

A proposed solution to problems in learning the knowledge needed by self-driving vehicles

J Gerard Wolff*

September 13, 2021

Abstract

Three problems in learning knowledge for self-driving vehicles are: how a finite sample of information about driving, \mathbf{N} , can yield an ability to deal with the infinity of possible driving situations; the problem of generalising from \mathbf{N} without over- or under-generalisation; and how to weed out errors in \mathbf{N} . A theory developed with computer models to explain a child's learning of his or her first language, now incorporated in the *SP System*, suggests: compress \mathbf{N} as much as possible by a process that creates a grammar, \mathbf{G} , and an encoding of \mathbf{N} in terms of \mathbf{G} called \mathbf{E} . Then discard \mathbf{E} which contains all or most of the errors in \mathbf{N} , and retain \mathbf{G} which solves the first two problems.

This paper is about how the knowledge about driving that is needed by self-driving vehicles (SDVs) may be developed, taking account of the following problems: the finite body of information, \mathbf{N} , which is the basis for learning must give rise to a capability for dealing with the infinite range of possible situations that may be encountered in driving; the problem of generalising from \mathbf{N} without over-generalisation (aka under-fitting), and without under-generalisation (aka over-fitting); and the problem of learning knowledge that is largely or completely free of any errors or 'dirty data' that may be in \mathbf{N} .

The proposed solution to these three problems originated in research developing computer models of the learning of a first language or languages by children [13], and is now a part of the *SP System* (SPS), meaning the *SP Theory of Intelligence* and its realisation in the *SP Computer Model* ([14]—a book, [15]—a paper, [17, Section 3]—one section of a paper). The proposed solution will be referred to as the 'SPS solution'.

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

A point to bear in mind is that the SPS is work in progress, with some shortcomings as it stands now. This paper assumes that these weaknesses have been overcome, as described in “A roadmap for the development of the SP Machine for artificial intelligence” in [11].

To be clear, this paper is about theory, it is not reporting experiments with sections like ‘Method’, ‘Results’, and ‘Discussion’. The paper is about the potential of a theory developed with computer models of the learning of a first language, and later incorporated in the SP System, may be applied in the development of driving-related knowledge for SDVs.

1 Learning a first language

The language learning problem may be conceived in abstract terms as shown by the Venn diagram in Figure 1.

In the figure, the smallest envelope represents the large but nevertheless finite set of ‘utterances’ from which a young child learns—let us call her ‘Jane’. Here, an ‘utterance’ is any spoken sound including gurgles and burps as well as ordinary words, phrases, and sentences.

The middle-sized envelope represents the infinite set of utterances in a mature knowledge of **L**.

We know that a mature knowledge of a language like English encompasses infinitely many utterances because it provides for recursive structures like “This is the horse and the hound and the horn, That belonged to the farmer sowing his corn, ... That killed the rat that ate the malt, That lay in the house that Jack built.” and that, in principle, there is no limit to the depth of the recursion. In a similar way, since there is in principle no limit to the lengths of grammatical sentences, and since there are normally many shorter versions of long sentence, the number of potential sentences that may be created in any natural language is infinite. We shall return to this point in connection with learning to drive.

The largest envelope in the figure represents the infinite set of all possible utterances, including all the utterances in **L**, all utterances in the rest of the world’s languages, both living and dead, and utterances that may not appear in any language.

Notice that the smallest envelope, representing the finite set of utterances from which Jane learns, overlaps the largest envelope, meaning that part of the information from which she learns is not in **L**. It has been labelled ‘dirty data’ because, with respect to **L**, those data may be seen as corrupted information.

