

Quantum Mathematics in Artificial Intelligence

Dominic Widdows

IonQ, Inc., 4505 Campus Drive, College Park, MD 20740, USA

WIDDOWS@IONQ.COM

Kirsty Kitto

University of Technology Sydney, PO Box 123, Broadway, NSW 2007, Australia

KIRSTY.KITTO@UTS.EDU.AU

Trevor Cohen

University of Washington, Box 358047, 850 Republican St, Seattle, WA 98109, USA

COHENTA@UW.EDU

Abstract

In the decade since 2010, successes in artificial intelligence have been at the forefront of computer science and technology, and vector space models have solidified a position at the forefront of artificial intelligence. At the same time, quantum computers have become much more powerful, and announcements of major advances are frequently in the news.

The mathematical techniques underlying both these areas have more in common than is sometimes realized. Vector spaces took a position at the axiomatic heart of quantum mechanics in the 1930s, and this adoption was a key motivation for the derivation of logic and probability from the linear geometry of vector spaces. Quantum interactions between particles are modelled using the tensor product, which is also used to express objects and operations in artificial neural networks.

This paper describes some of these common mathematical areas, including examples of how they are used in artificial intelligence (AI), particularly in automated reasoning and natural language processing (NLP). Techniques discussed include vector spaces, scalar products, subspaces and implication, orthogonal projection and negation, dual vectors, density matrices, positive operators, and tensor products. Application areas include information retrieval, categorization and implication, modelling word-senses and disambiguation, inference in knowledge bases, decision making, and semantic composition.

Some of these approaches can potentially be implemented on quantum hardware. Many of the practical steps in this implementation are in early stages, and some are already realized. Explaining some of the common mathematical tools can help researchers in both AI and quantum computing further exploit these overlaps, recognizing and exploring new directions along the way.

1. Introduction

Vector space models have solidified a position at the forefront of artificial intelligence (AI), and are a common building block for many kinds of systems. From early uses in information retrieval and image processing, they have spread to semantic modeling in many systems, reaching even greater prominence during the past decade due to their use in deep learning. However, vector spaces have a long and illustrious history, and as an analytical technique, vector spaces were established in the mid-1800s. They also provide the foundation for the underlying mathematical model for quantum theory.

This paper aims to support readers who are familiar with the use of vectors in AI, and interested in mathematical concepts from quantum theory, to make use of this promising set of results and novel techniques in the field of AI. Section 2 summarizes the use of vector spaces in AI and quantum mechanics, including an explanation of how the representation of wavefunctions leads to the study of vector spaces and linear operators in formal quantum mechanics. Section 3 explains the quantum model for information retrieval, including the use of the quantum logic of subspaces for modelling concepts and query operations. Section 4 looks in more detail at models for categories and hypernyms, motivating the use of density matrices and positive operators from quantum mechanics. Section 5 surveys vector models for semantic composition, in particular the use of tensor products

Table 1: A term–document matrix

	Doc ₁	Doc ₂	Doc ₃	...	Doc _m
Term ₁	M_{11}	M_{12}	M_{13}	...	M_{1m}
Term ₂	M_{21}	M_{22}	M_{23}	...	M_{2m}
Term ₃	M_{31}	M_{32}	M_{33}	...	M_{3m}
...
Term _n	M_{n1}	M_{n2}	M_{n3}	...	M_{nm}

and related operators. Section 6 presents some of the early implementations of these methods on physical quantum hardware, adding fresh vitality to this area of research.

Our hope is that AI practitioners with a grounding in linear algebra will find that they already know much of the mathematics they need to understand large parts of quantum theory, and to begin exploring ways in which they might apply methods inspired by quantum mathematics to the generation of new results in AI.

2. Vector Space Models in AI and Quantum Physics

Vectors have their roots in the beginning of Euclidean geometry, with the notion that a straight line can be drawn between any two points and indefinitely extended (*Elements of Geometry*, Book I, Axioms 1 and 2). By the mid-1840s, 3-dimensional vectors and 1-dimensional scalars were thoroughly explored in Hamilton’s work on quaternions; Cauchy and Hamilton collaborated on the development of matrices and linear algebra; and the algebra of vectors or ‘extended quantities’ (German *Ausdehnungslehre*) and their products in any number of dimensions is fully-formed in the work of Grassmann (1862), a summary of which is made available by Fearnley-Sander (1979).

