Short proposal for the development of the SP machine*

J Gerard Wolff† Vasile Palade‡

July 5, 2016

Abstract

This is a short version of a detailed proposal to create a facility to develop the *SP machine*, based on the *SP theory of intelligence*, and with many potential benefits and applications. The first version of the SP machine would be derived directly from the existing SP computer model and developed as a high-parallel, open-source software virtual machine, hosted on a high-performance computer. This would be a means for researchers anywhere in the world to explore what can be done with the SP machine and to create their own versions of it. The new facility would also provide a computational version of SP-neural, a realisation of the SP theory in terms of neurons and their interconnections. As with the main SP machine, this would be a means for researchers to see what can be done with the system and to create new versions of it.

1 Introduction

This is a short version of a proposal to create a facility for the development of the *SP machine*, based on the *SP theory of intelligence*, and with many potential benefits and applications. The more detailed version [13] may be downloaded via bit.ly/1zZjjIs.

*An electronic copy of this document may be downloaded via bit.ly/1SKAjhZ.

†Dr Gerry Wolff, BA (Cantab), PhD (Wales), CEng, MIEEE, MBCS (CITP); CognitionResearch.org, Menai Bridge, UK; jgw@cognitionresearch.org; +44 (0) 1248 712962; +44 (0) 7746 290775; *Skype*: gerry.wolff; *Web*: www.cognitionresearch.org.

‡Dr Vasile Palade, MEng (Tech. Univ. of Bucharest), PhD (University of Galaţi), IEEE Senior Member; Reader in Pervasive Computing, Department of Computing, Coventry University, UK; vasile.palade@coventry.ac.uk; +44 (0) 2477 659190; +44 (0) 7912 043982; *Skype*: v_palade; *Web*: bit.ly/1xNM5xr, bit.ly/1Buyk9E.

In brief, it is envisaged that the SP machine will be created as an open-source, high-parallel, software virtual machine, running on an existing high-performance computer. This would be a means for researchers everywhere to explore what can be done with the SP machine and to create new versions of it.

2 The SP theory and the SP machine

The *SP theory of intelligence*, and its realisation in the SP computer model, is the product of an extended programme of research designed to simplify and integrate observations and concepts relating to 'intelligence', broadly construed. The SP computer model may be regarded as the first version of the SP machine.

The SP system is described most fully in [3] and more briefly in [5].

The SP system can be thought of as a versatile intelligent database system. *It is the versatility of the system, with capabilities across several different areas of intelligence (outlined below), which is perhaps the main strength of the SP system.*

The SP system may also be thought of as a model of human perception, thinking, and learning.

A central idea in the SP system is a powerful concept of *multiple alignment*, borrowed and adapted from bioinformatics. Multiple alignment could be the 'double helix' of intelligence—as significant for 'intelligence' in a broad sense as DNA is for biological sciences.

In accordance with the goal of simplification and integration, the SP system, with multiple alignment centre stage, has strengths in several different areas including [3]: unsupervised learning; the analysis and production of natural language; pattern recognition at multiple levels of abstraction and robust in the face of errors of omission, commission, and substitution; information retrieval; computer vision and modelling aspects of natural vision; information retrieval; several kinds of reasoning (including one-step 'deductive' reasoning, abductive reasoning, reasoning with probabilistic networks and trees, reasoning with if-then rules, nonmonotonic reasoning, 'explaining away', and causal diagnosis); planning; problem solving; and information compression.

The SP theory also provides an interpretation for several aspects of mathematics ([3, Chapter 10], [9]).

The abstract concepts in the SP theory may be mapped into an interpretation of the theory—*SP-neural*—in terms of neurons and their interconnections ([3, Chapter 11], [11]).

3 Advantages of the SP system, and potential benefits and applications

Distinctive features of the SP system, and its advantages compared with other AI-related systems, are described in some detail in [12]. In particular, the SP system has the potential to overcome several problems with 'deep learning in neural networks' [12, Section V], a type of AI system that is attracting a lot of attention at present.

Potential benefits and applications of the SP system include:

- *Development of artificial intelligence.* Almost certainly, the SP system provides a firmer foundation than deep learning and other AI-related alternatives for the development of human-like versatility and adaptability in artificial systems [12].

- *Big data.* The SP system may help to solve nine problems associated with big data, including the potential for big cuts in energy consumption by computers, for big gains in efficiency in the transmission of data, and for taming the great diversity in computing of formalisms and formats for knowledge [8]. Big data is one of the "eight great technologies" identified as priorities by the Government.[1]

- *Robotics and autonomous systems.* The SP system may also help in the development of human-like versatility and adaptability in autonomous robots [7]—another of the Government's "eight great technologies".[2]

- *Pattern recognition, computer vision and natural vision.* The SP system provides a framework for the development of pattern recognition [3, Chapter 6] and computer vision [6], and their integration with other aspects of intelligence. At the same time, it may help to throw light on aspects of natural vision [6].

