, the second connects with the SP-pattern for ‘person’ ‘C’ in column 9, and the third connects with the SP-pattern for ‘person’ ‘B’ in column 13. That a given Old SP-pattern may appear two or more times in an SP-multiple-alignment is not an error. It is intrinsic to the concept of ‘SP-multiple-alignment’, as described in [61, Section 3.4.6].

⁶As a step in the hierarchy of classes which includes ‘vertebrates’, ‘animals’, and ‘living things’, each of the appearances of the SP-pattern ‘person name #name mammal #mammal ... #person’ in Figure 7, together with the SP-pattern for a specific person, provides a simple example of how the SP System works well in the representation and processing of class hierarchies, noted in Section 2.1.1 and illustrated in Appendices B.1, B.2, and B.4.

<i>Column</i>	<i>What</i>	<i>Connections</i>
1	This SP-pattern represents the concept of ‘seeing’.	It is a feature of ‘mammals’ (in column 2) and thus, via the connections noted in row 2 of the table, it is a feature of person ‘A’.
2	The concept of a ‘mammal’.	
3	A specific person (‘pn1’) with a ‘name’ (‘A’).	This SP-pattern connects with person ‘A’ (in column 3) via the concept of a ‘person’ (in column 4). Connections amongst 3, 4, 5, 6, 7, and 8 mean that person ‘A’ is a witness to the murder.
4	The concept of ‘person’ (see also columns 10 and 14).	
5	An SP-pattern for an ‘event’.	
6	‘What’ the ‘event’ is about.	Connections amongst 5, 6, and 11 mean that the ‘event’ is the ‘murder’ identified in column 11.
7	The concept of ‘evidence’ (for the ‘event’).	Person ‘C’ is the ‘victim’ of the same ‘murder’ because of connections amongst 9, 10, 11, and 12.
8	The concept of a ‘witness’ (as a source of ‘evidence’).	
9	A specific ‘person’ (‘pn3’) with a ‘name’ (‘C’).	
10	The concept of a ‘person’ (see also column 4 and 14).	
11	The concept of a ‘murder’ with an ‘agent’ and a ‘victim’.	Person ‘B’ is the ‘agent’ of the ‘murder’ (ie he or she is the murderer) because of connections amongst 13, 14, 15. and 11.
12	The concept of a ‘victim’ (of the ‘murder’).	
13	A specific ‘person’ (‘pn2’) with a ‘name’ (‘B’).	
14	The concept of a ‘person’ (see also column 4 and 10).	
15	The concept of an ‘agent’ (for the ‘murder’).	
16	An SP-pattern showing that if the ‘witness’ had a good view of the ‘murder’, the ‘evidence’ is ‘strong’.	

Table 3: A table showing the Old SP-patterns in the SP-multiple-alignment in Figure 7 and what they mean. *Key:* ‘Column’ is short for ‘the SP-pattern in the given column’; ‘What’ is short for ‘what that SP-pattern represents’; the ‘Connections’ column shows connections amongst columns in the figure, and what the connections mean.

5. Also, we can see that the SP-pattern in column 16 is about ‘evidence’ for the ‘event’, not merely because it begins and ends with ‘evidence ... #evidence’ but because those two SP-symbols are aligned with the same two SP-symbols that define the slot for ‘evidence’ within the SP-pattern for ‘event’ in column 5.

5.2 The evidence is weak that B actually did murder C if the murder was committed in the dark

As readers will see, the SP-multiple-alignment in Figure 8 is very similar to the one in Figure 7 (Section 5.1). As with that figure, this one may be magnified and remain sharp at any convenient size, and this may be done with the copy of the figure in the file ‘spsrk2_figures.pdf’ so that everything else in the main paper need not be magnified.

With this figure, the ‘illumination’ of the ‘murder’ is ‘dark’, as can be seen from the SP-symbols ‘event murder ... illumination dark #illumination ... #murder ... #event’ in the New SP-pattern in column 0.

In the SP-multiple-alignment, the three SP-symbols ‘illumination dark #illumination’ in the New SP-pattern are aligned with the same three SP-symbols in the SP-pattern ‘sees ss1 illumination dark #illumination vision poor #vision #sees’ in column 1, leading to the inference that, in witnessing the ‘murder’, the ‘vision’ of ‘A’ was ‘poor’.

Then, because the three SP-symbols ‘vision poor #vision’ in column 1 are aligned with the same three SP-symbols in the SP-pattern ‘evidence enc2 witness vision poor #vision #witness strength weak #strength #evidence’ in column 16, we may infer from the following three symbols in that column, ‘strength weak #strength’, that the ‘strength’ of the ‘evidence’ is ‘weak’.

5.3 Discussion

The analyses that have been presented in Sections 5.1 and 5.2 show how inferences that the ‘evidence’ is ‘strong’ or ‘weak’ may be made with the SP System. But of course, with any with any real-life crime investigation, there would be many more details, and SP-multiple-alignments would be even bigger than the ones that have been shown.

In the long run, in line with Murray Shanahan’s conclusions [46, p. 142] noted in Sections 8 and 3.4.1, the best way to ensure that realistic levels of detail are recorded more accurately than can be done by hand coding, is to develop unsupervised learning in the SP system to the point where it may assimilate the diverse aspects of the world as they exist.

6 How the horse’s head scene in *The Godfather* may be interpreted by the SP System

This Section describes some of the CSRK complexity of the horse’s head scene in the *The Godfather* film, discussed in [13, p. 93], and describes how the SP System may model some of that complexity.

In summary, the relevant part of the plot is this:

“Johnny Fontane, a famous singer and godson to Vito [Corleone—the Godfather], seeks Vito’s help in securing a movie role; Vito dispatches his consigliere, Tom Hagen, to Los Angeles to talk the obnoxious studio head, Jack Woltz, into giving Johnny the part. Woltz refuses until he wakes up in bed with the severed head of his prized stallion.” (Adapted from “The Godfather”, *Wikipedia*, bit.ly/2c5YZAy, retrieved 2019-06-04.)

Instead of trying to understand the example from the perspective of a cinema audience, the analysis here will focus on how Jack Woltz might interpret the unpleasant experience of finding a horse’s head in his bed.

Although recognition and inference are intimately related (Appendix B.4), it seems that those two phases may usefully be distinguished in Woltz’s thinking:

- *Phase 1: Recognition.* The ‘recognition phase’, to be discussed in Section 6.1, may be seen to comprise three elements:
 - In order to make sense of the event, the first step is that Woltz must recognise the horse’s head as what it is. This may seem too easy and simple to deserve comment but that should not disguise the existence of this first step or its complexity.
 - The next step, which may again seem too simple to deserve comment, is that Woltz would make the very obvious inference that the horse’s head had been part of a horse.
 - Woltz would also recognise that the horse was his prized stallion which, we shall suppose, was called “Lightning Force”. We shall suppose also that a white flash on the horse’s forehead is distinctive for the stallion, although indirect inferences might also lead to the same identification.
- *Phase 2: Inference.* Why should the head of Lightning Force have appeared in Woltz’s bed? Some possibilities, to be discussed in Section 6.2, are summarised here:

- It could have been some kind of accident, although it is much more likely that it was the deliberate act by some person.
- Assuming that it was a deliberate act, what was the motivation? Here, Woltz’s knowledge of the Mafia would kick in: killing things is something that the Mafia do as a warning or means of persuading people to do what they want. The person to be persuaded must have an emotional attachment to the person or animal that is killed.⁷
- Woltz also knows that Tom Hagen is a member of the Mafia and that Hagen wants Woltz to give Johnny Fontane a part in a movie. From that knowledge and his knowledge of how the Mafia operate, Woltz can make connections with the killing of Lightning Force.

In principle, there could be one SP-multiple-alignment for both the recognition and the inference phases but this would have been too big to show on one page. So it has been convenient to split the analysis into two SP-multiple-alignments corresponding to the posited two phases in Woltz’s thinking.

6.1 Modelling the recognition phase via the creation of an SP-multiple-alignment

Figure 9 shows how Phase 1 in the scheme above (the recognition phase) may be modelled via the creation of an SP-multiple-alignment created by the SP Computer Model.

In this example, the computer model has been supplied with a set of Old SP-patterns describing various aspects of horses, mammals, and of Lightning Force. It has also been supplied with New information describing some of the features of the severed head that Woltz saw.

That set of features includes one that indicates that the horse is dead. Of course, this is a simplification of the way in which physical signs would have shown that the severed head, and thus the whole horse, was dead.