Since then vectors have played a ubiquitous role in scientific advances across a great range of fields, including differential geometry, electromagnetism, relativity, machine learning: and of particular interest in this paper, artificial intelligence and quantum theory.

2.1 Vector Models in AI

The use of vector spaces in AI goes back at least to information retrieval in the 1960’s (Switzer, 1965), the vector model of the SMART retrieval system being a particularly influential example dating from this pioneering period (Salton & McGill, 1983). Introductions to the vector model for information retrieval are given by Widdows (2004, Ch 5) and Manning, Raghavan, and Schütze (2008, Ch 14). In such a system, a document collection is processed into a large table or matrix M whose rows represent terms, whose columns represent documents, and where the entry M_{ij} contains a number measuring the extent to which term i appears in document j (as is depicted in Table 1).

The rows (or columns) in any such matrix can be added together by adding the corresponding coordinates, and multiplied by a scalar by multiplying each coordinate, so it is immediately apparent that the rows (or columns) of any such matrix can be thought of as vectors in a vector space, whose dimension is the number of columns (or rows).

Such a general description is naturally available as a model for many other situations that are relevant to AI. Common examples include:

- The adjacency matrix of any weighted graph (Bollobás, 1998, Chapter VIII).
- A collection of grayscale images, each of the same width and height, where a row represents an image, columns represent individual pixels, and each coordinate represents the shade for that pixel in that image (Géron, 2019, Ch 3).

- A dataset from a medical study where each row corresponds to an individual patient and columns correspond to a measurable property or “vital sign” such as height, weight, age, or blood pressure values. Many statistical learning techniques begin by assuming their inputs can be modelled in such a fashion (Hastie, Tibshirani, & Friedman, 2001).
- Any collection of data projected into a lower-dimensional space, for example, the ‘best fit’ subspace produced by a decomposition technique such as Singular Value Decomposition. In the case of textual data, the technique of projecting a term-document matrix onto such a subspace is referred to as Latent Semantic Analysis (Landauer & Dumais, 1997), and today it has many variants, surveyed from an AI point of view by Turney, Pantel, et al. (2010) and from a machine learning point of view by Géron (2019, Ch. 8).
- The activity states of a connectionist network (Hinton, 1990; Smolensky, 1990, Definition 2.1). Work on connectionist representations in AI contributed greatly to the development of today’s powerful neural network models (Géron, 2019, Ch 10).
- The output of a nonlinear training / learning algorithm, for example, the feature-weights learned by a neural network. Such models have become increasingly prominent in the last few years: techniques that use several intervening layers in the network are often referred to as *deep learning*, and whether or not a network is deep in this sense, the inputs, outputs, and learned parameters for intervening layers are typically represented as vectors (Géron, 2019, Ch. 10). When applied to textual data the resulting vector models are often referred to as *word embeddings* in more recent papers since Mikolov, Chen, Corrado, and Dean (2013).

Thus there are many related techniques for deriving collections of vectors from empirical observations, and these techniques have become a crucial cornerstone of AI. Key mathematical benefits include expressing addition of vectors as a pairwise-sum of the corresponding coordinates (for vectors $x = (x_1, \dots, x_n)$ and $y = (y_1, \dots, y_n)$, their vector sum $x + y = (x_1 + y_1, \dots, x_n + y_n)$), and measuring similarity (for example, using a scalar product where $x \cdot y = \sum_{i=1}^n x_i y_i$). For a beginner’s introduction to measuring similarity and distance, and the relationship between scalar products, cosine similarity, and Euclidean distance, see Widdows (2004, Ch 4, 5).

The graded way similarity can be expressed in vector models contrasts with classical Boolean logic: for example, the introduction of cosine-similarity between query vectors and document vectors as a continuous (hence graded) measure of relevance in information retrieval enabled systems to return the *most relevant* documents ranked in order. This property became vital once document collections became large enough that simply marking documents as ‘relevant’ or ‘not relevant’ left too many ‘relevant’ documents for a user to read.

The use of vector sums for generating graded probabilities in this manner has been termed ‘superposition’ in physics. It enables a combination that is ‘some of each’, which contrasts with both Boolean set intersection which is sometimes too specific, and Boolean set union which is sometimes too general. Historically, this modeling requirement was recognized very early in the development of quantum theory, where the same vector sum technique is used to effectively represent particles that can be measured in more than one way (this usage is explained in more detail in Section 2.3).

