

Introducing Scalable Quantum Approaches in Language Representation

Peter Wittek and Sándor Darányi

Swedish School of Library and Information Science
Göteborg University & University of Borås
Allégatan 1, 50190 Borås, Sweden

Abstract. High-performance computational resources and distributed systems are crucial for the success of real-world language technology applications. The novel paradigm of general-purpose computing on graphics processors (GPGPU) offers a feasible and economical alternative: it has already become a common phenomenon in scientific computation, with many algorithms adapted to the new paradigm. However, applications in language technology do not readily adapt to this approach. Recent advances show the applicability of quantum metaphors in language representation, and many algorithms in quantum mechanics have already been adapted to GPGPU computing. SQUALAR aims to match quantum algorithms with heterogeneous computing to develop new formalisms of information representation for natural language processing in quantum environments.

1 Introduction

Quantum mechanics is a very successful scientific theory for making predictions about systems with inherent ambiguity in them. That natural language bears similarities with such a system is at least plausible. Recent advances in theory and experimentation to apply quantum mechanics to non-quantum domains include the use of quantum algorithms to address, or to more efficiently solve, problems in such domains (including contrasts between classical vs. quantum methods), such as applications of artificial intelligence, information retrieval, and language modelling.

The quantum metaphor promises improved methodologies to capture the subtleties and ambiguities of human language, resulting in optimised algorithms for text processing. The purpose of SQUALAR is to investigate methods borrowed from the field of quantum mechanics in a wide range of large-scale language technology applications by seeking a match between quantum algorithms and heterogeneous computing.

To this end, a scalable environment is a must. Latest trends indicate the rise of a heterogeneous platform in which multi-core central processing

units (CPUs) and graphics processing units (GPUs) work together in a distributed-memory parallelism. CPU-based parallelism has been utilized for decades, and while not without its own problems, it is a mature field and multicore CPUs enable developing faster algorithms with reasonable effort. In this paradigm, there is a considerable overhead on dividing the problem, distributing the bits along a small number of CPU cores, then collecting and merging results. This type of parallelism is available in a wide range of programming languages, although the source code needs to be modified to some extent. GPU-based parallelism is a completely different approach. The overhead of splitting the work is minimal, the number of cores is massive, but the kind of computations that can be split is limited to a simple, single-pass operation. This heterogeneous computing environment has to be studied at different levels to find scalable implementations: low-level linear algebra, numerical methods, kernel methods and manifold learning are candidates for testing, as well as higher level load distribution such as MapReduce [1]. The constraints are as follows:

- Text processing is typically a data-intensive task, and several distributed algorithms have been proposed to deal with large-scale collections on a grid or in a cloud computing environment. MapReduce¹ was originally developed to this end, and mature libraries, such as Cloud9, are readily available [2]. Other libraries, such as Mahout², facilitate the development of complex language technology applications.
- General-purpose computing on the GPU requires considerable effort from developers. Initial results in text processing, however, indicate that the improvement in execution time can be considerable [3–7].
- Quantum methods, on the other hand, rely on linear algebra and other numerical libraries, many of which have already been optimized to utilize the power of GPUs [8–11].

SQUALAR intends to bring the best of two worlds together. By bridging data-intensive text processing with sophisticated quantum modelling of languages, we expect to see major advances in language technology.

The challenges, however, are far from trivial. The major frameworks of GPGPU programming, CUDA and OpenCL, require wrapping in Java, which is the environment of Hadoop, the most mature open source MapReduce implementation. This paper offers an insight on the initial stage of our ongoing investigation.

¹ <http://hadoop.apache.org/mapreduce/>

² <http://mahout.apache.org>

This paper is organized as follows. Section 2 defines what we mean by heterogeneous computing: a distributed system of nodes which are equipped with multicore CPUs and GPUs. Section 3 gives a very short overview of quantum approaches in language processing, with a focus on methods that have the potential for acceleration. Section 4 discusses how we intend to bridge heterogeneous computing and these quantum approaches, and finally Section 5 concludes our paper.

2 Heterogeneous computing

Heterogeneous computing aims to combine the parallelism of traditional multicore CPUs and GPU accelerator cores to deliver unprecedented levels of performance [12]. While the phrase typically refers to single node, a distributed environment may be constructed from such heterogeneous nodes.