A point to emphasise here is that, in learning a first language, dirty data means haphazard errors like stumbling over the pronunciation of a word, getting one’s words in the wrong order, stopping in the middle of a phrase, clause, sentence, or

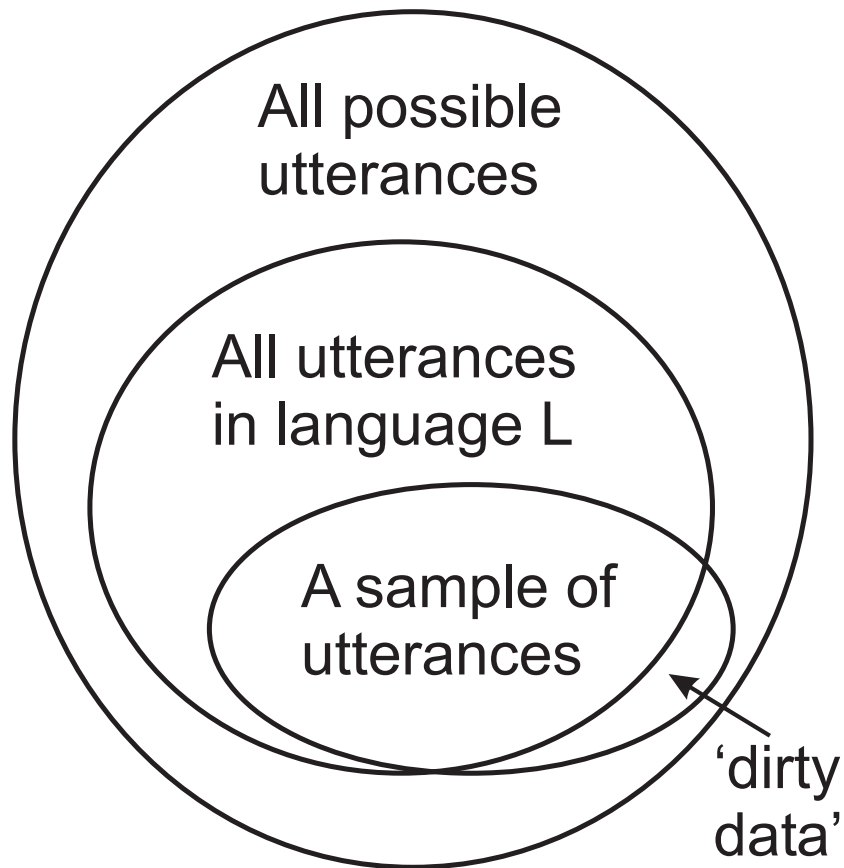


Figure 1: Categories of utterances involved in the learning of a first language, L . In ascending order of size, they are: the finite sample of utterances from which a child learns; the (infinite) set of utterances in a mature knowledge of L ; and the (infinite) set of all possible utterances. Adapted from Figure 7.1 in [13], with permission.

other grammatical solecisms. It does not mean hate speech, phishing, trolling, or the like.

1.1 The unsupervised learning of a first language

There is good evidence that Jane may learn (a mature knowledge of) her first language, **L**, in a manner that may be entirely ‘unsupervised’, meaning that there is no need for ‘teaching’ in a conventional sense, or rewards or punishments, or the labelling of data, or the correction of errors that Jane may make. While Jane may benefit from aids like that, the weight of evidence is that, apart from the opportunity to hear **L** spoken, she needs no assistance for successful learning of her first language.

Probably the best evidence for a child’s ability to learn a first language without correction of errors and the like are children like Christy Brown who achieved a good knowledge of English mainly by listening to people around him, but because of his cerebral palsy his speech was largely unintelligible so there was little opportunity to correct his language.

When he gained control over his left foot, his mother taught him to write and spell, so that eventually he could demonstrate his excellent knowledge of English by typing his autobiography (*My Left Foot* [7]), and other books. It is fairly clear that Brown’s mother concentrated on the writing and spelling of English, and that learning the language itself was his achievement.

This paper concentrates on unsupervised learning because:

- As just described, there is evidence that learning a first language may be achieved without any kind of teaching, the correction of errors, or the provision of labelled examples, and so on.
- There is potential in unsupervised learning for a relatively straightforward and effective means of building the knowledge needed by SDVs.