In the SP-multiple-alignment shown in Figure 9, the New information, which describes what Woltz saw and which appears in column 0, makes connections with various parts of the Old SP-patterns in columns 1 to 5. The SP-multiple-alignment shows that the horse’s head, represented by the SP-pattern in column 4, has been recognised, that it connects with the ‘head’ part of an SP-pattern representing the structure of mammals (in column 2), that this SP-pattern connects with an SP-pattern representing horses (in column 3), and that this in turn connects with an

⁷This is a little different from the interpretation by DM: “... it is clear Tom Hagen is sending Jack Woltz a message—if I can decapitate your horse, I can decapitate you; cooperate, or else.” [13, p. 93] but, arguably, equally valid.


```

0           1           2           3           4           5
                                                    <lf>
                                                    lf1
                                                    Lightning
                                                    Force
                                                    <horse>
<horse> ----- <horse>
  h1
<mammal> ----- <mammal>
  m1
<head> ----- <head>
                                                    h2
                                                    <marks> ----- <marks>
white-flash ----- white-flash
                                                    </marks> ----- </marks>
long-snout ----- long-snout
large-teeth ----- large-teeth
                                                    </head> ----- </head>
<body> ----- <body>
  hindgut
  fermentation
</body> ----- </body>
<legs> ----- <legs>
  odd-toed
</legs> ----- </legs>
<vital-signs> -- <vital-signs>
  v2
dead ----- dead
</vital-signs> - </vital-signs>
</mammal> ----- </mammal>
</horse> ----- </horse>
                                                    </lf>
0           1           2           3           4           5

```

Figure 9: An SP-multiple-alignment, created by the SP Computer Model, for the recognition phase in the horse's head example, as discussed in the text.

SP-pattern representing Lightning Force (in column 5). As mentioned above, we shall assume that Woltz recognises his prized stallion by the distinctive white flash on its forehead, and it is this feature which brings the SP-pattern for Lightning Force into the SP-multiple-alignment (in column 5).

6.2 Modelling the inference phase via via the creation of an sp-multiple-alignment

Figure 10 shows an SP-multiple-alignment for Phase 2 in the scheme above (the inference phase), ignoring the possibility (the second point under Phase 2) that the horse's head in Woltz's bed was the result of some kind of accident.




Figure 10: An SP-multiple-alignment for inference in the horse's head example, as discussed in the text.

Probably the most important feature of the SP-multiple-alignment shown in Figure 10 is the SP-pattern shown in column 3 which describes a supposed feature of how the Mafiosi operate. This is, reading from the top, that if x loves z , a member of the Mafiosi may persuade x to do something (an ‘action’ in the SP-pattern) by killing z (the thing that x loves). This is, no doubt, a distortion and oversimplification of how the Mafiosi operate but it is perhaps good enough for present purposes.

Other features of this SP-multiple-alignment concept include:

- The SP-pattern for Tom Hagen (in column 7) connects with the SP-pattern for Mafiosi (in column 3) and thus inherits their modes of operation.
- The SP-pattern in column 13 shows that Jack Woltz (with the reference code ‘psn2’ in the SP-pattern for Jack Woltz in column 12) ‘loves’ Lightning Force (with the reference code ‘lf1’ in the SP-pattern for Lightning Force in column 10).
- That fact (that Jack Woltz loves Lightning Force) connects with “ x loves z ” in the SP-pattern in column 3.
- Reading from the top, the SP-pattern in column 8 records the fact that Tom Hagen (with the reference code ‘psn3’) is seeking to ‘persuade’ Jack Woltz (with the reference code ‘psn2’) to perform a particular ‘action’ (with the reference code ‘ac2’). That action is to “Give Johnny the part”.

The analysis of the horse’s head scene that has been presented in this section is certainly not the last word, but I believe it suggests a possible way forward. Ultimately, robust capabilities will be needed for the unsupervised learning of CSK in realistic settings so that CSR may operate with relatively large and well-structured bodies of knowledge.

7 Winograd schemas

In a sentence like *The city councilmen refused the demonstrators a permit because they feared violence*, most people have no difficulty in understanding that “they” refers to the city councilmen, whereas in a sentence like *The city councilmen refused the demonstrators a permit because they advocated revolution*, it is easy to see that “they” refers to the demonstrators. But for AI systems, demonstrating how people make these judgements can be challenging.

This pair of sentences, first presented by Terry Winograd [58, p. 33], and many other ‘Winograd schemas’ like it that have been described subsequently,⁸ are seen

⁸See bit.ly/2MPm64B.

as examples of CSR [26]. Because their interpretation can be so challenging for artificial systems, it has also been suggested [25] that they could be alternatives to the Turing test for artificial intelligence [54], with possible advantages over that test.

How the SP System may help in the interpretation of Winograd schemas is described in [70], so there is no need for any detail here. In brief, it is suggested that, with Winograd’s example sentences:

- Any mature AI system for unsupervised learning—including future versions of the SP System—that has had the opportunity to learn about “the world”, would know that city councilmen, like other politicians, are generally in favour of peace and that they abhor violence amongst the general population.
- With knowledge of that association, the SP Computer Model may determine the referent of “they” in *The city councilmen refused the demonstrators a permit because they feared violence.*
- In a similar way, knowledge of the fact that some demonstrators favour revolution, allows the SP Computer Model to determine what “they” refers to in the sentence *The city councilmen refused the demonstrators a permit because they advocated revolution.*

Other Winograd schemas may be disambiguated in a similar way.

8 Cracking an egg into a bowl

Reasoning about the process of cracking an egg into a bowl is another CSRK problem contributed by Ernest Davis to the ‘Commonsense Reasoning Problem Page’ (bit.ly/2qjdMBj) on the website about ‘Commonsense Reasoning’ (bit.ly/2CPMWbq). Davis describes the problem like this:

“Characterize the following: A cook is cracking a raw egg against a glass bowl. Properly performed, the impact of the egg against the edge of the bowl will crack the eggshell in half. Holding the egg over the bowl, the cook will then separate the two halves of the shell with his fingers, enlarging the crack, and the contents of the egg will fall gently into the bowl. The end result is that the entire contents of the egg will be in the bowl, with the yolk unbroken, and that the two halves of the shell are held in the cook’s fingers.” (retrieved, 2018-11-09).

By contrast with Sections 5 to 7, which describe examples of CSRK problems where the SP System can demonstrate some success, this is a problem where the

SP System as it is now would not work, but where planned developments in the SP System may help solve the problem.

Despite the apparent simplicity of cracking an egg into a bowl, three different attempts at formalising the process using VLSs [28, 35, 46] show that it is remarkably complicated.

In connection with the surprising complexity of this seemingly simple task, and with reference to the “Naive Physics Manifesto” papers by Pat Hayes [20, 21], Murray Shanahan makes two interesting points [46, p. 142]:

- “First, it no longer seems plausible that a useable body of common sense knowledge about the physical world can be coded by hand.”
- “Second, the idea that researchers can make significant progress on the problem from their armchairs, that is to say without the ‘sanity check’ of having to deploy their formalisations on a robot, looks ridiculous.”

Those two points seem to be both right and important: 1) with anything but the simplest kind of problem, trying to handcraft the necessary knowledge is both excessively time consuming and prone to many errors; 2) it is very important to test any proposed solution by running it on a computer. This can quickly reveal any shortcomings that may exist in an idea. Although it may be necessary to begin with simplified versions of any CSRK problem, it is important to progress as far as possible towards realistic versions of such problems.

Owing mainly to shortcomings in the SP Computer Model as it is now (Appendix A.7), the egg-cracking problem is well beyond what may now be tackled now by the SP System. In particular,

- Without the ability to model structures in two or three dimensions, and with their motion and speed, it would be difficult or impossible to model with any accuracy an egg, a bowl, the cook’s hands, and the dynamic aspects of cracking the egg into the bowl.
- Because of problems in identifying low-level features in images, and because of shortcomings in unsupervised learning in the SP System, it is not yet feasible to learn the process of cracking an egg into a bowl from a live demonstration or from a video of a cook performing that task.
- Because of the shortcomings in SP System’s modelling of quantitative concepts including time, it is likely to be difficult or impossible to model the dynamics of cracking an egg into a bowl.

Planned developments of the SP system, described in [41], may overcome these problems.

9 Going forward

This section considers briefly some aspects of the research with a bearing on its future development.

9.1 The SP Computer Model

Readers who would like to work with the SP Computer Model and examples like those described in this paper, may download the source code, Windows executable code, example files, and related files, via the “SP62.zip” link under the heading “SOURCE CODE” near the bottom of the web page www.cognitionresearch.org/sp.htm. SP62 is the version of the SP Computer Model which was used for the examples in this paper, a version that excludes unsupervised learning.