CPUs excel in running single-threaded processes, or in multithreaded applications in which a thread often consists of fairly complicated sequential code. Graphics processors are ideally suited for computations that can be run on numerous data elements simultaneously in parallel. This typically involves arithmetic on large data sets (such as matrices) where the same operation can be performed across thousands of elements at the same time. This is actually a requirement for good performance: the software must use a large number of threads. The overhead of creating new threads is minimal compared to CPUs that typically take thousands of clock cycles to generate and schedule, and a low number of threads will not perform well on GPU [13]. The decomposition and scheduling of computation among CPU cores and GPUs are not trivial even on a single node [14–16], and the task is even more complicated for clusters [17]. In order to issue work to several GPUs concurrently, a program needs the same number of CPU threads, each with its own context. All inter-GPU communication takes place via host nodes. Threads can be lightweight (pthreads, OpenMP, etc. [18]) or heavyweight (MPI [19]). Any CPU multi-threading or message-passing API or library can be used, as CPU thread management is completely orthogonal to GPGPU programming. For example, one can add GPU processing to an existing MPI application by porting the compute-intensive portions of the code without changing the communication structure [20]. However, the efficient utilisation of all CPU and GPU cores remains an open question.

While research is being carried out to develop the formal foundations of efficient scheduling and decomposition in multiple heterogeneous

nodes, GPU-based clouds are becoming available ³⁴, and initial investigations have been carried out to develop an efficient MapReduce framework [21, 22]. Like OpenMP and MPI, MapReduce provides an abstraction, a means to distribute computation without burdening the programmer with the details of distributed computing; however, the level of granularity is different [2]. These frameworks are mostly designed to deal with processor-intensive problems and have only rudimentary support for dealing with very large amounts of input data. The strength of MapReduce is data-intensive distributed parallel processing on a massive scale [1]. The potential of combining a data-intensive cloud-based approach with the compute-intensive GPGPU paradigm for sophisticated, large-scale natural language processing is enormous.

3 Quantum approaches in language processing

Metaphors of quantum theory in linguistic applications arose over the last decade [23–28]. The vector space model of information retrieval was first adopted largely because it allowed for a naturally continuous ‘relevance score’ by using the cosine dissimilarity, as opposed a mere binary decision between relevant and irrelevant documents. In a similar fashion quantum mechanics yields a continuous probability that a particular event will be observed, a feature making it useful to reflect on possible similarities with natural language [24]. Moreover, it appears likely that quantum interaction would be of a type where the context of the interaction itself must be incorporated into the model. For example, a measurement in a quantum-scale system will have an impact on the result. If the system is displaying contextual behaviour such as natural languages, then a quantum approach often incorporates this behaviour very naturally [29].

Quantum phenomena in languages may be present at different levels. At sub-word level, terms and documents can be regarded as linear combinations of their semantic features [30], which can account for semantic priming [31].

At word level, a word in semantic space may be likened to a quantum particle. In the absence of context it is in a superposed state, it is a collection of all the possible meanings of the word: $\rho = p_1\rho_1 + \dots + p_m\rho_m$, where ρ is the word in the semantic space as a density matrix, and each i is a basis state representing one of the m senses of the word and the

³ <http://www.hoopoe-cloud.com/>

⁴ <http://aws.typepad.com/aws/2010/11/new-ec2-instance-type-the-cluster-gpu-instance.html>

probabilities p_i sum to unity. Encountering the word in context, however, gives rise to a ‘collapse’ of potential meanings onto an actual one. The context is modelled a projection operator which is applied to a given density matrix corresponding to the state of a word meaning resulting in its ‘collapse’ [32].

Turning to combinations of words, at least two approaches offer solutions. One uses the operator algebra of quantum theory to construct a ‘semantic calculus’ [26, 33]. The other approach encodes word order relying on random indexing [34, 35], using either permutation [36, 37] or circular convolution [31, 37]. The order can also be encoded by tensor product [25, 38].