The adoption of similar methods in AI is more recent, but stems from a similar motivation, as expressed for example by Smolensky (1986):

In the subsymbolic paradigm, the semantically interpretable entities are activation patterns, and these combine by *superposition*: activation patterns superimpose upon each other, the way that wave-like structures do in physical systems. This difference is another manifestation of moving the formalization from the discrete to the continuous ...

The core motivation for using vector spaces in AI still has its roots in the benefits of using continuous representations to complement and sometimes improve upon discrete Boolean models.

These representations can then be put to good use in problems involving classification (Manning et al., 2008, Ch 14), (Géron, 2019, Ch 3-5) and logical reasoning (Widdows & Cohen, 2015). Machine learning techniques for building such representations have improved considerably, with the use of deep neural networks, sequence modelling, transformers and attention-based techniques (for a survey, see Géron (2019, Ch 16)). Vector models are used throughout AI and related fields, all the way from pioneering work in information retrieval to today’s neural network language models. So it makes sense to ask whether vector methods derived from other areas of science are useful for learning and reasoning with such models.

2.2 Why Quantum Theory?

Quantum physics appears often and prominently in the science and technology news. Two popular quantum physics topics are immediately relevant to AI: firstly, the suggestion that natural intelligence may involve quantum physics directly, put forth most famously by Penrose (1999); secondly, the increasing technological progress being made in quantum computing, described by e.g., Rieffel and Polak (2014), Bernhardt (2019). The popular appeal of quantum physics and computing is easy to appreciate: it challenges established norms; it is at forefront of innovation; it is real yet mysterious (Aaronson, 2013).

It is commonly said that quantum theory applies only to particularly small physical systems, because many of the discrete quantization and interference effects explained by quantum mechanics are not significantly observable at large scale. However, the mathematics developed as part of vector models in quantum theory can also be used to describe larger human-scale phenomena, such as individuals or groups making decisions or searching for information. One of the pioneers and leaders in this area has been the physicist Diederik Aerts (see Aerts, Durt, Grib, Van Bogaert, and Zapatrin (1993) and onwards). The techniques and models proposed by Aerts and others are sometimes called quantum-inspired, quantum-like, or generalized quantum models (Khrennikov, 2010). This makes a case that quantum *mathematics* is worth investigating in situations outside the traditional microscopic domain of quantum *physics*, without answering the challenging ontological questions that arise when asking how quantum physics might affect larger systems. In this context, the main argument in favor of quantum mathematical models is utilitarian: many solutions to known problems have been improved in a number of different ways. Books and papers discussing the application of quantum mathematics are available in several fields including: information retrieval (Melucci, 2015); machine learning (Wittek, 2014); cognitive science (Busemeyer & Bruza, 2012); and economics (Khrennikov, 2010).

Within AI itself, the book by Wichert (2020) provides a graduate-level text on quantum computing and AI, grounded in computational theory, physics, and quantum algorithms. In contrast, this paper provides a much more thorough introduction to tensors and their uses. Potential quantum approaches in AI are well-surveyed in the article by Ying (2010), which gives a thorough introduction to qubits and quantum gates. In comparison, our work here focusses upon two decade’s hindsight and progress in the use of vector models and related mathematics including subspaces, projections, and tensor products which has flourished in AI.

Simultaneously, the increasing power of quantum computers has turned quantum mechanics from a fascinating mystery to an increasingly practical opportunity, and quantum advantages such as the promise of exponential memory savings are bound to attract the attention of researchers in AI. Some recent advances in AI and quantum computing are presented in Section 6.

A question remains: *Why* does the mathematics of quantum theory appear to support these types of scenarios? Many different explanations are possible depending upon the philosophical stance that one adopts, but for now we will remain agnostic. We will return to this important question in Section 7.6, with more insights from the mathematics of quantum theory to support our proposed answer to this question. For now we will continue with a brief introduction to the topic of quantum theory itself.