A more comprehensive but more experimental version of the model, which includes procedures for unsupervised learning that are still immature, may be downloaded via the “SP71.zip” link.

9.2 How the research advances the current state of understanding

The main thing learned from the relative successes of the SP System, described in Sections 4, 5, 6, and 7, has been that the descriptive and explanatory Power of the SP System has been increased without any offsetting decrease in its Simplicity, meaning an improved ratio of Simplicity to Power (Appendix A.1.1). This should mean increased confidence:

- In the SP System as a foundation for the development of human-like AI.
- More specifically, that the SP System may prove fruitful in the development of systems for commonsense reasoning and the representation of commonsense knowledge.

A possible objection to the points just made is that the SP Computer Model has proved to be a non-starter in modelling the apparently simple task of cracking an egg into a bowl (Section 8). But this failure of the SP Computer Model appears to be largely due to known shortcomings of the model (Section 3.5), and better results are to be expected when the model has been more fully developed.

9.3 Why the advance matters

New insights in the development of human-like AI, and more specifically in CSRK, should facilitate future research in these areas.

In terms of practicalities, there are many potential benefits from the SP System, summarised in Appendix A.8.2.

10 Conclusion

Commonsense reasoning (CSR) and commonsense knowledge (CSK) (together abbreviated as CSRK) are areas of study concerned with aspects of human reasoning and knowledge which seem trivial to adults—such as how to make a cup of tea or how to go shopping—but which have proved to be challenging for artificial systems.

The *SP System*—meaning the *SP Theory of Intelligence* and its realisation in the *SP Computer Model*—has several features that appear to be favourable for modelling aspects of CSRK. These include:

- The SP System is the product of a unique programme of research aiming to simplify and integrate observations and concepts across artificial intelligence, mainstream computing, mathematics, and human learning, perception, and cognition. The quest for simplification and integration across a wide area has been largely successful, yielding, *inter alia*, the powerful concept of *SP-multiple-alignment*.
- SP-multiple-alignment is largely responsible for the versatility of the SP System: in several aspects of human intelligence including several kinds of reasoning, and in the representation of diverse kinds of knowledge. Because these things all flow from one relatively simple construct—SP-multiple-alignment—there is clear potential for the seamless integration of diverse aspects of aspects of intelligence and the representation of diverse kinds of knowledge, in any combination.
- With some further development, it is likely that the SP Computer Model will have the generality of a universal Turing machine, but, unlike a universal Turing machine, with strengths and potential in AI.
- Information compression, which is central in the workings of the SP System, offers the potential for: generality in inference and in the calculation of probabilities; generality in the representation of knowledge; and for the potential for learning the kinds of ‘natural’ structures which are likely to be important in CSRK.

- Compared with a relative lack of integration amongst the four areas of success with CSRK that have been discussed by DM [13, pp. 94–97]—taxonomic reasoning; temporal reasoning; action, events, and change; and qualitative reasoning—the SP System has potential to provide simplification and integration across those areas and beyond.

But the SP System is work in progress and, as it stands now, it has shortcomings in modelling aspects of CSRK, which are mainly: the representation and processing of information in two or more dimensions; the recognition of low-level perceptual features in speech and visual images; unsupervised learning; and the representation and processing of numbers and magnitudes (Section 3.5).

The main sections of the paper include four describing some successes with the SP System in modelling aspects of CSRK:

- “*Strength of evidence*”. Section 5 is about a ‘strength of evidence’ problem described like this: “A says that he witnessed B murdering C. Infer that the evidence that B actually did murder C is stronger if the murder was well lit than if it was in the dark.” With the SP Computer Model, relevant facts may be presented to the system as a New SP-pattern with Old SP-patterns representing relevant aspects of the world. Then the best SP-multiple-alignment created by the system when the murder was well lit shows a two-step inference that the strength of the evidence is strong. And the best SP-multiple-alignment created by the system when the murder was in the dark shows in a similar way that the strength of the evidence is weak.
- “*The meaning of noun phrases*”. Section 4, about “the meaning of noun phrases”, describes how the SP Computer Model may contribute to a child’s learning of the meaning of phrases like ‘water bird’. Early guesses about what such a phrase may mean may be progressively refined by observing associations between surface structures and contexts in their use by adults. At later stages, ambiguities in interpretation may be resolved via frequencies of occurrence or disambiguating contexts, or both those things, as demonstrated via the SP Computer Model.
- *How the horse’s head scene in ‘The Godfather’ may be interpreted via the SP System*. Section 6 discusses the horse’s head scene in *The Godfather* film, described by DM as a particularly challenging example of a CSRK problem. From the perspective of the character Jack Woltz, who discovers the head of his prize stallion in his bed, CSRK may be divided into two stages: 1) Jack Woltz’s recognition of what is in his bed; and 2) his inferences about what the Mafia may mean by it. How each of these stages may be modelled

is described, with accompanying SP-multiple-alignments, one for each stage, from the SP Computer Model.

- *Winograd schemas*. Section 7 discusses briefly ‘Winograd schemas’, meaning pairs of sentences like *The city councilmen refused the demonstrators a permit because they feared violence*, and *The city councilmen refused the demonstrators a permit because they advocated revolution*, where a pronoun like “they” in each sentence has a different referent in the two cases. The SP Computer Model provides a means of modelling the semantics in each case, including the two different referents for the critical pronoun. A much fuller presentation of this work may be found in [70].

By contrast with the four areas of relative success just described, Section 8 describes the CSRK problem of cracking an egg into a bowl which, despite its apparent simplicity, is too difficult for the SP System as it is now. This is mainly because of shortcomings of the SP System as it is now, summarised in Section 3.5. But planned future developments promise more success.

Appendices

A Outline of the SP System

This Section provides an outline description of the SP System. More information, listed here in increasing levels of detail, may be found in [71], [62], and [61]. Other papers in the SP programme of research, including several about potential benefits and applications of the system, are detailed with download links near the top of www.cognitionresearch.org/sp.htm.

A.1 Origins of the SP System and some of the thinking behind it

Within this section, several subsections describe relevant aspects of how the SP System originated, and some of the thinking behind it.

A.1.1 Aiming for conceptual *Simplicity* and explanatory or descriptive *Power*

The SP Theory of Intelligence and its realisation in the SP Computer Model is the product of *a unique programme of research*, seeking to simplify and integrate observations and concepts across artificial intelligence, mainstream computing, mathematics, and human learning, perception, and cognition (HLPC).

That focus on simplification and integration means the goal of discovering or inventing a system that combines conceptual *Simplicity* with high levels of descriptive or explanatory *Power*,⁹ a goal which is itself a version of Ockham’s razor,

That in turn can mean overcoming the longstanding problem of fragmentation in AI research, 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.” [33, p. 417].

It can also mean overcoming the more general problem of fragmentation in computing research, described by John Kelly and Steve Hamm, both of IBM: “... there’s a strong tendency [for researchers] to view each sensory field in isolation as specialists focus only on a single sensory capability. Experts in each sense don’t read journals devoted to the others senses, and they don’t attend one another’s conferences. Even within IBM, our specialists in different sensing technologies don’t interact much.” [24, p. 74].

In this context, it is relevant to note that, since the SP System and the concept of SP-multiple-alignment have been developed to simplify and integrate observations and concepts across several fields, it will of course have points of resemblance to many other systems. But any such resemblance does not mean that the SP System is “nothing but X”, or “nothing but Y”, and should not be a distraction from the importance of simplification and integration in IT systems, and the relative success of the SP System in combining conceptual Simplicity with high levels of descriptive and explanatory Power.

A.1.2 The SP programme of research is based on earlier research on language learning

The SP programme of research has been inspired in part by an earlier programme of research developing computer models of language learning, summarised in [60]. A key idea in that earlier research was learning via the identification of recurrent ‘chunks’ of information [34], including the identification of chunks containing other chunks, leading to the creation of hierarchical (tree-structured) kinds of procedural knowledge.

⁹This is one of two reasons for the name ‘SP’. The second reason is given in a footnote to Appendix A.1.4.

A.1.3 Seeking a more general model than hierarchical chunking

With the new goal of the SP research—simplification and integration of observations and concepts across a broad canvass—hierarchical chunking like that in the research on language learning—would not do. The aim has been to discover or create a conceptual framework that would accommodate a wide variety of aspects of intelligence and a wide variety of kinds of knowledge, including both tree-structured and non-tree-structured kinds of knowledge, both procedural and static.