Using different units of analysis, quantum approaches find their way to applications, most notably:

- Information retrieval: Vector space logic and quantum logic (Neumann algebra) are very similar [26]. In particular, negation has been investigated in depth in [24]. These models may allow new types of queries and also inference [39].
- Memory models: Two schools of memory models are relevant to quantum theory: matrix memory [40, 41, 32], and convolution-correlation memory (holographic-like) [42, 43, 31]. Matrix models are not directly related to QT, but there can be a connection through Heisenberg’s matrix mechanics, which was the first complete and correct definition of quantum mechanics. It is equivalent to the Schrödinger wave formulation of quantum mechanics, and is the basis of Dirac’s bra-ket notation for the wave function. Matrix models can incorporate hierarchical sentence and paragraph representation [44], bridging distributional and symbolic approaches [45], encode location [46], or include sense in a term-space approach [47, 48]. The other approach, convolution memory is particularly useful to encode syntactic information [49, 37].
- Semantic regions: Regions meant to solve the problem to be able to say that apple is a kind of fruit (apple is part of the fruit region), as opposed to modelling that apples and fruit have something to do with one another. Separating hyperspaces may define a semantic region [33]. As an alternative, [50] measures the distance between subspaces spanned by documents by projecting them into one another.
- Spectral theory in mathematics is key to the success of as diverse application domains as quantum mechanics and methods of latent semantic analysis (LSA, [51]) for language processing. In other words, both quantum mechanics and LSA rely on eigenvalue decomposition

for the localization of their respective entities in observation space. This neglected fact, supported by a high number of papers in different disciplines describing the dynamic behaviour of documents and their index terms over time, points at some implicit “energy” inherent in them and in need of quantification. Prominently, theories of word meaning (contextual [52, 53] and referential [54, 55]), currently used in applications trying to capture and exploit semantic content, fall back on quantities of qualities, but quite possibly miss the underlying framework. LSA is just one spectral approach in language representation: [32] demonstrate the quantum collapse of meaning using the hyperspace analogue to language (HAL, [56]).

4 Methods and planned outcomes

With the above plethora of approaches available for testing, the fundamental task of SQUALAR is bridging scalable linear algebra and numerical methods that are widely used in scientific computing with the emerging theories in quantum interaction to enable practical, real-world language technology applications.

The hardware and basic software infrastructure is what we described in section 2: a distributed system consisting of heterogeneous nodes which combine multicore CPUs and GPUs (top part of Figure 1). Since hardware virtualization is already at consumer level, the distributed system can be either a privately owned cluster or grid, or a high-performance computing cloud provided by a third-party.

Without going into details, algorithms in linear algebra are the most obvious candidates for acceleration on graphics hardware (middle part of Figure 1, left). Vector space models of semantics can be implemented by accelerated BLAS libraries [8, 10], including operator algebra for semantic inference [24, 26]. Matrix decompositions and dimension reduction that also play an important role in understanding semantics are currently limited to matrices of limited sizes [11]. Convolution, which plays an important part in encoding term positions [31, 37], can be mapped to the frequency domain by Fourier transformation, where the operation simplifies to a simple multiplication. Fast Fourier transformation on GPUs is a classical area for acceleration [57]. More complex examples in accelerated quantum methods [58, 59] and related visualization [60] are awaiting appropriate metaphors in language processing.

Approaching from existing language processing algorithms, if a sufficient metaphor cannot be found or if the method does not lend itself