- *The SP system as a DBMS with intelligence.* The SP system may function as a database management system that provides for the integration, in any combination, of several different formats for knowledge, including: the relational model (with retrieval of information in the manner of query-by-example, creating a join between two or more tables, and aggregation);

[1] See "The 'eight great technologies' which will propel the UK to future growth receive a funding boost", a speech given by the Rt Hon David Willetts MP, former Minister for Universities and Science, *GOV.UK*, 2013-01-23, bit.ly/1wKCqII.

[2] See the previous footnote.

object-oriented models (with class-inclusion hierarchies, part-whole hierarchies and their integration, inheritance of attributes, cross-classification and multiple inheritance); and hierarchical and network models (with discrimination networks) [4]. In addition, the system has strengths in several aspects of intelligence as summarised in Section 2.

- *Medical diagnosis.* The SP system may serve as a repository for medical knowledge and a vehicle for the diagnosis of diseases from incomplete and erroneous information in the manner of pattern recognition, with potential for diagnosis via causal reasoning [2].

- *SP-neural.* As already mentioned (Section 2), multiple alignment and other abstract concepts in the SP theory (*SP-abstract*) may be recast in terms of neurons and their interconnections. This *SP-neural* form of the theory has potential to cast light on the representation and processing of knowledge in the brain, and to suggest lines of investigation in neuroscience [11]. It may also provide the basis for a radically new architecture for computers, outlined below.

- *Other potential benefits and applications.* The SP theory promises deeper insights and better solutions in several other areas of application, including: natural language processing; bioinformatics; structuring of documents; software engineering including automatic programming; information compression; the economical transmission of data; the semantic web; the detection of computer viruses; and data fusion [10].

- *Integration of structures and functions.* The theory promises seamless integration of structures and functions within and between different areas of application, with an overall simplification of computing systems, including software (*ibid.*).

- *Radically new architectures for computers.* SP-abstract and SP-neural both have potential as the basis for radically new architectures for computers with benefits and applications outlined above.

3.1 Potential benefits and applications on relatively short timescales

Some of the potential applications of the SP system may be developed on relatively short timescales using existing high-performance computers or even ordinary computers. These include the SP system as an intelligent database system, and applications in such areas as medical diagnosis; pattern recognition; information

compression; highly-economical transmission of information; bioinformatics; and natural language processing.

3.2 Potential value of the SP concepts

Given its wide scope, the potential impact of the SP theory on the ICT industry is very large. If, as a conservative estimate, the SP theory were to add 5% to the value of ICT products that are sold each year, the potential value of the theory would be at least $190 billion each year, and increasing with the continuing growth of ICT [10, Section 8].

3.3 Endorsements and beneficiaries

The need for this facility has been kindly confirmed by several people with appropriate qualifications and experience, listed in Appendix A.

We believe this research may yield results of value for defence in several of the areas listed above especially: solving problems associated with big data; helping to develop human-like versatility and adaptability in autonomous robots; pattern recognition and computer vision; the SP system as a DBMS with intelligence; software engineering including automatic programming; and data fusion.

As noted in Appendix A, Mr Richard Biers of the Defence Science and Technology Laboratory, Programme Manager Homeland Security (previously Programme Leader Futures and Innovation), has kindly suggested that the SP system is potentially of interest to several areas relevant to security applications.

We also believe that there is potential value in the SP system for national intelligence in several areas including: pattern recognition that is robust in the face of errors of omission, commission, and substitution; several forms of reasoning; solving problems associated with big data; computer vision and its integration with other aspects of intelligene; the SP system as a DBMS with intelligence; software engineering including automatic programming; and natural language processing.

There are many other potential beneficiaries of this research in industry, commerce, and public administration.

4 What is needed to advance these ideas

The SP system has now developed to the point where there are far more avenues to be explored than could be tackled by any one research group. There is a pressing need to open things up to researchers everywhere, very much in the spirit of this quote from the book *Smart Machines* by John Kelly and Steve Hamm, both of IBM:

> "The creation of this new era of [cognitive] computing is a monumental endeavor ... no company can take on this challenge alone. So we look to our clients, university researchers, government policy makers, industry partners, and entrepreneurs—indeed the entire tech industry—to take this journey with us." [1, Preface].