As outlined in Appendix A.8, this quest has been largely successful, with the discovery and development of the powerful concept of *SP-multiple-alignment*, borrowed and adapted from the concept of ‘multiple sequence alignment’ in bioinformatics. The concept of SP-multiple-alignment (described in Appendix A.3, below) is largely responsible for the versatility of the SP System (Section 3.1). It has the potential to be as significant for an understanding of intelligence in a broad sense as is DNA for biological sciences. It could prove to be the ‘double helix’ of intelligence.

A.1.4 Compression of information in HLPC

The SP research, like the earlier research on language learning, has been inspired in part by a body of research, pioneered by Fred Attneave [2], Horace Barlow [4, 5], and others, with a focus on the importance of information compression (IC) in HLPC.¹⁰ A review of relevant evidence may be found in [73].

A.1.5 The intimate relation between IC and concepts of inference and probability

Another significant strand of thinking in the development of the SP System is the intimate relation between IC and concepts of inference (Appendix B.3), and between IC and concepts of probability [47, 48, 49, 27]. Those two things are combined in the system so that, although the SP System is dedicated to IC, it has strengths the making of inferences and in the calculation of their probabilities ([62, Section 4.4], [61, Section 3.7]).

The emphasis on IC in the SP System accords with research in the tradition of Minimum Length Encoding (see, for example, [27]), with the qualification that most research relating to MLE assumes that the concept of a universal Turing

¹⁰A second reason for the name ‘SP’ (additional to that given in a footnote to Appendix A.1.1) is that the SP System, in its operation, is dedicated to IC, and IC may be seen to be a process of maximising the ‘Simplicity’ of a body of information, **I**, by extracting redundancy from **I**, whilst retaining as much as possible of its non-redundant descriptive ‘Power’.

machine provides the foundation for theorising, whereas the SP System is founded on concepts of ICMUP and SP-multiple-alignment.

A.1.6 The SP Computer Model

The SP Theory and the SP Computer Model have been developed together, with the computer model helping to reduce vagueness in the theory, and providing a means of testing the theory and demonstrating what it can do. Many seemingly plausible ideas have been discarded as a result of testing in a long succession of versions of the SP Computer Model over a period of about 17 years. The current SP Computer Model provides a relatively robust expression of the SP Theory, validated via its performance with a variety of kinds of data.

Source code and Windows executive code for the SP Computer Model may be downloaded from below the heading “SOURCE CODE” on www.cognitionresearch.org/sp.htm.

A.1.7 The SP Machine

It is envisaged that the SP Theory and the SP Computer Model will provide the basis for an industrial-strength *SP Machine*, as shown schematically in Figure 11.




Figure 11: Schematic representation of the development and application of the SP machine. Reproduced from Figure 2 in [62], with permission.

Things to be done in the development of such an SP Machine are described in [41].

A.2 Organisation and Workings of the SP System

In broad terms, the SP System is a brain-like system that takes in *New* information through its senses and stores some or all of it in compressed form as *Old* information, as shown schematically in Figure 12.




Figure 12: Schematic representation of the SP System from an ‘input’ perspective. Reproduced, with permission, from Figure 1 in [62].

All kinds of knowledge are represented in the SP System with arrays of atomic *SP-symbols* called *SP-patterns*, in one or two dimensions. At present, the SP Computer Model works only with one-dimensional SP-patterns but it is envisaged that the model will be generalised to work with SP-patterns in two dimensions.

In view of evidence reviewed in [73], all kinds of information processing in the SP System are done via the compression of information—via a search for patterns that match each other, and via the merging or ‘unification’ of patterns (or parts of patterns) that are the same. As noted in Section 1.1, the expression ‘information compression via the matching and unification of patterns’ may be abbreviated as ‘ICMUP’.

More specifically, all kinds of processing in the SP System is done via IC via the concept of *SP-multiple-alignment*, described in Appendix A.3, next.

A.3 SP-multiple-alignment

Central in the SP Computer Model is the building of SP-multiple-alignments like the two shown in Figure 13. These may be seen as two alternative syntactic parsings of the ambiguous sentence *fruit flies like a banana*.¹¹

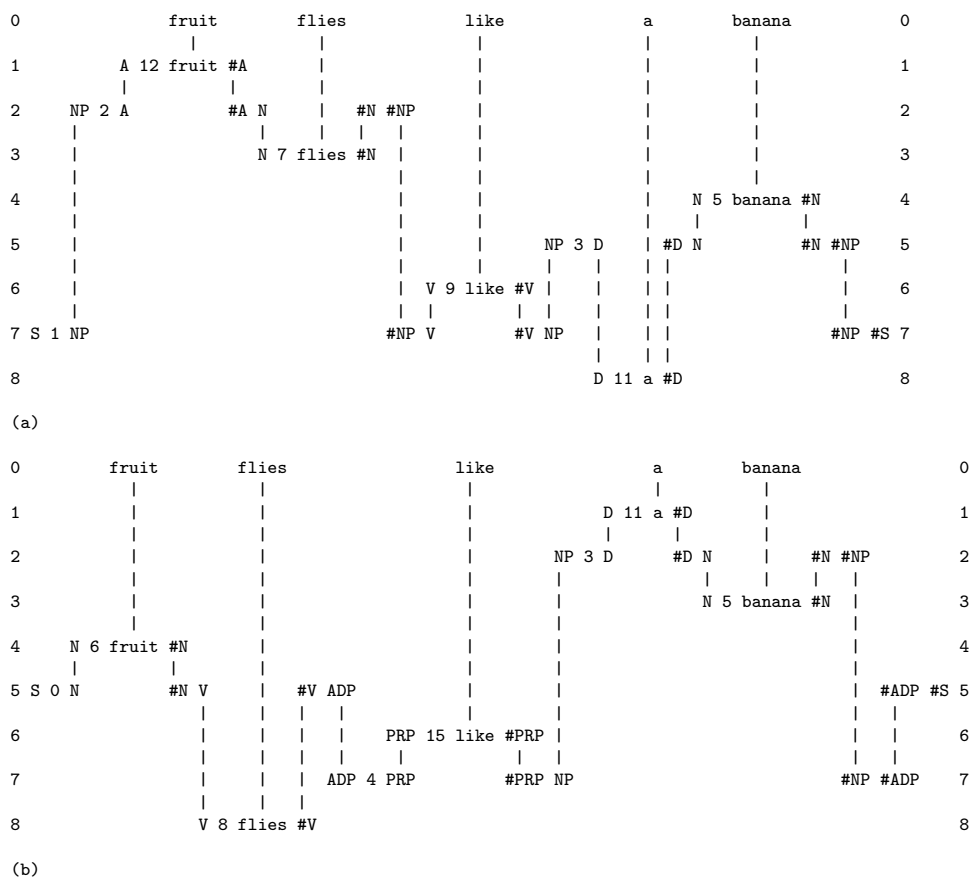


Figure 13: The two best SP-multiple-alignments created by the SP Computer Model showing two different parsings of the ambiguous sentence *Fruit flies like a banana* in terms of SP-patterns representing grammatical categories, including words. Here, SP-multiple-alignments are evaluated in terms of economical encoding of information as outlined in the text. Adapted from Figure 5.1 in [61], with permission.

In row 0 of each SP-multiple-alignment in the figure, there is a New SP-pattern, ‘fruit flies like a banana’, representing a sentence to be parsed, with five SP-symbols, each one corresponding to a word. By convention, all New SP-patterns

¹¹The sentence in each of these two SP-multiple-alignments is the second part of *Time flies like an arrow. Fruit flies like a banana*, attributed to Groucho Marx.

are shown in row 0 of each SP-multiple-alignment. Normally there is only one New SP-pattern in each SP-multiple-alignment but, as described in Appendix B.1, there can be more.

In each of rows 1 to 8 of each of the two SP-multiple-alignments in the figure, there is an Old SP-pattern which represents a grammatical category, such as the determiner, ‘a’, between the SP-symbols ‘D’ and ‘#D’, the noun, ‘banana’, between the SP-symbols ‘N’ and ‘#N’, noun phrases, each one between the SP-symbols ‘NP’ and ‘#NP’, and whole sentences, each one between the SP-symbols ‘S’ and ‘#S’. By convention, Old SP-patterns are always shown in rows numbered 1 and above, and there is always just one Old SP-pattern per row. The order of the Old SP-patterns across the rows is entirely arbitrary, with no special significance.