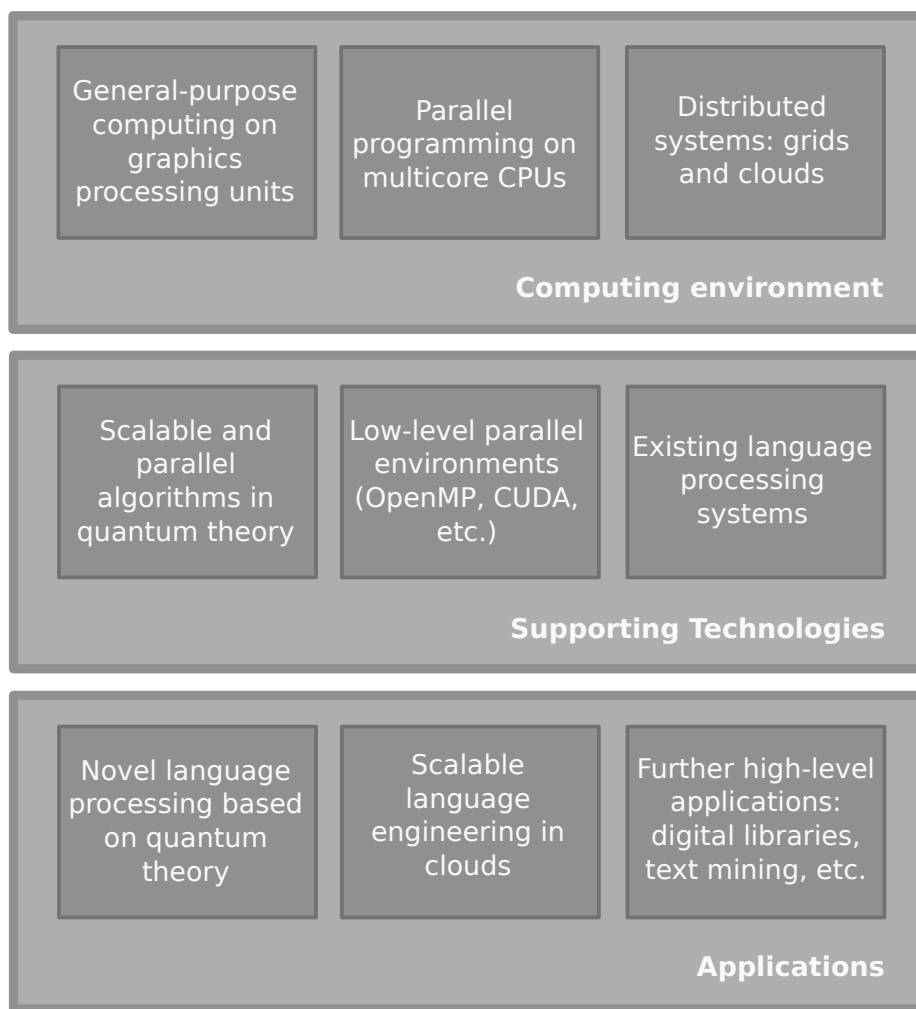


Fig. 1. An overview of the SQUALAR framework

easily to any of the methods described above, lower level libraries can be used for developing multithreaded, GPU-based implementations (middle part of Figure 1, right and middle).

If we focus on a single computer, we will be able to perform operations several folds faster, gaining new insights on language technology (bottom part of Figure 1, left). By providing a high-level load balancing mechanism, the potential of compute and data-intensive processing can be released in a distributed environment for web-scale applications (bot-

tom part of Figure 1, middle). Some machine learning algorithms, such as support vector machines, have already been adopted to graphics hardware [61]. Combining these with the above, we gain powerful text mining applications (bottom part of Figure 1, right). Since Information Retrieval has already begun experimenting with a wide range of quantum theory based metaphors, this field has the most to benefit.

5 Conclusion

Whether language to some extent shares a conceptual framework with quantum mechanics, and if thereby some linguistic phenomena could be eventually modelled on physical ones, is a research question of interest to SQUALAR. We trust that by better mastering the match between quantum algorithms and GPU computing, web-scale applications will become feasible.

The fundamental tasks and challenges of the project are the following:

- Rephrasing natural language processing and text mining algorithms in quantum domain to use compute-intensive heterogeneous programming model;
- Data and compute-intensive distributed and cloud computing applications with heterogeneous hardware;
- Performance evaluation of heterogeneous hardware for natural language processing tasks;
- Trade-offs of using scalable quantum models in language engineering;
- Exploiting heterogeneous architectures to accelerate sophisticated language processing.

6 Acknowledgement

We would like to thank Lance de Vine (Queensland University of Technology) for discussions related to ideas presented in this paper. This work was also supported by Amazon Web Services.

References

1. Dean, J., Ghemawat, S.: MapReduce: Simplified data processing on large clusters. In: Proceedings of OSDI-04, 6th International Symposium on Operating Systems Design & Implementation, San Francisco, CA, USA (December 2004)
2. Lin, J., Dyer, C.: Data-Intensive Text Processing with MapReduce. Morgan & Claypool (2010)