With other kinds of research, all that would be needed would be to encourage researchers to explore issues arising from the research, as described in published papers. But the SP programme of research is different. Published papers are important but to engage fully with the research, researchers need to study and fully understand the SP machine, they need to be able to run it, and they need to be able to create new versions of it.

It is not without interest that this approach to the conduct of research—providing a computational facility to enable researchers anywhere to investigate a given area—is now being promoted by at least the following organisations:

- *The Human Brain Project.* Starting 30 March 2016, the scientific community worldwide can begin exploring the initial versions of the six ICT Platforms of the Human Brain Project.[3]

- *OpenAI.* The OpenAI research organisation is now promoting the "OpenAI Gym" (bit.ly/1Y3cVhc), a platform for artificial intelligence research, including "a site and API allowing people to meaningfully compare performance of their trained agents".[4]

- *Google.* Google has announced:

  - That "TensorFlow", its open-source platform for machine learning, will be accessible by anyone with an internet connection.[5]
  - That some of its software for the processing of natural language will be be made freely available to researchers and others.[6]

- *IBM.* IBM has announced that it is making its 5-qubit quantum processor available to the public via its cloud computing facility.[7]

[3] See "Platforms overview" of the Human Brain Project, bit.ly/1Uj4LxJ.

[4] See "Elon Musks artificial intelligence group opens a 'gym' to train AI", *Popular Science*, 2016-04-27), bit.ly/1VH6fYm.

[5] See "How Google aims to dominate AI", *Popular Science*, 2016-04-29) bit.ly/1klAiF2.

[6] See "Google's artificial-intelligence tool is offered for free", *The Wall Street Journal*, 2016-05-12, on.wsj.com/1qgYAlJ.

[7] See "IBM makes quantum computing available in the cloud", *Computerworld*, 2016-05-04, bit.ly/1TtAlZd.

- *VisionLabs.* A global open-source computer vision project has been set up by VisionLabs.[8]

The following subsections expand on what is required for the SP research facility.

4.1 Parallel processing

Like most AI-related programs, the SP computer model is quite hungry for computer power. The SP machine, like the human brain, needs high levels of parallel processing. The SP computer model, which currently runs as a single-threaded program, needs to be enhanced so that it can run as a software virtual machine with parallel processing on an existing high-parallel computer cluster or supercomputer. With the added speed of processing, the SP machine will be a much more effective and user-friendly tool for research.

It is envisaged that the software virtual machine, and all derivatives from it, will be open-source and governed by an appropriate licence from the Open Software Foundation.

4.2 User-friendly user interface

The user interface for the existing SP computer model is poor. To smooth the path for other researchers, the user interface for the SP machine needs to be very much more 'friendly' and easy to use.

4.3 Permanent

For several important reasons, the facility needs to be permanent in the sense that it will be available for the foreseeable future. Please see [13, Section 3.4] for a description of those reasons.

In case this commitment seems unduly expensive, it should be born in mind: 1) That the proposed facility will be extraordinarily cheap compared with some other research facilities; and 2) The cost of maintaining the system for the foreseeable future is likely to be very much less than the value of research that is undertaken with the system.

Of course, if people stop using the facility, it would make sense to save on development and maintenance costs. But the software and its execution environment should be preserved in accordance with Vint Cerf's arguments outlined in

[8] See "Developer teams up with Facebook and Google to make 'machines see'", *TechWorm*, , 2016-06-05, bit.ly/1ZK1XOv.

[13, Section 3.4], And it would be useful if people could still access the software and run it.

4.4 International

In accordance with established principles and long traditions in science, development of the SP theory and the SP machine should be international and worldwide. Hence, the research facility should be available to researchers anywhere in the world.

It may seem perverse for UK research funding to be providing a research facility for researchers all around the world. But instead of viewing it as an unreasonably generous form of charity, it would be better to view it as a way of getting potentially valuable research done for nothing.

The reason that science has been and still is a highly international enterprise is because of the valuable synergies that arise from the sharing of ideas and results. The cost of the proposed facility is likely to be much smaller than the value of opening up the research to researchers anywhere in the world—meaning a net gain in value for the UK.

4.5 Free

With the possible exception of computationally-intensive uses of the facility, there should be no charge to researchers for exploring what can be done with the SP machine or creating their own versions of it.

The reasons for this proposed feature of the facility are:

- Charging for the use of the facility will discourage researchers from using it. Given that most researchers, in universities and in industry, are under intense pressure, any such discouragement may cause the entire project to fail.
- Compared with, for example, the Hubble space telescope or the Large Hadron Collider, the proposed facility will be extremely cheap. The value of research conducted using the facility may easily exceed the cost of the facility (Sections 3, 4.3, and 4.4).
- The administrative costs of charging for the use of the facility are likely to outweigh the value of moneys received.