A.3.1 Display of SP-multiple-alignments

As we shall see, in Figures 2, 3, and elsewhere, SP-multiple-alignments may also be rotated by 90°, with SP-patterns arranged in columns instead of rows, and alignments between matching symbols shown in rows instead of columns. The choice between these two ways of displaying SP-multiple-alignments depends largely on what fits best on the page.

Another point about the display of SP-multiple-alignments is that they are often too big to display on one page. In some cases there are workarounds, such as splitting an SP-multiple-alignment into two or more parts, as has been done with the example in Section 6. But with Figures 7 and 8 in this paper, any such solution would make the SP-multiple-alignment unreasonably difficult to understand. Accordingly, another solution has been adopted.

Each of those two figures have been prepared in vector graphic format in a PDF file and shrunk to a size which allows them to be included in this paper, which is itself in a PDF file. Provided the paper is read in electronic form, the figures may be magnified and remain sharp at any convenient size. For the convenience of readers, those two figures have been provided in a separate file (‘spcsrk2_figures.pdf’) so that the figures may be magnified without magnifying all the text and the other figures in the paper.

A.3.2 Generality of the SP-multiple-alignment concept

A key point about the SP-multiple-alignment concept is that, although it accommodates hierarchical chunking in an efficient manner, as can be seen in Figure 13, it is well suited to several other forms of knowledge as well (Section 3.1.3). In this respect, it contrasts sharply with systems that can accommodate efficiently only one or two forms of knowledge, like the hierarchical chunking systems mention in Section A.1.3.

A.4 The building of SP-multiple-alignments

The process of building SP-multiple-alignments provides for the modelling of several different aspects of intelligence including: the analysis and production of natural language; pattern recognition at multiple levels of abstraction that is robust in the face of errors in data; best-match and semantic kinds of information retrieval; several kinds of reasoning; planning; and problem solving ([62, Sections 8 to 12], [61, Chapters 5 to 8]). Unsupervised learning in the SP System is somewhat different, as outlined in Appendix A.5.

In order to create SP-multiple-alignments like the two shown in Figure 13, the SP Computer Model must be supplied with the New SP-pattern representing the sentence to be parsed, and a set of Old SP-patterns representing a variety of grammatical structures. That set of Old SP-patterns would normally be much larger than the relatively few Old SP-patterns shown in the figure.

The overall aim is to create one or more SP-multiple-alignments where the New SP-pattern may be encoded economically in terms of the Old SP-patterns in the SP-multiple-alignment. How the *compression score* of each SP-multiple-alignment is calculated is described in [62, Section 4.1] and [61, Section 3.5]. Any SP-multiple-alignment with a high compression score may be described as ‘good’.

Normally, the building of good SP-multiple-alignments is much too complicated to be achieved by any kind of exhaustive search amongst the many possibilities. Instead, it is necessary to use heuristic search, building the SP-multiple-alignments in stages, and at each stage discarding all but the best partial structures. This approach cannot guarantee to find the best possible answer, but with enough computational resources, it can be good enough to find SP-multiple-alignments that are acceptably good.

A.5 Unsupervised Learning

In the SP System, learning is ‘unsupervised’, deriving structures from incoming sensory information without the need for any kind of ‘teacher’, or anything equivalent (*cf.* [19]). In this research, unsupervised learning is seen as a likely foundation for other kinds of learning, including ‘supervised’ learning (learning from examples of input-output pairs), ‘reinforcement’ learning (learning via rewards, and perhaps punishments), learning by being told, learning via imitation, and so on.

In the SP System, unsupervised learning incorporates the building of SP-multiple-alignments but there are other processes as well. In brief, the system creates Old SP-patterns from complete New SP-patterns and also from partial matches between New and Old SP-patterns.

When all the New SP-patterns have been processed like that, the system creates 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. As with the building of SP-multiple-alignments, the process is normally too complex to be done by exhaustive search so heuristic methods are needed. This means that the system builds SP-grammars incrementally and, at each stage, it discards all but the best SP-grammars.

The SP Computer Model 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 ([62, Appendix 5], [61, Chapter 9]). But there are two shortcomings in the system, outlined in Section A.7.

A.6 SP-Neural

Abstract concepts in the SP Theory of Intelligence map quite well into structures and processes represented in terms of neurons and their interconnections [67]. In developing this version of the SP Theory, as with the non-neural version of the theory, it is envisaged that the building and testing of computer models will play a central role, to maintain precision in the theory.

A.7 Unfinished business

In brief, the four main shortcomings in the SP System as it is now, are:¹²

- *Structures in two and three dimensions.* No attempt has yet been made to generalise the SP Computer Model to work with SP-patterns in two dimensions. This means that it would be difficult or impossible with the SP System to represent and process such things as photographs, diagrams, paintings, and the like. It is envisaged that this shortcoming will be remedied with further development of the SP System [41, Section 9].

Without an ability to represent and process structures in two dimensions, the SP System cannot easily represent and process information in three dimensions via the use of 2D SP-patterns as described in [63, Sections 6.1 and 6.2].

- *Recognition of low-level features in sound and vision.* Attention is needed for how the system may recognise low-level features in sound and visual images. Work is planned to remedy this deficiency [41, Section 10].
- *Unsupervised learning.* Although unsupervised learning in the SP System shows promise, there are two main deficiencies in the system as it is now:

¹²This is an updated version of [62, Section 3.3].

1) It cannot discover intermediate levels of structure such as phrases and clauses; and 2) It cannot discover dependencies in knowledge—such as number dependencies in the syntax of English or gender dependencies in the syntax of French—dependencies which are often ‘discontinuous’ because they may bridge other kinds of structure, and those intervening structures may be quite large. It is anticipated that these shortcomings in unsupervised learning in the SP System will be overcome with further development of the system [41, Section 12].

- *Magnitudes.* Although the SP System has things to say about the nature of mathematics ([72], [61, Chapter 10]), the SP Computer Model is not yet good at processing numbers, or magnitudes such as speed, time, length, area, volume, and the like. It is anticipated that deficiencies in areas like these may be remedied with further development of the SP System [41, Sections 9.3 and 14].

With regard to the frequency of occurrence of each Old SP-pattern, ‘neural’ versions of the theory (Appendix A.6) may use relatively imprecise biochemical or physiological measures. And in Appendix C, it is suggested that such imprecise measures may serve for recording such variables as speed, time, length, and so on.

A.8 Strengths and potential of the SP System

How the SP System may solve or help to solve several problems in AI research—described in Martin Ford’s book *Architects of Intelligence* [17]—is described in [74]. An earlier and slightly less up-to-date account of the distinctive features and advantages of the SP System is described in [68]. Section V in that paper describes 13 problems with deep learning in artificial neural networks and how, in the SP System, those problems may be overcome.

A.8.1 Combining conceptual Simplicity with high levels of descriptive and explanatory Power

In keeping with the goal of simplifying and integrating observations and concepts across a wide area (Appendix A.1.1), the main strength of the SP System is its combination of conceptual Simplicity with high levels of descriptive and explanatory Power. This is described in Section 3.1 with summaries of its versatility in aspects of intelligence (Section 3.1.1), including versatility in kinds of reasoning (Section 3.1.2), versatility in the representation of diverse forms of knowledge (Section 3.1.3), and its potential for the seamless integration of diverse aspects of intelligence with diverse forms of knowledge, in any combination (Section 3.1.4).

It appears that the SP System exhibits a more favourable combination of conceptual Simplicity with descriptive or explanatory Power, than any of the several attempts to develop ‘unified theories of cognition’ (see, for example, [39, 38]) or ‘artificial general intelligence’ (AGI, see, for example, [22]).

Although the AGI initiative is welcome, the difficulty of reaching agreement on a comprehensive framework for general, human-like AI is suggested by: 1) the following observation by the editors of the proceedings of the 2018 conference on AGI [22]: “Despite all the current enthusiasm in AI, the technologies involved still represent no more than advanced versions of classic statistics and machine learning.” [23, Locations 43–52]; by 2) the fact that none of the systems described in [22]) are plausible as paths to human-like AI; and 3) the editors of those proceedings seem to confirm the fragmentation in AI (noted in Appendix A.1.1) that the AGI initiative has aimed to solve: “Behind the scenes, however, many breakthroughs are happening on multiple fronts: in unsupervised language and grammar learning, deep-learning, generative adversarial methods, vision systems, reinforcement learning, transfer learning, probabilistic programming, blockchain integration, causal networks, and many more.” [23, Location 52].