3. Cavanagh, J., Potok, T., Cui, X.: Parallel latent semantic analysis using a graphics processing unit. In: Proceedings of GECCO-09, 11th Annual Conference Companion on Genetic and Evolutionary Computation Conference: Late Breaking Papers, Montreal, QC, Canada (July 2009) 2505–2510
4. Ding, S., He, J., Yan, H., Suel, T.: Using graphics processors for high performance IR query processing. In: Proceedings of WWW-09, 18th International Conference on World Wide Web, Spain, Madrid (April 2009) 421–430
5. Zhang, Y., Mueller, F., Cui, X., Potok, T.: Large-scale multi-dimensional document clustering on GPU clusters. In: Proceedings of IDPDS-10, 24th International Parallel and Distributed Computing Symposium, Atlanta, GA, USA (April 2010)
6. Byna, S., Meng, J., Raghunathan, A., Chakradhar, S., Cadambi, S.: Best-effort semantic document search on GPUs. In: Proceedings of GPGPU-10, 3rd Workshop on General-Purpose Computation on Graphics Processing Units, New York, NY, USA (2010) 86–93
7. Wei, Z., JaJa, J.: A fast algorithm for constructing inverted files on heterogeneous platforms. In: Proceedings of IPDPS-11, 25th International Parallel and Distributed Computing Symposium, Anchorage, AK, USA (May 2011)
8. Krüger, J., Westermann, R.: Linear algebra operators for GPU implementation of numerical algorithms. In: Proceedings of SIGGRAPH-05, 32nd International Conference on Computer Graphics and Interactive Techniques, Los Angeles, CA, USA (July 2005) 234–242
9. Galoppo, N., Govindaraju, N., Henson, M., Bondhugula, V., Larsen, S., Manocha, D.: Efficient numerical algorithms on graphics hardware. In: Proceedings of EDGE-06, Workshop on Edge Computing Using New Commodity Architectures, Chapel Hill, NC, USA (May 2006)
10. Barrachina, S., Castillo, M., Igual, F., Mayo, R., Quintana-Orti, E.: Evaluation and tuning of the level 3 CUBLAS for graphics processors. In: Proceedings of IPDPS-08, 22nd International Symposium on Parallel and Distributed Processing, Miami, FL, USA (April 2008) 1–8
11. Lahabar, S., Narayanan, P.: Singular value decomposition on GPU using CUDA. In: Proceedings of IPDPS-09, 23rd International Symposium on Parallel and Distributed Processing, Rome, Italy (May 2009)
12. Brodtkorb, A., Dyken, C., Hagen, T., Hjelmervik, J., Storaasli, O.: State-of-the-art in heterogeneous computing. *Scientific Programming* **18**(1) (2010) 1–33
13. Kirk, D., Hwu, W.: Programming massively parallel processors: A hands-on approach. (2009)
14. Jiménez, V., Vilanova, L., Gelado, I., Gil, M., Fursin, G., Navarro, N.: Predictive runtime code scheduling for heterogeneous architectures. *High Performance Embedded Architectures and Compilers* (2009) 19–33
15. Lee, S., Min, S., Eigenmann, R.: OpenMP to GPGPU: a compiler framework for automatic translation and optimization. In: Proceedings of PPOPP-09, 14th Symposium on Principles and Practice of Parallel Programming. (February 2009) 101–110
16. Luk, C., Hong, S., Kim, H.: Qilin: Exploiting parallelism on heterogeneous multiprocessors with adaptive mapping. In: MICRO-42, 42nd Annual IEEE/ACM International Symposium on Microarchitecture, New York, NY, USA (December 2009) 45–55
17. Phillips, J., Stone, J., Schulten, K.: Adapting a message-driven parallel application to GPU-accelerated clusters. In: Proceedings of SC-08, 21st Conference on Supercomputing, Austin, TX, USA (November 2008) 1–9