4.6 Ease of access

To smooth the path for researchers and to minimise the chance that anyone will be discouraged from using the facility, registration to use the facility, and signing in for any one session, should be very simple.

4.7 Facilitation of software development

Most supercomputers are set up mainly to run computing-intensive programs and make few if any provisions for the needs of software developers: unrestricted access to the machine for frequent short runs of the program that is under development, with good facilities for interactive debugging of programs.

For those reasons, the facility probably needs to be hosted on its own dedicated machine. In the early stages, this can be relatively small and correspondingly inexpensive, with perhaps 200 to 300 cores. With increasing demand, the machine may be progressively expanded.

4.8 Modern development environment

To facilitate the process of development, a modern development environment would be helpful, such as Microsoft's Visual C++, part of Microsoft's Visual Studio.

4.9 A computational version of SP-neural

Although it probably would not require parallel processing, at least initially, it is envisaged that the new facility would provide a preliminary version of an SP-neural computer model. As with the main SP machine, researchers may explore what can be done with this model and they may create their own versions of it.

5 Barriers to progress and how they may be overcome

While it is true that valuable long-range research has been done in research laboratories of commercial companies, such research can easily be terminated if or when a company falls on hard times. As a general rule, commercial companies concentrate on near-market research. While it is true that the SP programme may yield applications on relatively short timescales (Section 3.1), most of what needs to be done should be treated as long-range research. As such, it is best if the risk can be shared widely, as is the case with public funding of research or, sometimes, funding from charitable foundations.

We have not yet discovered any appropriate source of funding from a charitable foundation. Hence, it appears that funding for the development of the proposed research facility would need to come from the Government or comparable sources.

It appears that the only source of funding that might be suitable is the UK's Engineering and Physical Sciences Research Council (EPSRC). However, while we have received positive comments from Mr Liam Blackwell of the EPSRC and also Professor Cliff Brereton, Director of the Hartree Centre (see Appendix A), it appears that none of the EPSRC schemes for funding new research facilities meet the needs that we have identified (Section 4):

- It is clear from discussions with EPSRC staff that it may be difficult for the EPSRC to meet the following requirements: that the facility would be permanent; that the facility would be available to researchers anywhere in the world; and that, with the possible exception of computing-intensive applications, there would be no charges for users.

- Also, our discussions with staff at the Archer supercomputing facility suggest that such a facility would not be a suitable vehicle for the development of the SP machine, because the machine is largely geared to the running of computing-intensive programs and because procedures for gaining access to the machine are not as straightforward as one might wish.

Public funding that meets the needs of this research would by-pass these several barriers to progress.

6 Costs

The current estimate of the non-recurrent costs of setting up the proposed facility is £929,500 [13, Section 4]. The estimated annual cost of maintaining and administering the facility is £261,800 (*ibid.*).

7 Conclusion

We believe the creation of the proposed facility would be a very cost-effective way of advancing the SP system, both in terms of theory and in terms of a range of potential benefits and applications.

The SP system has clear advantages compared with other AI-related systems, and it has many potential benefits and applications. Potential beneficiaries include the MoD, the intelligence agencies, and many other organisations and people in industry, commerce, and public administration.

A Endorsements of the SP programme

The following people, listed alphabetically, have kindly endorsed the proposed development, confirming that they believe the research facility would be useful:

- Professor Pieter Adriaans, Professor of Machine Learning and Artificial Intelligence, University of Amsterdam, The Netherlands.

- Mr Richard Biers, Programme Manager Homeland Security (previously Programme Leader Futures and Innovation), Defence Science and Technology Laboratory, Salisbury, UK.

  In a letter to Dr Gerry Wolff, dated 2015-06-01, he wrote: "I am writing to record my interest in your proposal to develop an open-source software virtual machine, derived from the existing SP Computer Model. I am not an expert in this area, but from our discussions I think that the multiple alignment approach that SP Theory is based upon is promising. As we discussed, I am not in a position to provide funding for your proposal. However if the SP Machine demonstrates significant benefits in addressing 'real-world' problems, I think this would be of interest to several areas relevant to security applications."

- Dr Dorrit Billman, Senior Research Scientist, The Institute for the Study of Learning and Expertise, Computational Learning Laboratory (CSLI), Stanford University, USA.

- Mr Liam Blackwell, Lead, Information & Communications Technologies (Capability) in the EPSRC.