2.3 Formal Quantum Mechanics

To appreciate the appeal of quantum-inspired techniques in vector space models we must work to bridge a gap between the physical results of quantum theory (as predicted by quantum mechanics), and its mathematical formulation (in terms of state vectors and operators on Hilbert space). In the process, we hope to show that topics considered to be philosophically challenging in quantum physics, such as non-locality, quantum disjunction, and entanglement, are mathematically quite straightforward. Many of them rely upon simple operations such as projecting lines onto planes, finding the plane spanned by two lines, and multiplying vectors to generate square matrices which can be added together. The mathematical formulation of quantum mechanics mainly uses standard linear algebra, which means that many of its techniques are already familiar to those working in AI. An excellent introduction is provided by Isham (1995) and of course Dirac (1930). Of particular note, the formalism of quantum theory requires us to move from a model of reality that assumes localised particles, to one that assumes waves to be fundamental, and these can be represented as vectors using Fourier series.

Quantum theory was invented to describe the behavior of subatomic particles, and has successfully explained and predicted many behaviors that could not be adequately accounted for by late 19th century classical conceptualizations. These classical models understood the universe as made up of particles, imagined to consist of Euclidean ‘points’ (that is, things that have no spatial extent of their own, and can be located anywhere in a continuous space), and vibrating waves such as light.¹ In the early decades of the 20th century, particle-like behavior was observed in light using experiments that increased our understanding of the polarization and interference of photons (Dirac, 1930, Ch 1). More experiments revealed that matter exhibited wave-like as well as particle-like behaviour. A new theory was needed, one in which both ‘matter’ such as atoms, electrons, and protons, and ‘radiation’ such as light exhibited both wave-like properties (such as interference) and particle-like properties (such as coming in discrete chunks).

2.4 Waves as Vectors Using Fourier Series

Mathematically, any wave can be represented as a function. For example, a standing wave on a string of length L is represented as a function $f : x \in [0, L] \rightarrow \mathbb{R}$, and a dynamic (moving) wave is represented as a function of two variables $f : (x \in [0, L], t \in \mathbb{R}) \rightarrow \mathbb{R}$. Real functions can be multiplied by any real number, and added together pairwise, and with these natural definitions, the standard vector space axioms are satisfied: thus the set of all such real functions forms a vector space. Functions to the reals are vectors.

Moreover, piecewise smooth functions of one variable can be represented as weighted sums of their basic constituents. Such representations were introduced by Joseph Fourier (1768–1830), who demonstrated that under certain continuity conditions a real function f over the interval $[0, 2\pi]$ can be uniquely represented as a sum of harmonic functions,

$$f(x) = \sum_0^{\infty} a_k \sin(kx) + b_k \cos(kx).$$

This sum is called the Fourier series of f , and the a_k and b_k are called its Fourier coefficients. The example in Figure 1 shows the first few harmonic functions that sum up to form closer and closer approximations to a square wave, the expansion in this case being $\sum_{n=1,3,5,\dots} \sin(nx)/n$.

1. This is an accepted if a somewhat simplified view of physics before 1900. For example, Sir Isaac Newton postulated the existence of distinct corpuscles of light, and both Democritus and Plato postulated the existence of particles with some sort of shape or spatial extent.

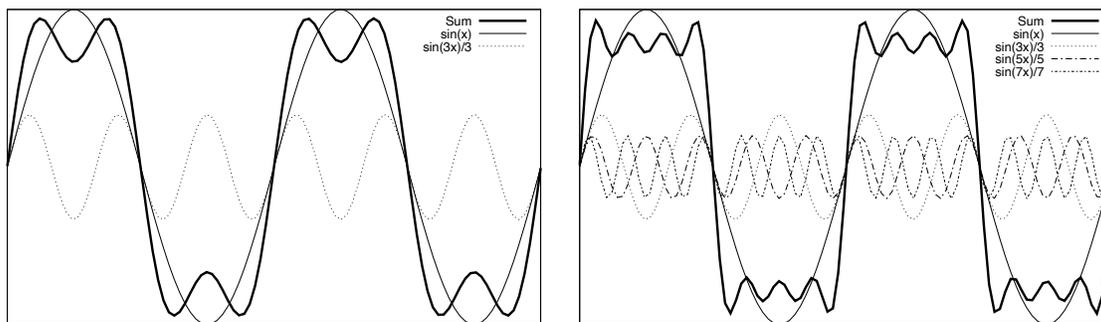


Figure 1: Fourier series approximations for a square wave, showing contributions from the first two and a the first four nonzero harmonics.