18. Kuhn, B., Petersen, P., O'Toole, E.: OpenMP versus threading in C/C++. *Concurrency: Practice and Experience* **12**(12) (2000) 1165–1176
19. Koop, M., Sur, S., Gao, Q., Panda, D.: High performance MPI design using unreliable datagram for ultra-scale InfiniBand clusters. In: *Proceedings of ISC-06, 21st Annual International Conference on Supercomputing*, Dresden, Germany (June 2006) 180–189
20. : NVida Compute Unified Device Architecture Best Practices Guide 3.2 (2010)
21. Shirahata, K., Sato, H., Matsuoka, S.: Hybrid map task scheduling on GPU-based heterogeneous clusters. In: *Proceedings of CloudCom-10, The 2nd International Conference on Cloud Computing*, Indianapolis, IN, USA (December 2010)
22. Stuart, J., Owens, J.: Multi-GPU MapReduce on GPU clusters. In: *Proceedings of IPDPS-11, 25th International Parallel and Distributed Computing Symposium*, Anchorage, AK, USA (May 2011)
23. Aerts, D., Aerts, S., Broekaert, J., Gabora, L.: The violation of bell inequalities in the macroworld. *Foundations of Physics* **30**(9) (2000) 1387–1414
24. Widdows, D., Peters, S.: Word vectors and quantum logic: Experiments with negation and disjunction. In: *Proceedings of MoL-03, 8th Mathematics of Language Conference*. Volume 8., Bloomington, IN, USA (June 2003) 141–154
25. Aerts, D., Czachor, M.: Quantum aspects of semantic analysis and symbolic artificial intelligence. *Journal of Physics A: Mathematical and General* **37** (2004) L123–L132
26. van Rijsbergen, C.J.: *The Geometry of Information Retrieval*. Cambridge University Press, New York, NY, USA (2004)
27. Widdows, D.: *Geometry and meaning*. CLSI Publications (2004)
28. Bruza, P., Widdows, D., Woods, J.: A quantum logic of down below. In Engesser, K., Gabbay, D., Lehmann, D., eds.: *Handbook of Quantum Logic and Quantum Structures*. Volume 2. (2009)
29. Kitto, K.: Why quantum theory? In: *Proceedings of QI-08, 2nd International Symposium on Quantum Interaction*, Oxford, UK (March 2008) 11–18
30. Lyons, J.: *Semantics*. Cambridge University Press, New York, NY, USA (1977)
31. Jones, M., Mewhort, D.: Representing word meaning and order information in a composite holographic lexicon. *Psychological review* **114**(1) (2007) 1–37
32. Bruza, P., Woods, J.: Quantum collapse in semantic space: interpreting natural language argumentation. In: *Proceedings of QI-08, 2nd International Symposium on Quantum Interaction*, Oxford, UK (March 2008)
33. Widdows, D.: Semantic vector products: Some initial investigations. In: *Proceedings of QI-08, 2nd International Symposium on Quantum Interaction*, Oxford, UK (March 2008)
34. Kanerva, P., Kristofersson, J., Holst, A.: Random indexing of text samples for latent semantic analysis. In: *Proceedings of CogSci-00, 22nd Annual Conference of the Cognitive Science Society*. Volume 1036., Philadelphia, PA, USA (2000)
35. Sahlgren, M.: An introduction to random indexing. In: *Proceedings of TKE-05, Methods and Applications of Semantic Indexing Workshop at the 7th International Conference on Terminology and Knowledge Engineering*, Copenhagen, Denmark (August 2005)
36. Sahlgren, M., Holst, A., Kanerva, P.: Permutations as a means to encode order in word space. In: *Proceedings of CogSci-08, 30th Annual Meeting of the Cognitive Science Society*, Washington, DC, USA (July 2008)
37. De Vine, L., Bruza, P.: Semantic oscillations: Encoding context and structure in complex valued holographic vectors. In: *Proceedings of QI-10, 4th Symposium on*