A.8.2 Potential benefits and applications of the SP System

The SP system has several potential benefits and applications, described in several papers, details of which, with download links, may be found on www.cognitionresearch.org/sp.htm. These potential benefits and applications include: helping to solve nine problems with big data; helping to develop intelligence in autonomous robots; development of an intelligent database system; application of the SP System in medical diagnosis; the development of computer vision and research in natural vision; suggesting avenues for investigation in neuroscience, in commonsense reasoning, and more.

B Examples to illustrate versatility in intelligence of the SP System

As noted in Section 3.1.1, this appendix presents some examples showing how the SP Computer Model may demonstrate some of the aspects of intelligence mentioned in Section 3.1.1. They have been taken out of the main text to avoid disrupting the main presentation.

B.1 Pattern recognition

Figure 14 shows how, via the building of an SP-multiple-alignment, the SP computer system may model the recognition of an unknown plant.¹³

¹³As noted in Appendix A.3.1, SP-multiple-alignments like that shown in Figure 14, and others in the main text, compared with the SP-multiple-alignments in Figures 13 and 15, is rotated by 90°—with SP-patterns arranged in columns instead of rows and alignments between matching symbols shown in rows instead of columns. The choice between these two ways of displaying SP-multiple-alignments depends purely on what fits best on the page.

```

0          1          2          3          4          5          6
<species>
acris
<genus> -----<genus>
Ranunculus ----- Ranunculus
-----<family> -----<family>
----- Ranunculaceae ----- Ranunculaceae
-----<order> -----<order>
----- Ranunculales ----- Ranunculales
-----<class> -----<class>
----- Angiospermae ----- Angiospermae
-----<phylum> -----<phylum>
----- Plants ----- Plants
has_chlorophyll -----<feeding>
----- has_chlorophyll
----- photosynthesises
-----<feeding>
-----<structure> -----<structure>
-----<shoot>
-----<stem> -----<stem>
<stem> -----<stem>
----- hairy ----- hairy
-----</stem> -----</stem>
-----<leaves> -----<leaves>
----- compound
----- palmately_cut
-----</leaves> -----</leaves>
-----<flowers> -----<flowers>
-----<arrangement>
----- regular
----- all_parts_free
-----</arrangement>
-----<sepals> -----<sepals>
-----<sepal> -----<sepal>
----- not_reflexed
-----</sepal> -----</sepal>
-----<petals> -----<petals> -----<petals>
-----<number> -----<number>
----- five
-----</number> -----</number>
-----<colour> -----<colour>
----- yellow ----- yellow
-----</colour> -----</colour>
-----</petals> -----</petals> -----</petals>
-----<hermaphrodite>
-----<stamens> -----<stamens>
----- numerous ----- numerous
-----</stamens> -----</stamens>
-----<pistil>
----- ovary
----- style
----- stigma
-----</pistil>
-----</hermaphrodite>
-----</flowers> -----</flowers>
-----</shoot>
-----<root> -----<root>
-----</root>
-----</structure> -----</structure>
-----<habitat> -----<habitat> -----<habitat>
----- meadows ----- meadows
-----</habitat> -----</habitat>
-----<common_name> -----<common_name>
----- Meadow
----- Buttercup
-----</common_name> -----</common_name>
-----<food_value> -----<food_value>
----- poisonous
-----</food_value> -----</food_value>
-----</phylum> -----</phylum>
-----</class> -----</class>
-----</order> -----</order>
-----</family> -----</family>
-----</genus> -----</genus>
-----</species>
0          1          2          3          4          5          6

```

Figure 14: The best SP-multiple-alignment created by the SP Computer Model, with a small set of New SP-patterns (in column 0) that describe some features of an unknown plant, and a set of Old SP-patterns, including those shown in columns 1 to 6, that describe different categories of plant, with their parts and sub-parts, and other attributes. Reproduced with permission from Figure 16 in [62].

For the creation of the SP-multiple-alignment shown in Figure 14, the SP Computer Model was supplied with:

- Five New SP-patterns, shown in column 0 of the figure: the one-symbol SP-pattern ‘has_chlorophyll’, and four multi-symbol SP-patterns, ‘<stem> hairy </stem>’, ‘<petals> yellow </petals>’, ‘<stamens> numerous </stamens>’, and ‘<habitat> meadows </habitat>’, which describe the features of some unknown plant. These New SP-patterns may be supplied to the SP Computer Model in any order, not only the order shown in column 0 of the figure.
- A relatively large set of Old SP-patterns including those shown in columns 1 to 6. These describe the structures and attributes of different kinds of plant.

As with the example discussed in Appendix A.4, the SP System tries to find one or more SP-multiple-alignments where, in each one, the New SP-pattern or SP-patterns may be encoded economically in terms of the Old SP-patterns in the given SP-multiple-alignment. Broadly speaking, this means finding plenty of matches between New SP-patterns and Old SP-patterns, and a good number of matches between Old SP-patterns.

In the same way that the New SP-pattern or SP-patterns are always shown in row 0 of ‘horizontally’ arranged SP-multiple-alignments like those shown in Figure 13, with Old SP-patterns, one per row, in the remaining rows, the New SP-pattern or SP-patterns of ‘vertically’ arranged SP-multiple-alignments like that shown in Figure 14 are, by convention, always shown in column 0, and the Old SP-patterns are shown in the remaining columns, one SP-pattern per column. Otherwise, the order of SP-patterns across the columns has no significance.

As described in Appendix A.4, the aim in creating SP-multiple-alignments is to find ones which are ‘good’ in terms of the economical encoding of the New SP-pattern(s) in terms of the Old SP-patterns in that SP-multiple-alignment. In the process of searching for such SP-multiple-alignments, the SP System creates many SP-multiple-alignments or partial SP-multiple-alignments and discards most of them. Eventually, it is left with a few SP-multiple-alignments which are good, often as few as one or two.

The SP-multiple-alignment in Figure 14 is the best of those created by the SP Computer Model. It identifies the unknown plant with the SP-pattern shown in column 1: it is the species *acris* and it has the common name ‘Meadow Buttercup’.

B.2 Class-inclusion relations and part-whole relations

A feature of Figure 14 that has not so far been mentioned is that the unknown plant is not only identified as the species *acris* with the common name ‘Meadow

Buttercup’ (in column 1) but it is also recognised as belonging to the genus *Ranunculus* (in column 6). And the unknown plant is also recognised as belonging to the family Ranunculaceae (in column 5), which is in the order Ranunculales (in column 4), in the class Angiospermae (in column 3), and in the phylum Plants (in column 2).

In short, the SP System provides for the representation of class-inclusion hierarchies, and for their being an integral part of the recognition process, providing for recognition at multiple levels of abstraction, as mentioned near the beginning of Section 3.1.1. Notice how different New SP-symbols may be matched with Old SP-symbols at different levels in the class hierarchy: the feature ‘<stamens> numerous </stamens>’ is matched at the ‘family’ level (in column 5), the feature ‘has_chlorophyl’ is matched at the ‘Plants’ level (in column 2), the feature ‘<stem> hairy </stem>’ is matched at the ‘species’ level (in column 1), and so on.

Although it is not shown in Figure 14, a feature of the SP-multiple-alignment construct is that it can accommodate cross-classification as easily as simple hierarchies of classes.

Another important feature of Figure 14 is that it shows how the SP System not only supports the representation of class-inclusion hierarchies but it also supports the representation of part-whole hierarchies. For example, attributes which have a part-whole hierarchical structure include: ‘flowers’ in column 3; broken down into ‘sepals’, ‘petals’, ‘stamens’, and other attributes in column 5, with a further breakdown of ‘petals’ into the attributes ‘<number> </number>’ and ‘<colour> </colour>’. Actual values for the latter two attributes are, in this example, ‘<number> five </number>’ in column 6, and ‘<colour> yellow </colour>’ in column 1.

The figure also illustrates an important feature of the SP System: that there can be seamless integration of class-inclusion relations with part-whole relations, as discussed under ‘Categories and properties’ in Section 2.1.1 below.

With regard to the learning such structures, there is reason to believe that, when unsupervised learning in the SP System is more fully developed, class-inclusion relations and part-whole relations will fall within its scope. This is because IC via SP-multiple-alignments is central in how the SP System learns, and because both those kinds of relations can be represented within the SP-multiple-alignment framework, and because they provide powerful means of compressing information.

B.3 Inference via partial matching

The key to the making of inferences in the SP System is the partial matching of one structure with another which may be seen, for example, in ‘inheritance of

attributes’ (Appendix B.4). Thus in Figure 14, because the several matches with the SP-pattern in column 2 has ensured its place in the SP-multiple-alignment, and because the SP-symbol ‘**photosynthesises**’ has not been matched by anything, we may infer that the unknown plant photosynthesises. Likewise, because the SP-symbol ‘**poisonous**’ in the SP-pattern in column 5 is not matched with anything, we may infer that the unknown plant is poisonous.