  In an email dated 2014-12-9 he said: "I think this is an interesting area and can see research into intelligent systems of the kind you describe as being one of several priorities for the ICT Theme in the next EPSRC Delivery Plan."

- Professor Cliff Brereton, Director, Hartree Centre, Science & Technology Facilities Council, Daresbury Laboratory, Warrington, UK.

  In an email dated 2015-04-13, he said "I agree that the development would be useful."

- Professor Nick Chater, Professor of Behavioural Science, Warwick Business School, University of Warwick, UK.

- Ms Lorraine Dodd, Research Director, Complex Adaptive Systems, Cranfield Defence and Security, Cranfield University, UK.

- Professor Gordana Dodig-Crnkovic, Professor of Computer Science, School of Innovation, Design and Engineering. Mlardalen University, Sweden.
- Professor Zeno Geradts, Special Chair in Forensic Data Science at the University of Amsterdam, Amsterdam, The Netherlands.
- Professor John Goldsmith, Edward Carson Waller Distinguished Service Professor, Departments of Linguistics and Computer Science, University of Chicago, Chicago, USA.
- Professor Adrian Hopgood, Pro Vice-Chancellor and Dean of the Sheffield Business School at Sheffield Hallam University, Sheffield, UK.
- Professor Pat Langley, Professor and Head of the Computational Learning Laboratory, and Consulting Professor of Symbolic Systems, Stanford University, USA.
- Dr Wim Melis, Senior Lecturer at the University of Greenwich, Greenwich, UK.
- Sergio Navega, Director of Intelliwise Research and Training, Portugal.
- Dr Pierre Perruchet, Lecturer, Laboratoire d'Etude de l'Apprentissage et du Développement (LEAD), Université de Bourgogne, Dijon, France.
- Dr John Pickering, Lecturer in Cognitive Science, Department of Psychology, University of Warwick, Coventry, UK.
- Professor Tim Porter, Professor Emeritus of Mathematics, School of Computer Science, University of Bangor, UK.
- Professor Emmanuel Pothos, Professor of Psychology, Department of Psychology, City University, London, UK.
- Professor David Powers, Professor of Computer Science, Flinders University, Adelaide, Australia.
- Dr Menno van Zaanen, Assistant Professor, Department of Communication and Information Sciences, University of Tilburg, The Netherlands.

References

[1] 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.

[2] J. G. Wolff. Medical diagnosis as pattern recognition in a framework of information compression by multiple alignment, unification and search. *Decision Support Systems*, 42:608–625, 2006. bit.ly/XE7pRG, arXiv:1409.8053 [cs.AI].

[3] 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.

[4] J. G. Wolff. Towards an intelligent database system founded on the SP theory of computing and cognition. *Data & Knowledge Engineering*, 60:596–624, 2007. bit.ly/Yg2onp, arXiv:cs/0311031 [cs.DB].

[5] J. G. Wolff. The SP theory of intelligence: an overview. *Information*, 4(3):283–341, 2013. bit.ly/1hz0lFE.

[6] 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. bit.ly/1scmpV9.

[7] J. G. Wolff. Autonomous robots and the SP theory of intelligence. *IEEE Access*, 2:1629–1651, 2014. bit.ly/1zrSemu.

[8] J. G. Wolff. Big data and the SP theory of intelligence. *IEEE Access*, 2:301–315, 2014. bit.ly/1jGWXDH. This article, with minor revisions, is to be reproduced in Fei Hu (Ed.), *Big Data: Storage, Sharing, and Security (3S)*, Taylor & Francis LLC, CRC Press, 2016, pp. 143–170.

[9] J. G. Wolff. Information compression, intelligence, computing, and mathematics. Technical report, CognitionResearch.org, 2014. bit.ly/1jEoECH, arXiv:1310.8599 [cs.AI]. Submitted for publication.

[10] J. G. Wolff. The SP theory of intelligence: benefits and applications. *Information*, 5(1):1–27, 2014. bit.ly/1lcquWF.

[11] J. G. Wolff. The SP theory of intelligence and the representation and processing of knowledge in the brain. 2016. In preparation. bit.ly/1WDr6Ld, arXiv:1604.05535 [cs.AI].

[12] J. G. Wolff. The SP theory of intelligence: its distinctive features and advantages. *IEEE Access*, 4:216–246, 2016. bit.ly/21gv2jT.

[13] J. G. Wolff and V. Palade. Proposal for the creation of a research facility for the development of the SP machine. Technical report, CognitionResearch.org, 2015. bit.ly/1zZjjIs, arXiv:1508.04570 [cs.AI].