- Quantum Informatics for Cognitive, Social, and Semantic Processes, Arlington, VA, USA (November 2010) 11–13
38. Mitchell, J., Lapata, M.: Vector-based models of semantic composition. In: Proceedings of ACL-08, 46th Annual Meeting of the Association for Computational Linguistics, Columbus, Ohio (June 2008) 236–244
 39. Song, D., Lalmas, M., van Rijsbergen, C., Frommholz, I., Piwowarski, B., Wang, J., Zhang, P., Zuccon, G., Bruza, P., Arafat, S., et al.: How quantum theory is developing the field of Information Retrieval. In: Proceedings of QI-10, 4th Symposium on Quantum Informatics for Cognitive, Social, and Semantic Processes, Arlington, VA, USA (November 2010) 105–108
 40. Humphreys, M., Bain, J., Pike, R.: Different ways to cue a coherent memory system: A theory for episodic, semantic, and procedural tasks. *Psychological Review* **96**(2) (1989) 208–233
 41. Wiles, J., Halford, G., Stewart, J., Humphreys, M., Bain, J., Wilson, W.: Tensor models: A creative basis for memory retrieval and analogical mapping. In Dartnall, T., ed.: *Artificial intelligence and creativity*. Kluwer Academic (1994) 145–159
 42. Plate, T.: Holographic reduced representations: Convolution algebra for compositional distributed representations. In: Proceedings of IJCAI-91, 12th International Joint Conference on Artificial Intelligence, Sydney, Australia (August 1991) 30–35
 43. Plate, T.: Holographic reduced representations. *IEEE Transactions on Neural Networks* **6**(3) (1995) 623–641
 44. Antonellis, I., Gallopoulos, E.: Exploring term-document matrices from matrix models in text mining. In: Proceedings of SDM-06, Text Mining Workshop in conjunction with the 6th SIAM International Conference on Data Mining, Bethesda, MD, USA (April 2006)
 45. Rudolph, S., Giesbrecht, E.: Compositional matrix-space models of language. In: Proceedings of ACL-10, 48th Annual Meeting of the Association for Computational Linguistics, Uppsala, Sweden (July 2010) 907–916
 46. Rölleke, T., Tsirikas, T., Kazai, G.: A general matrix framework for modelling information retrieval. *Information Processing & Management* **42**(1) (2006) 4–30
 47. Swen, B.: A sense matrix model for information retrieval. Technical report, Technical report TR-2004-2 of ICL-PK (2004)
 48. Novakovitch, D., Bruza, P., Sitbon, L.: Inducing shades of meaning by matrix methods: a first step towards thematic analysis of opinion. In: Proceedings of SEMAPRO-09, 3rd International Conference on Advances in Semantic Processing, Sliema, Malta (October 2009) 86–91
 49. Jones, M., Kintsch, W., Mewhort, D.: High-dimensional semantic space accounts of priming. *Journal of memory and language* **55**(4) (2006) 534–552
 50. Zuccon, G., Azzopardi, L., Rijsbergen, C.: Semantic spaces: Measuring the distance between different subspaces. In: Proceedings of QI-09, 3rd International Symposium on Quantum Interaction, Saarbruecken, Germany (March 2009) 225–236
 51. Deerwester, S., Dumais, S., Furnas, G., Landauer, T., Harshman, R.: Indexing by latent semantic analysis. *Journal of the American Society for Information Science* **41**(6) (1990) 391–407
 52. Wittgenstein, L.: *Philosophical Investigations*. Blackwell Publishing, Oxford, UK (1967)
 53. Harris, Z.: Distributional structure. In Harris, Z., ed.: *Papers in structural and transformational Linguistics*. Formal Linguistics. Humanities Press, New York, NY, USA (1970) 775–794

54. Peirce, C.: Logic as semiotic: The theory of signs. In Peirce, C., Buchler, J., eds.: *Philosophical Writings of Peirce*. Dover Publications (1955) 98–119
55. Frege, G.: Sense and reference. *The Philosophical Review* **57**(3) (1948) 209–230
56. Lund, K., Burgess, C.: Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods Instruments and Computers* **28** (1996) 203–208
57. Govindaraju, N., Lloyd, B., Dotsenko, Y., Smith, B., Manferdelli, J.: High performance discrete Fourier transforms on graphics processors. In: *Proceedings of SC-08, 21st Conference on Supercomputing*, Austin, TX, USA (November 2008)
58. Ufimtsev, I., Martínez, T.: Graphical processing units for quantum chemistry. *Computing in Science & Engineering* **10**(6) (2008) 26–34
59. Watson, M., Olivares-Amaya, R., Edgar, R., Aspuru-Guzik, A.: Accelerating correlated quantum chemistry calculations using graphical processing units. *Computing in Science & Engineering* **12**(4) (2010) 40–51
60. Stone, J., Saam, J., Hardy, D., Vandivort, K., Hwu, W., Schulten, K.: High performance computation and interactive display of molecular orbitals on GPUs and multi-core CPUs. In: *Proceedings of GPGPU-09, 2nd Workshop on General Purpose Processing on Graphics Processing Units*, Washington, DC, USA (March 2009) 9–18
61. Catanzaro, B., Sundaram, N., Keutzer, K.: Fast support vector machine training and classification on graphics processors. In McCallum, A., Roweis, S., eds.: *Proceedings of ICML-08, 25th Annual International Conference on Machine Learning*, Helsinki, Finland (July 2008) 104–111