2 Three problems in learning a first language

In view of the evidence that the learning of a first language is normally unsupervised and does not depend on the correction of errors or the like, there seem to be three main things to be explained. They are essentially the same problems as were mentioned in the introduction in connection with learning how to drive:

1. *INF*. How can Jane translate **N**, the large but finite sample of language **L** which is the basis for her learning, into a body of knowledge, **G**, which

provides the means of interpreting any of the infinite range of sentences in the language \mathbf{L} ?

2. *GEN*. How can Jane generalise from \mathbf{N} without over-generalisation or under-generalisation?
3. *DD*. In the knowledge that Jane learns, how can she avoid the corrupting effect of any errors or dirty data that may be in \mathbf{N} ?

Answers to these questions and the proposed SPS solution for both the learning of a first language and for learning how to drive is given later.

3 Three problems in learning to drive

At first sight, it may seem that the best way to develop driving skills in an SDV is to teach them in much the same way that a driving instructor teaches a human pupil.

This would make sense if the SDV had all the knowledge and intelligence of a person but this is not yet the case. In particular, without programming, SDVs (i.e. computers) do not have human abilities to model infinity via finite information, to generalise without over- or under-generalisation, and to deal with dirty data.

As indicated at the beginning, the proposed SPS solution is essentially the same as is employed by children in the unsupervised learning of a first language and may yield the same benefits in terms of the three problems mentioned earlier.

An advantage here of using unsupervised learning is that it has potential to create detailed driving-related knowledge automatically or nearly so, with potential to eliminate the many errors that can creep in via ordinary programming.

In summary, this is how things may work in the automation of learning to drive:

- *INF*. With the learning of knowledge needed by SDVs, \mathbf{L} would be the entire ‘language’ comprising initiatives by the artificial driver (instructed by human passengers), incoming information as they drive along in a wide variety of situations, and responses by the SDV to changes in that incoming information. \mathbf{N} would be a large but finite sample of \mathbf{L} .

With the learning of knowledge needed by SDVs, \mathbf{N} would be compiled from long drives with good human drivers in a wide variety of situations.

The *INF* problem is how, from \mathbf{N} , to create an ability to model an infinitely large \mathbf{L} .

- *GEN*. As with learning a first language, the GEN problem is how to generalise from **N** without over- or under-generalisations.
- *DD*. As with learning a first language, the DD problem is how the learning process can filter out corrupted or dirty data in **N**. Here, dirty data will mean errors of various kinds that may creep into **N**, including what one hopes will be rare errors made by any of the human drivers.

Notice that the many examples of bad driving that will be encountered on the roads will not count as dirty data. They are merely examples of driving situations that will need appropriate responses.

An important point about the SPS is that, with any given body of stored knowledge, precisely the same software may serve both the analysis of incoming information and the production of outgoing information. This is explained and discussed in [14, Sections 3.5 and 3.8] and [15, Section 4.5].

In an SDV, this feature of the SPS is potentially useful because: any one body of learned knowledge may serve both the analysis of incoming information and the creation of responses; in the learning process, there is no need to differentiate between incoming information and responses by the SDV, so that both of them may be built into a single body of knowledge.

4 The state of play

In the light of what has been said so far, what is the state of SDV research on INF, GEN, and DD? A search of recent work in this area shows that, even in papers where one would expect the INF, GEN, and DD issues to be discussed (eg in connection with safety), those concepts are largely invisible. For example, there is no mention of them in:

- “Self-driving cars: a survey” [3];
- A special issue of *IEEE Computer* on “Self-driving cars” introduced in [8];
- “An improved safety algorithm for artificial intelligence enabled processors in self driving cars” [9];
- “The key technology toward the self-driving car” [19].
- “Query-efficient imitation learning for end-to-end autonomous driving” [18].

Some other papers about SDVs do consider generalisation, but the treatment of the subject is disappointing:

- “On a formal model of safe and scalable self-driving cars” [12]. Although this paper focuses on safety, there are only a few brief mentions of generalisation, beginning with:

“The challenge with scenario-based approaches has to do with the notion of ‘generalisation’, in the sense of the underlying assumption that if the AV [autonomous vehicle] passes the scenarios successfully then it is likely to pass other similar scenarios as well. The danger, just as in machine learning, is ‘overfitting’ the system to pass the test. Even if extra care is taken not to overfit, the arguments of generalisation are weak at best.”