In vector space terms,² the sine and cosine functions form a basis for the vector space of real functions, and the coefficients a_k , b_k are the coordinates of the vector f in this basis. For example, in the case of the square wave of Figure 1, the a_k coordinates are $(0, 1, 0, \frac{1}{3}, 0, \frac{1}{5}, 0, \frac{1}{7}, \dots)$ and the b_k coordinates are all zero. To represent an arbitrary function with perfect precision, an infinite number of such coordinates may be needed, so the vector space is infinite dimensional in principle. Infinite dimensions are needed in quantum *mechanics* to represent potential values of the position and momentum operators (Khrennikov, 2010, Ch 2), but in quantum *computing*, the most typical vector space building block is the 2-dimensional complex vector space \mathbb{C}^2 which is used to represent the state of a single qubit (Nielsen & Chuang, 2002, §1.2) (Ying, 2010, §2).

The vector space of piecewise smooth functions has a natural scalar product: for two real functions f and g defined over the interval $[a, b]$, we define their scalar product to be $\int_a^b f(x)g(x) dx$. (This scalar product is often called the inner product and is typically written $\langle f, g \rangle$ in functional analysis.) Crucially for the Fourier theory, for all natural numbers n, m , the inner product $\int_0^{2\pi} \sin(nx)\sin(mx) dx$ evaluates to 0, unless $n = m$ in which case it evaluates to 1. The same result holds for products of cosine functions, and all products of a sine and a cosine function. Thus these sine and cosine functions form not only a basis, but an orthonormal basis for the vector space of smooth real functions. (An orthonormal basis is one where each basic vector has unit length and each pair of distinct basis vectors are orthogonal to one another. Note that in machine learning, ‘one-hot encodings’ form a kind of orthonormal basis (Géron, 2019, Ch 2).

Representing functions using coordinates via Fourier series expansion is only one option. Many other such expansions were later found, for single-variable functions over the real numbers, and also for functions over other spaces such as the sphere, the so-called spherical harmonic functions. The basis of spherical harmonic functions is crucial to the solution of the Schrödinger equation for the hydrogen atom, and in fact each possible energy level for an electron corresponds to a particular spherical harmonic function. Because of this, the number and distribution of spherical harmonic functions accounts directly for the structure of the periodic table of elements in chemistry (Dirac, 1930, §39), but they are not just limited to physical systems. Interestingly, this mathematical formalism has been fruitfully used in the rendering of images, a point that was first recognised by Sloan, Kautz, and Snyder (2002). Spherical harmonics can be seen as a way to encode information over a sphere, in general, beyond their historical origins in solving partial differential equations in mathematical physics.

2. Introductions to vector spaces — basis vectors, coordinates, dimensions, and linear dependence — can be found in texts on linear algebra (Strang, 1993), quantum computing (Nielsen & Chuang, 2002, Ch 2), natural language processing (Widdows, 2004, Ch 5), and in many online tutorials.

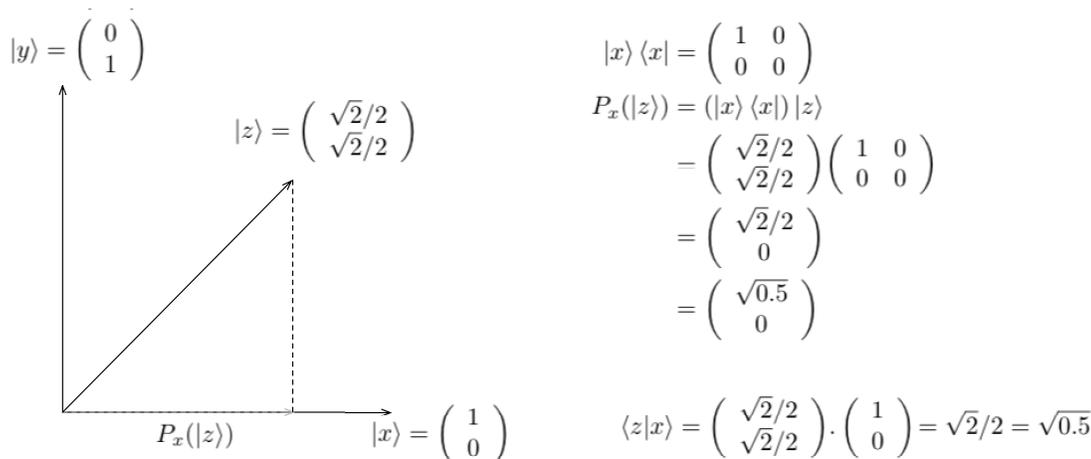


Figure 2: Projection using the outer product of the x axis with itself, $|x\rangle\langle x|$.