In the SP System, inference via partial matching may be seen in: inheritance of attributes (as above); recognition in the face of errors of omission, commission, and substitution (Appendix B.5); and in probabilistic reasoning ([62, Section 10], [61, Chapter 7]).

In general, SP-multiple-alignments provide a foundation for a generalised and powerful version of “prediction by partial matching”, described in, for example, [53].

B.4 Inheritance of attributes

An important aspect of recognition in the SP System with class-inclusion hierarchies and with part-whole hierarchies is that both kinds of hierarchy provide a means of making a type of inference called “inheritance of attributes” that is bread-and-butter in everyday reasoning and everyday thinking.

To see inheritance of attributes in action, we may infer from the SP-multiple-alignment shown in Figure 14 that, in the plant that has been identified as *Ranunculus acris*, there are sepals that are not reflexed and leaves that are compound and palmately cut (the ‘**species**’ level in column 1). As noted in Appendix B.3, we may also infer that the plant nourishes itself via photosynthesis (the ‘**phylum**’ level in column 2), and that it is poisonous (the ‘**family**’ level in column 5). If there was more detail in the SP-patterns in the example, many more such inferences would be possible.

As we saw in Appendix B.3, inheritance of attributes in the SP-multiple-alignment framework, may be seen as ‘inference via partial matching’.

The intimate relation between pattern recognition and inference via inheritance of attributes illustrates a general truth about the SP System: that there is potential for the seamless integration of different aspects of intelligence, as discussed in Section 3.1.4, below.

B.5 Recognition in the face of errors of omission, commission, or substitution

An aspect of pattern recognition via the SP System that is not illustrated in Figure 14 is that, like people, the system has a robust ability to recognise patterns

despite errors of omission, commission, and substitution in the pattern or patterns that are to be recognised. An example of this aspect of pattern recognition, mentioned as “fuzzy” pattern recognition near the beginning of Section 3.1.1, is shown in the differences between Figures 15 (a) and (b).




Figure 15: (a) The best SP-multiple-alignment created by the SP Computer Model with a store of Old SP-patterns like those in rows 1 to 8 (representing grammatical structures, including words) and a New SP-pattern (representing a sentence to be parsed) shown in row 0. (b) As in (a) but with errors of omission, commission and substitution, and with same set of Old SP-patterns as before. Figure (a) reproduced from Figures 1 in [66], with permission.

As with inheritance of attributes (Section B.4), error corrections like those illustrated in Figure 15 (b) may be seen as examples of ‘inference via partial matching’ (Appendix B.3).

C How non-numerical magnitudes may be represented in the SP System

A remedy for the SP System’s current deficiency in the representation and processing of numbers (Appendix A.7) would make it relatively straightforward for the SP System—as an AI system—to represent magnitudes like speed, time, length, area, and volume.

But the SP system is not merely an AI system. It is intended, in addition, to model HLPC, and that would include young children before they have learned anything about numbers but appear to have a good approximate sense of magnitudes. As evidence for that last point, “... children as young as three years old notice and react negatively to an unfair distribution, particularly when they receive less than their partner.” [29, Abstract]).

In a similar way, it appears that anyone engaging in sport, or driving a car, assesses magnitudes such as distance and speed without the use of numbers. And a squirrel jumping from branch to branch of a tree must have some kind of mental measure of distance, without the use of numbers or arithmetic as they are taught in school.

How could these kinds of magnitude be assessed in the SP System in that kind of informal, approximate way, without the use of numbers? The suggestion here is that it may be possible to exploit a feature of the SP System that is integral to how it works, at least in the ‘neural’ version of the system, SP-Neural (Appendix A.6).

In the SP Computer Model as it is now, every Old SP-pattern has an associated measure of its frequency of occurrence which is expressed directly with an integer. But in SP-Neural or any biological version of the SP System, it is more plausible to suppose that each frequency measure would be expressed approximately as the concentration of some biological chemical associated with a given (neural) SP-pattern or, with each such entity, the strength some physiological variable.

Similar principles may apply to the representation of variables such as speed, time, length, and so on. Any such variable may be represented via the concentration or strength of some biochemical or physiological indicator.

References

- [1] S. Aditya, Y. Yang, C. Baral, C. Fermuller, and Y. Aloimonos. Visual commonsense for scene understanding using perception, semantic parsing and reasoning. In *Papers from the 2015 AAAI Spring Symposium: Logical Formalizations of Commonsense Reasoning*, pages 9–16, 2015.

- [2] F. Attneave. Some informational aspects of visual perception. *Psychological Review*, 61:183–193, 1954.
- [3] F. Baarder, I. Horrocks, and U. Sattler. Description logics. In F. van Harmelin, V. Lifschitz, and B. Porter, editors, *Handbook of Knowledge Representation*, pages 135–179. Elsevier, Amsterdam, first edition, 2008.
- [4] 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.
- [5] 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.
- [6] C. A. Becker. Semantic context and word frequency effects in visual word recognition. *Journal of Experimental Psychology: Human Perception and Performance*, 5(2):252–259, 1979.
- [7] G. M. Birtwistle, O-J Dahl, B. Myhrhaug, and K. Nygaard. *Simula Begin*. Studentlitteratur, Lund, 1973.
- [8] J. A. Blass and K. D. Forbus. Analogical chaining with natural language instruction for commonsense reasoning. In *Proceedings of the Thirty First AAAI Conference on Artificial Intelligence (AAAI-17)*, pages 4357–4363, 2017.
- [9] M. Brysbaert, P. Mandera, and E. Keuleers. The word frequency effect in word processing: an updated review. *Current Directions in Psychological Science*, 27(1):45–50, 2018.
- [10] E. Cambria, S. Poria, D. Hazarika, and K. Kwok. SenticNet 5: Discovering conceptual primitives for sentiment analysis by means of context embeddings. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI-18)*, pages 1795–1802, 2018.
- [11] E. Davis. *Representations of Commonsense Knowledge*. Morgan Kaufmann, San Mateo, California, 1990.
- [12] E. Davis. Logical formalizations of commonsense reasoning: a survey. *Journal of Artificial Intelligence Research*, 59:651–723, 2017.
- [13] E. Davis and G. Marcus. Commonsense reasoning and commonsense knowledge in artificial intelligence. *Communications of the ACM*, 58(9):92–103, 2015.

- [14] E. Davis and G. Marcus. The scope and limits of simulation in automated reasoning. *Artificial Intelligence*, 233:60–72, 2016.
- [15] A. Ferrein, C. Fritz, and G. Lakemeyer. Using Golog for deliberation and team coordination in robotic soccer. *Kuntzliche Intelligenz*, 19(1):24–30, 2005.
- [16] M. Fisher. Temporal representation and reasoning. In F. Van Harmelen, V. Lifschitz, and B. Porter, editors, *Handbook of Knowledge Representation*, pages 513–550. Elsevier, Amsterdam, 2008.
- [17] M. Ford. *Architects of Intelligence: the Truth About AI From the People Building It*. Packt Publishing, Birmingham, UK, Kindle edition, 2018.
- [18] A. Freitas, J. C. P. da Silva, E. Curry, and P. Buitelaar. A distributional semantics approach for selective reasoning on commonsense graph knowledge bases. In *Proceedings of the International Conference on Applications of Natural Language to Data Bases/Information Systems NLDB 2014: Natural Language Processing and Information Systems*, pages 21–32, 2014.
- [19] M. Gold. Language identification in the limit. *Information and Control*, 10:447–474, 1967.
- [20] P. J. Hayes. The naive physics manifesto. In D. Michie, editor, *Expert Systems in the Microelectronic Age*. Edinburgh University Press, Edinburgh, 1979.
- [21] P. J. Hayes. The second naive physics manifesto. In J. R. Hobbs and R. C. Moore, editors, *Formal Theories of the Commonsense World*, pages 1–36. Ablex, New York, 1985.
- [22] M. Iklé, A. Franz, R. Rzepka, and B. Goertzel, editors. *Artificial General Intelligence*, volume 10999 of *Lecture Notes in Computer Science*, Heidelberg, 2018. Springer.
- [23] M. Iklé, A. Franz, R. Rzepka, and B. Goertzel. Preface. In *Proceedings of the 11th International Conference, AGI 2018, Prague, Czech Republic, August 22-25, 2018* [22], pages Locations 43–72.
- [24] J. E. Kelly and S. Hamm. *Smart machines: IBM’s Watson and the era of cognitive computing*. Columbia University Press, New York, Kindle edition, 2013.
- [25] H. J. Levesque. The Winograd Schema Challenge. In *Proceedings of the Tenth International Symposium on Logical Formalizations of Commonsense Reasoning (Commonsense-2011)*, 2011. Part of the AAAI Spring Symposium Series at Stanford University, March 21-23, 2011.