Collectively, the remarks are a long way from a satisfactory analysis of generalisation, and there is no mention of INF or DD.

- “Rethinking self-driving: multi-task knowledge for better generalisation and accident explanation ability” [20]. With regard to generalisation in the development of SDVs, this paper notes that in studies by other researchers there has been “Poor generalization ability of unobserved driving environment given limited diversity of training scenerios.” (p. 1). The proposed solution is “... a new driving system for better generalization ... by enabling it to do simpler driving-related perception task before generating commands for difficult driving task.”

The main problem here seems to be that the proposed solution is entirely *ad hoc* with no attempt to relate it to any other conceptual framework or theory.

- “Synthetic examples improve generalization for rare classes” [6]. This relatively general paper about generalization includes discussion of the issue in relation to SDVs: “The ability to detect and classify rare occurrences in images has important applications—[in] for example, ... detecting infrequent traffic scenarios that pose a danger to self-driving cars.” [6, Abstract]. Using simulated data, the authors conclude, as one might expect, that: “... as the amount of simulated data is increased, accuracy on the target class improves.” but that, with larger amounts of data, there can be distortions in the identification of other classes [6, Abstract]. Also “the variation of simulated data generated is very important, and maximum variation provides maximum performance gain.” [6, p. 870].

These results with simulated data are much as one might expect but seem not to provide much help in, for example, dealing with traffic scenarios that occur only rarely. Nevertheless, the study serves a useful purpose in drawing attention to the issue of rare scenarios. Of course, this matters

most if those rare scenarios are dangerous. This issue is discussed briefly in Section 6.5, below.

- “Deep fully convolutional networks with random data augmentation for enhanced generalization in road detection” [10]. In this paper: “It is demonstrated that significant generalization gains in the learning process are attained by randomly generating augmented training data using several geometric transformations and pixelwise changes, such as affine and perspective transformations, mirroring, image cropping, distortions, blur, noise, and color changes.” [10, Abstract]. It appears that the ‘generalization’ that has been achieved in this research is not generalisation in the sense of going beyond the information that has been given, but is merely the addition of some variety to the data from which the model learns.

This sampling of recent studies suggests that the important subjects of INF and GEN is not well understood, and the DD problem has not been considered at all.

In none of the papers above is there any mention of the possibility that, in SDVs, one body of knowledge may serve both the analysis of incoming information and the production of responses.

5 Solving the INF-GEN-DD problems via information compression

This main section describes how the problems that have been described may be solved via IC in unsupervised learning in the SPS solution.

In its broad structure, this method is the same for learning driving-related knowledge as it is for learning a first language, although the details may differ in the two cases.

Here is the proposed solution:

1. Start with **N**, a large sample of the language **L** which includes the kinds of dirty data that one would normally find in such a sample.
2. Then compress **N** as far as possible using a relatively mature version of the SP Computer Model which we may call ‘M-SPCM’.
3. With M-SPCM, the original data, **N**, would be reduced to a grammar, **G**, and an encoding of **N** in terms of **G**, where the encoding is called **E**.

Apart from encodings of **N**, **E** is likely also to contain anything else in **N** which cannot be encoded at all, including haphazard features of **N** which may be seen as dirty data.

4. The next step is to discard **E** and retain **G**. This may solve the INF-GEN-DD problems as follows:

- *INF*. **G** is likely to express, or, in the jargon of theoretical linguistics, ‘generate’ an infinite set of patterns, even though it was derived from the finite sample, **N**. This is because **G** would normally contain recursive structures or simple iterations, either of which would yield an infinite set of possibilities.

If **N** is large enough, **G** would generate the majority of recurrent patterns in **L** but there may still be unfamiliar patterns that are missing. However, **G** would normally include rules for all the individual symbols in the symbol set, eg individual letters, digits, and punctuation symbols. This would mean that, where necessary, any possible sequence of those symbols could be learned.