By the end of the 19th century this rich mathematical theory was fully mature, before quantum physics was developed. Functions, series, limits, conditions for convergence, and the uses of different bases of functions to represent different physical systems (typically arising from solutions to partial differential equations) were well-understood. Around this time David Hilbert added to the earlier vector spaces of Hermann Grassmann (1862), working specifically with vector spaces with a well-defined scalar product and where every convergent sequence has a limit. Such a space is today called a Hilbert space.

Later, Paul Dirac (1930) introduced the bra-ket notation, designed especially for calculations involving wavefunctions of quantum systems represented as vectors in such a Hilbert space. In Dirac notation, the wavefunction ψ is written as a ‘ket’ vector $|\psi\rangle$, and then the scalar product of $|\phi\rangle$ and $|\psi\rangle$ is written as $\langle\phi|\psi\rangle$. The object $\langle\phi|$ in this expression has its own interpretation as a *covector* or dual vector, called a ‘bra’ vector in Dirac notation: for any vector space V over the field F , the dual space V^* is the space of F -linear mappings from V to F , whereupon it is easy to show that V^* is isomorphic to V , and a particular scalar product (written as a bilinear form $\langle u, v \rangle$) induces one such isomorphism by the identity $u \rightarrow \langle u, _ \rangle$ (Nielsen & Chuang, 2002, §2.1.4). One of the key reasons for using Hilbert space as the setting for quantum mechanics is that the existence of such a scalar product makes it possible to switch between the ket vector $|\psi\rangle$ and the bra covector $\langle\psi|$ whenever this is convenient, and the similarity between the angle bracket notation for bilinear forms and the bra-ket notation makes this elegantly obvious. Introductions to Dirac notation are given by Nielsen and Chuang (2002, Ch 2), Van Rijnsbergen (2004, p. 104), Bernhardt (2019, Ch 2), and the original presentation in Dirac (1930) is still very readable.

Figure 2 shows the projection of a ket vector $|z\rangle$ onto the x -axis whose ket vector is $|x\rangle$. The projection operator is constructed as $|x\rangle\langle x|$, which represents the *outer product* of the vector $|x\rangle$ with itself. For $|x\rangle$, this outer product is a diagonal matrix with zero values except for a 1 in the top leftmost quadrant. Multiplying $|z\rangle$ by this matrix results in its projection onto the x axis, $P_x(|z\rangle)$. Of note, the length of this projected vector is $\sqrt{2}/2$, which is also the scalar product between $|z\rangle$ and $|x\rangle$, or $\langle z|x\rangle$ (as well as the cosine of the angle between these unit length vectors). This projection can be interpreted probabilistically — the length of the projection on $|x\rangle$ gives the *probability amplitude* of observing the superposition z in state x . Squaring this probability amplitude gives the probability itself, in this case 0.5.

It should also be noted that elements of the dual space V^* are linear operators on V and vectors in their own right: and elements of V are also linear operators on V^* . Matrices present another

such example. An $m \times n$ real matrix can be used to represent a linear mapping from \mathbb{R}^n to \mathbb{R}^m using matrix multiplication. At the same time, it is easy to see that that matrix addition adds the coordinates of two $m \times n$ matrices just like vector addition, and that the set of all $m \times n$ is also a vector space — matrices are both vectors and operators on vectors.

Another crucial definition related to dual vectors and linear operators is that of the *adjoint* operator. The adjoint of a linear operator A is the operator A^\dagger such that $\langle A^\dagger(\phi)|\psi\rangle = \langle\phi|A(\psi)\rangle$ for all $|\phi\rangle, |\psi\rangle$. The definition of a Hilbert space guarantees that adjoints exist and are unique, and that $(A^\dagger)^\dagger = A$ (Nielsen & Chuang, 2002, §2.1.6).