- [26] H. J. Levesque, E. Davis, and L. Morgenstern. The Winograd Schema Challenge. In *Proceedings of the Thirteenth International Conference on Principles of Knowledge Representation and Reasoning*, pages 552–561, 2012.
- [27] M. Li and P. Vitányi. *An Introduction to Kolmogorov Complexity and Its Applications*. Springer, New York, 4th edition, 2019.
- [28] V. Lifschitz. Cracking an egg: an exercise in commonsense reasoning. In *Common Sense 1998, Logical Formalizations of Commonsense Reasoning: Proceedings of the AAAI Spring Symposium, 1998*, pages 51–56, 1998.
- [29] V. LoBue, T. Nishida, C. Chiong, J. S. DeLoache, and J. Haidt. When getting something good is bad: even three-year-olds react to inequality. *Social Development*, 20(1):154–170, 2011.
- [30] Y. Ma, H. Peng, and E. Cambria. Targeted aspect-based sentiment analysis via embedding commonsense knowledge into an attentive LSTM. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI-18)*, pages 5876–5883, 2018.
- [31] G. Marcus and E. Davis. How robust are probabilistic models of higher-level cognition? *Psychological Science*, 24(12):2351–2360, 2013.
- [32] N. Maslan1, M. Roemmele, and A. S. Gordon. One hundred challenge problems for logical formalizations of commonsense psychology. In *Papers from the 2015 AAAI Spring Symposium on Logical Formalizations of Commonsense Reasoning*, pages 107–113, 2015.
- [33] 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. ISBN: 1-56881-205-1.
- [34] G. A. Miller. The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological Review*, 63:81–97, 1956.
- [35] L. Morgenstern. Mid-sized axiomatizations of commonsense problems: a case study in egg cracking. *Studia Logica*, 67(3):333–384, 2001.
- [36] L. Morgenstern, E. Davis, and C. L. Ortiz. Planning, executing, and evaluating the Winograd Schema Challenge. *AI Magazine*, 37(1):50–54, 2017.
- [37] E. T. Mueller. *Commonsense Reasoning: An Event Calculus Based Approach*. Morgan Kaufmann, Waltham, MA, 02451, USA, second, Kindle edition, 2006.

- [38] A. Newell, editor. *Unified Theories of Cognition*. Harvard University Press, Cambridge, Mass., 1990.
- [39] A. Newell. Précis of *Unified Theories of Cognition*. *Behavioural and Brain Sciences*, 15(3):425–437, 1992.
- [40] D. Norris. Word recognition: Context effects without priming. *Cognition*, 22(2):93–136, 1986.
- [41] V. Palade and J. G. Wolff. A roadmap for the development of the ‘SP Machine’ for artificial intelligence. *The Computer Journal*, 62, 2019. arXiv:1707.00614, bit.ly/2tWb88M.
- [42] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Francisco, revised second printing edition, 1997.
- [43] S. L. Reed and D. B. Lenat. Mapping ontologies into Cyc. Technical report, AAI, 2002. From AAI Technical Report WS-02-11.
- [44] R. Reiter. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. MIT Press, Cambridge, MA, 2001.
- [45] M. Sap, R. Le Bras, E. Allaway, C. Bhagavatulay, N. Lourie, H. Rashkin, B. Roof, N. A. Smith, and Y. Choi. ATOMIC: an atlas of machine commonsense for *if-then* reasoning. Technical report, Allen Institute for Artificial Intelligence, 2019. arXiv:1811.00146v3 [cs.CL].
- [46] M. Shanahan. An attempt to formalise a non-trivial benchmark problem in common sense reasoning. *Artificial Intelligence*, 153(1–2):141–165, 2004.
- [47] C. E. Shannon and W. Weaver. *The Mathematical Theory of Communication*. University of Illinois Press, Urbana, 1949.
- [48] R. J. Solomonoff. A formal theory of inductive inference. Parts I and II. *Information and Control*, 7:1–22 and 224–254, 1964.
- [49] R. J. Solomonoff. The discovery of algorithmic probability. *Journal of Computer and System Sciences*, 55(1):73–88, 1997.
- [50] R. Speer, J. Chin, and C. Havasi. ConceptNet 5.5: An open multilingual graph of general knowledge. In *Proceedings of 31st AAAI Conference on Artificial Intelligence (AAAI-17)*, pages 4444–4451, 2017.

- [51] K. E. Stanovich and R. F. West. The effect of sentence context on ongoing word recognition: tests of a two-process theory. *Journal of Experimental Psychology: Human Perception and Performance*, 7(3):658–672, 1981.
- [52] M. Surdeanu. Overview of the TAC2013 knowledge base population evaluation: English slot filling and temporal slot filling. In *Proceedings of the 6th Text Analysis Conference (2013)*, 2013.
- [53] W. J. Teahan and K. M. Alhawiti. Preprocessing for PPM: compressing UTF-8 encoded natural language text. *International Journal of Computer Science & Information Technology*, 7(2):41–51, 2015.
- [54] A. M. Turing. Intelligent machinery. In B. J. Copeland, editor, *The Essential Turing: Seminal Writings in Computing, Logic, Philosophy, Artificial Intelligence, and Artificial Life: Plus The Secrets of Enigma*, pages 395–432. Oxford University Press, Oxford, 1948.
- [55] A. M. Turing. Computing machinery and intelligence. *Mind*, 59:433–460, 1950.
- [56] F. van Harmelin, V. Lifschitz, and B. Porter. *Handbook of Knowledge Representation*. Elsevier, Amsterdam, 2008.
- [57] C. S. Webster. Alan turing’s unorganized machines and artificial neural networks: his remarkable early work and future possibilities. *Evolutionary Intelligence*, 5:35–43, 2012.
- [58] T. Winograd. Understanding natural language. *Cognitive Psychology*, 3(1):1–191, 1972.
- [59] M. Witbrock, K. Pittman, J. Moszkowicz, A. Beck, D. Schneider, and D. Lenat. Cyc and the big C: Reading that produces and uses hypotheses about complex molecular biology mechanisms. In *Scholarly Big Data: AI Perspectives, Challenges, and Ideas: Papers from the 2015 AAAI Workshop*, pages 27–38. 2015.
- [60] 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.
- [61] 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.

- [62] J. G. Wolff. The SP Theory of Intelligence: an overview. *Information*, 4(3):283–341, 2013. arXiv:1306.3888 [cs.AI], bit.ly/1NOMJ6l.
- [63] 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.
- [64] 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.
- [65] 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, 2016, Chapter 6, pp. 143–170.
- [66] 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.
- [67] 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.
- [68] 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.
- [69] J. G. Wolff. Software engineering and the SP Theory of Intelligence. Technical report, CognitionResearch.org, 2017. Submitted for publication. arXiv:1708.06665 [cs.SE], bit.ly/2w99Wzq.
- [70] J. G. Wolff. Interpreting Winograd Schemas via the SP Theory of Intelligence and its realisation in the SP Computer Model. Technical report, CognitionResearch.org, 2018. Submitted for publication. bit.ly/2ME8DOA.
- [71] J. G. Wolff. Introduction to the SP Theory of Intelligence. Technical report, CognitionResearch.org, 2018. arXiv:1802.09924, bit.ly/2ELq0Jq.
- [72] J. G. Wolff. Mathematics as information compression via the matching and unification of patterns. Technical report, CognitionResearch.org, 2018. Submitted for publication. arXiv:1808.07004 [cs.AI], bit.ly/2LWbjtK.
- [73] J. G. Wolff. Information compression as a unifying principle in human learning, perception, and cognition. *Complexity*, 2019:38 pages, February 2019. Article ID 1879746. viXra:1707.0161v3, hal-01624595 v2.

- [74] J. G. Wolff. Unsolved problems in AI, described in the book ‘Architects of Intelligence’ by Martin Ford, and how they may be solved via the SP System. Technical report, CognitionResearch.org, 2019. Submitted for publication. viXra:1902.0220, hal.archives-ouvertes.fr/hal-02061171.
- [75] S. Zhang and P. Stone. CORPP: Commonsense reasoning and probabilistic planning, as applied to dialog with a mobile robot. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 1394–1400, 2015.