That last point means that, in effect, the recognition or analysis of already-known structures would be combined with the learning of new structures. That combining of recognition/analysis with unsupervised learning is an important feature of how the SPS works.

In short, the SPS provides for the expansion of **G** in the light of new driving experiences, after the initial period of learning. This would help to make up for any deficiencies in that initial learning.

- *GEN*. Since **G** normally contains all the generalisations from **N**, and since **G** is retained, this is likely to solve the GEN problem.
- *DD*. Since **E** normally contains all the dirty data in **N**, and since **E** is discarded, this is likely to solve the DD problem.

6 Evidence and arguments in support of the proposed SPS solution

What evidence or arguments are there in support of the ideas described above? Some possibilities are described here.

6.1 Evidence from computer models of the learning of English-like artificial languages

With the unsupervised learning of artificial English-like languages [13], there is informal evidence that the learning process, which is essentially the compression of **N**, can yield generalisations which are intuitively correct:

“An artificial text with no segmentation markers was prepared from a simple grammar but all instances of two of the (64) sentences generated by the grammar were excluded from the text. When the [SNPR] program was run on this text it successfully retrieved the original grammar despite the fact that the generative range of the grammar was not fully represented in the sample. In the course of building up the grammar it produced many wrong generalisations all of which were corrected. Every one of the correct generalisations, including those required to predict the missing sentences, were retained as permanent fixtures in the grammar.” [13, p. 197]

Here, the discarding of **E** is not made explicit, but from the focus on **G** with no mention of **E**, it is clear that **E** has been excluded from consideration.

A word of caution with this kind of evidence is that: “It is a mistake to allow one’s knowledge of English (say) to dictate what is right and wrong when one is dealing with a text which may look superficially like a subset of English but whose true structure may be significantly different from English.” (*ibid.*).

With regard to ‘dirty data’: “In practice, the programs MK10 and SNPR have been found to be quite insensitive to errors (of omission, addition, or substitution) in their data.” [13, p. 209]. Errors of omission and their corrections may be seen as generalisation, but the correction of errors of addition or substitution are more clearly the correction of dirty data.

6.2 Reasoning from the SPS solution itself

Another approach to the validation of the SPS solution is to consider the workings of the process itself:

- In the SPS solution, SP-patterns or parts of SP-patterns which are relatively frequent are stored in **G**. The assumption or meta-theory which lies behind the SPS solution is that anything that occurs relatively frequently is more likely than not to be part of the relatively stable structures which the learning process is attempting to reveal.

A point of interest is that, in the SPS solution, repeating patterns can include patterns that are discontinuous in the sense that they may be interleaved with other information, and they may be abstractions that contain references to lower-level patterns.

- By contrast with the grammar **G**, the file **E**—which is the encoding of **N** in terms of **G**—is a repository of things that are, individually, relatively rare. In this case, the assumption or meta-theory behind the SPS solution

is that features or SP-patterns that occur rarely are more likely than not to be dirty data or incidental features which are not part of the relatively stable structures which the learning process is designed to reveal. Thus discarding **E** has the desired effect of discarding all or most of the dirty data.

In short, generalisations are likely to be favoured in the creation and retention of **G**, while dirty data and other information with no long-term significance are likely to be eliminated via their assignment to **E** and its subsequent disposal.

6.3 Reasoning from evidence for the importance of information compression in human learning, perception and cognition

Another reason that we may have confidence in the proposed solution is that it is consistent with much evidence for the importance of IC in human learning, perception, and cognition. This comes from two main sources:

- *Direct empirical evidence.* Beginning with pioneering research by Fred Atneave [1, 2], Horace Barlow [4, 5], and others, there has been an accumulation of evidence for the importance of IC in the workings of brains and nervous systems. Much of that evidence is described in [16].
- *Indirect evidence.* The SP Computer Model works exclusively via IC and incorporates the concept of *SP-multiple-alignment* which is a powerful framework for the compression of information [17, Section 5.7].