2.5 Observation and Uncertainty

Experimental measurements on quantum systems are represented in quantum theory using an *observable*, which is a linear operator A that acts on the vector $|\phi\rangle$ to give a vector $|A\phi\rangle$. If this observation is performed twice in quick succession, the same state is always observed, so $|A\phi\rangle$ is invariant under the action of A . Such a situation occurs mathematically for states represented by eigenvectors of the operator A , which are those vectors for which $A|\psi_k\rangle = \lambda_k|\psi_k\rangle$ (i.e. vectors that do not change their orientation when the operator is applied), where λ_k is the k^{th} eigenvalue of A . The probability of observing the system $|\phi\rangle$ in the eigenstate $|\psi_k\rangle$ is given empirically by $|\langle\phi|\psi_k\rangle|^2$. As well as being probability amplitudes, in geometric terms the scalar products $\{\langle\phi|\psi_i\rangle\}$ are also the coordinates for the vector $|\phi\rangle$ when it is expanded in the basis given by the eigenvectors $\{|\psi_i\rangle\}$. So to understand the way A operates on $|\phi\rangle$, it is enough to write the vector $|\phi\rangle$ as an expansion $\sum \lambda_i|\psi_i\rangle$, where $\lambda_i = \langle\phi|\psi_i\rangle$. Such an expansion works only if the operator A has a suitable spectrum of eigenvectors with real eigenvalues (they must be real if their squares are to be interpreted as probabilities). It follows from linear algebra that A must be self-adjoint, that is, $A = A^\dagger$ or more explicitly, $\langle A\psi|\phi\rangle = \langle\psi|A\phi\rangle$ for all $|\phi\rangle, |\psi\rangle$. In terms of coordinates and matrices, an operator is self-adjoint if and only if it is represented by a matrix that is equal to its conjugate transpose: that is, $a_{ji} = \bar{a}_{ij}$ for all i, j . Such a matrix is called *Hermitian*.

Other results in quantum mechanics can now be derived relatively easily. For example, a *quantum superposition* arises whenever two (or more) state vectors are added together, a scenario that commonly occurs in many vector space AI models. Similarly, the Heisenberg Uncertainty Principle follows from the fact that linear operators do not usually commute with one another (as is well-known from matrix multiplication, AB and BA are not usually the same), and the size of the uncertainty is bounded by the commutator $AB - BA$. This follows directly from the mathematical representation, which as we have seen, follows from the basic requirement that the system be represented in a way that is amenable to linear superposition (hence the use of vectors), and that observations be represented in such a way that performing the same observation twice in succession gives the same answer.

Understanding the relationship between such mathematical implications and behaviors in the physical world is one of the key areas of research in quantum foundations. For example, on the Uncertainty Principle itself, Van Rijsbergen (2004, p. 114) pointed out “It is surprising that such a famous principle in physics is implied by the choice of mathematical representation for state and observable in Hilbert space.” While the philosophical status of this implication may still be debated, the claim that a change of basis leads to different and unpredictable measurements has gone from a paradox in the 1930s, to the design of encryption protocols such as BB84 in the 1980s (Bernhardt, 2019, Ch 3), to commercial offerings of quantum key distribution today. Other properties such as nonlocality and entanglement will be introduced and discussed mathematically in this paper: as described by authors including Coecke and Kissinger (2017), although some of these mathematical consequences were first perceived as awkward difficulties for realistic physics, they have now become cornerstones of how quantum computers work in practice.

3. The Quantum Formulation of Information Retrieval

Information retrieval was pioneering in the use of vector spaces for representing language and information, and the first area related to AI and computational linguistics to be described thoroughly from a quantum theoretical point of view (by Van Rijsbergen (2004)). Because of this head start, information retrieval and related word vector representations are a prominent source of examples in this paper. This section introduces these applications of vectors and quantum theory to information retrieval, explaining the way that cosine similarity can be seen as a projection operator, and how projection operators have been used to model other logical operations including conditionals, negation and disjunction in the logic of subspaces and projections which is still called ‘quantum logic’ following Birkhoff and von Neumann (1936).

The vector space model for information retrieval (summarized by Salton and McGill (1983) and Manning et al. (2008, Ch 6)) stands out as one of the first AI models that was given a distinctly quantum theoretical formulation, developed by Van Rijsbergen (2004) and Widdows (2004, Ch 7). Van Rijsbergen realized that the Hilbert-space formulation of quantum mechanics and the vector space model for information retrieval have much in common, and are in many ways identical. This paved the way for much work in understanding and exploiting commonalities between quantum mechanics and information retrieval. A thorough summary of progress in the subsequent decade is given by Melucci (2015).