The central importance of IC in the SP Computer Model, coupled with the versatility of the SP Computer Model in modelling several aspects of intelligence ([14, Chapters 5 to 10], [15, Sections 5 to 14]), is itself evidence for the importance of IC in human intelligence.

6.4 Reasoning from biology and engineering

Yet another reason for having confidence in the SPS solution is that it is consistent with biological arguments for the importance of IC in the workings of brains and nervous systems [16, Section 4].

Even if we knew nothing about specific animals, principles of evolution by natural selection should lead us to expect that IC would be important in every animal's systems for the storage and transmission of information. This is because natural selection would probably favour IC: by allowing more information to be stored in a given storage space, or by requiring a smaller storage space for a given amount of information; and by speeding up the transmission of a given body of

information along a given communication channel, or by requiring a smaller bandwidth for the transmission of a given body of information at a given speed.

Since IC is likely to be important in the workings of brains and nervous systems, it would be surprising if IC were not exploited (via the SPS solution) for the purposes of generalisation and the weeding out of dirty data.

Similar principles apply in engineering. Instead of natural selection we may invoke the artificial selection that applies in engineering when an inefficient product is displaced by something that is more efficient. It is true that the QWERTY phenomenon may sometimes apply—when there are factors that make it difficult to replace an inefficient system with something more efficient—but there is still a general tendency to favour more efficient systems and retire less efficient systems.

6.5 What to do about driving scenarios that occur only rarely, especially if they are dangerous?

As was mentioned earlier, the research by [6] has served a useful purpose in drawing attention to the question of what, if anything, should be done about driving scenarios that occur only rarely. This is clearly of most interest if the rare scenarios are also dangerous. In the light of the proposed solution, two tentative answers are offered here:

- *Make N as big and varied as possible.* A fairly obvious answer to the problem of rare and dangerous scenarios is to make N as big as possible and as varied as possible, aiming to take in possible dangers. This may seem extravagant but the costs associated with a large N will be spread across the many SDVs that may take advantage of the knowledge gleaned from that large N.
- *Take advantage of generalisations.* The whole point of adopting a robust and well-founded system for generalisation is to increase the range of scenarios that one may deal with effectively beyond what we may learn directly from N. If we know what to do about a tiger that has leapt into the back of our vehicle, we should know what to do about a lion or a leopard etc in the same kind of situation.

7 Conclusion

The proposal in this paper for the development of driving expertise in SDVs draws on previous research into the learning of a first language by a child, and also the SPS solution.

Key problems are: INF: How, from a finite sample of driving-related data, **N**, to create an ability to analyse an infinity of different possible driving situations; GEN: How to generalise correctly (without over-generalisations or under-generalisations) from a finite body of driving-related information; DD: How to learn knowledge which is correct, despite the inevitable ‘dirty data’ in the information which is the basis for learning.

In brief, the proposed solution is: Starting with **N**, compress it as much as possible by a learning process that creates a grammar, **G**, and an encoding of **N** in terms of **G** called **E**. Then: discard **E** which contains dirty data (DD), and retain **G** which provides for INF and GEN. Four justifications for the proposed solution are described.

In SDVs, a potentially useful feature of the SPS is that exactly the same software may serve both the analysis of incoming information and the creation of responses by the SDV.

In view of the importance of robust solutions for the INF-GEN-DD problems, and in view of the apparent shortcomings of published research in those area, it seems that *any SDV that has been developed **without** the SPS solution is unlikely to provide the levels of safety that will be demanded by politicians and the general public.*

8 Acknowledgements

I’m grateful to referees ...

9 Funding

There has been no external funding for this research.

10 Software

The ‘SP71’ version of the SP Computer Model may be downloaded via links under the heading ‘SOURCE CODE’ in tinyurl.com/a6h2arhr.

A very similar but slightly earlier version of the program is described, with pseudocode, in [14, Section 3.10 and Chapter 9].

References

- [1] F. Attneave. Some informational aspects of visual perception. *Psychological Review*, 61:183–193, 1954.
- [2] F. Attneave. *Applications of Information Theory to Psychology*. Holt, Rinehart and Winston, New York, 1959.
- [3] C. Badue, R. Guidolini, R. V. Carneiro, P. Azevedo, V. B. Cardoso, A. Forechi, L. Jesus, R. Berriel, T. M. Paixão, F. Mutz, L. de P. Veronese, T. Oliveira-Santos, and A. F. De Souza. Self-driving cars: A survey. *Expert Systems with Applications*, 6(3):113816, 2021.
- [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] S. Beery, Y. Liu, D. Morris, J. Piavis, A. Kapoor, N. Joshi, M. Meister, and P. Perona. Synthetic examples improve generalization for rare classes. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, March 2020.
- [7] C. Brown. *My Left Foot*. Vintage Digital, London, Kindle edition, 2014. First published in 1954.
- [8] M. Daily, S. Medasani, R. Behringer, and M. Trivedi. Self-driving cars. *IEEE Computer*, (December):18–23, 2017.
- [9] S. Manoharan. An improved safety algorithm for artificial intelligence enabled processors in self driving cars. *Journal of Artificial Intelligence and Capsule Networks*, 1(2):95–104, 2019.
- [10] J. Muñoz-Bulnes, C. Fernandez, I. Parra, David Fernández-Llorca, and M. A. Sotelo. Deep fully convolutional networks with random data augmentation for enhanced generalization in road detection. In *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, pages 366–371, 2017.

- [11] V. Palade and J. G. Wolff. A roadmap for the development of the ‘SP Machine’ for artificial intelligence. *The Computer Journal*, 62:1584–1604, 2019. <https://doi.org/10.1093/comjnl/bxy126>, arXiv:1707.00614, bit.ly/2tWb88M.
- [12] S. Shalev-Shwartz, S. Shammah, and A. Shashua. On a formal model of safe and scalable self-driving cars. 2017. arXiv:1708.06374 [cs.RO].
- [13] 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.
- [14] J. G. Wolff. *Unifying Computing and Cognition: the SP Theory and Its Applications*. CognitionResearch.org, Menai Bridge, 2006. ISBNs: 0-9550726-0-3 (ebook edition), 0-9550726-1-1 (print edition). Distributors, including Amazon.com, are detailed on bit.ly/WmB1rs. The print version is produced via print-on-demand from “INGRAM Lightning Source” and, via that technology, is unlikely to go out of print.
- [15] J. G. Wolff. The SP Theory of Intelligence: an overview. *Information*, 4(3):283–341, 2013. arXiv:1306.3888 [cs.AI], bit.ly/1NOMJ6l.
- [16] J. G. Wolff. Information compression as a unifying principle in human learning, perception, and cognition. *Complexity*, 2019:38, February 2019. Article ID 1879746. viXra:1707.0161v3, hal-01624595 v2.
- [17] J. G. Wolff. Mathematics as information compression via the matching and unification of patterns. *Complexity*, 2019:25, 2019. Article ID 6427493, Archives: vixra.org/abs/1912.0100 and hal.archives-ouvertes.fr/hal-02395680. This paper is reproduced in *New Ideas Concerning Science and Technology*, Vol. 13, 12 June 2021, Pages 132–169, doi.org/10.9734/bpi/nicst/v13/8449D.
- [18] J. Zhang and K. Cho. Query-efficient imitation learning for end-to-end autonomous driving. 2016. arXiv:1605.06450 [cs.LG].
- [19] J. Zhao, B. Liang, and Q. Chen. The key technology toward the self-driving car. *International Journal of Intelligent Unmanned Systems*, 6(1):2–20, 2018.
- [20] L. Zhihao, M. Toshiyuki, K. Sasaki, T. Ogata, and S. Sugano. Rethinking self-driving: multi-task knowledge for better generalization and accident explanation ability. 2018. arXiv:1809.11100 [cs.CV].