Van Rijsbergen (2004) began to join these areas using the simple observation that if we use the Euclidean norm to measure distance, then the vector spaces used in information retrieval are all trivially Hilbert spaces (since they are finite-dimensional, they are also complete). Typically, vector-based information retrieval systems use the cosine measure to give similarities between queries and documents, which gives the same ranking as using Euclidean distance with normalized vectors (as explained in detail by Widdows (2004, §5.5)), so while the strictly metric properties of Euclidean space are often not emphasized when calculating similarity, they are present mathematically. Given a query vector q and a document vector d (both of unit length), the cosine similarity $q \cdot d$ is also the length of the projection of d onto q . Using Dirac’s bra-ket notation, this can easily be written as $\langle q|d\rangle$.

This theory was then developed using the non-commutative behavior of projection operators to model the interaction of a document being relevant to a particular query, and about a particular topic. It can be exploited to model conditional logic and implication for IR in vector spaces (Van Rijsbergen, 2004, Ch 5). Thus there are ways to adapt conditional operators to vector space models. Importantly, much of the work involved can proceed (and historically has proceeded) quite independently of quantum mechanics. Powerful logical operations in Hilbert spaces are available, and whether or not they are inherently ‘quantum’ is partly a historical question. Here, we will argue that they are useful either way, but we will return to this this question in Section 7.6.

3.1 An Intuitive Primer: Quantum Logic as Vector Subspaces and Projections

Recognizing that similarity is implemented as a kind of projection operator takes the parallel between QM and IR further, because the observables in a quantum system are often represented as projection operators in a Hilbert space. The logical structure of such projections was discovered and analyzed in a seminal paper called *The Logic of Quantum Mechanics* (Birkhoff & von Neumann, 1936), and has since been called just ‘quantum logic’, or sometimes the ‘standard logic’ on a Hilbert space (Varadarajan, 1985, Ch 1). Geometrically speaking, the key to understanding the contrast between quantum logic and Boolean logic is to see that Boolean logic is modelled by *subsets*, whereas quantum logic is modelled by *subspaces*, that is, subsets that are themselves vector spaces. For example, in 3-dimensional space, any arbitrary collection of points can be considered a subset, but only the lines and planes are subspaces. Note also that with a scalar product, each subspace P can also be used to define the operation of *orthogonal projection* onto that subspace, and if $|p_1\rangle, \dots, |p_n\rangle$ is an orthonormal basis for P , this projection takes the particularly simple form $\pi_P = \sum_{i=1}^n |p_i\rangle \langle p_i|$

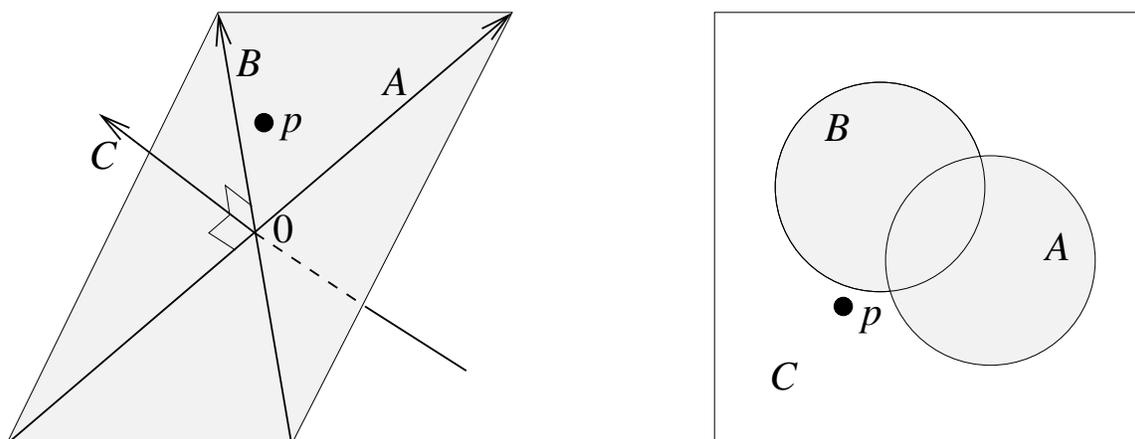


Figure 3: Vector Subspaces contrasted with Venn Diagrams of Sets. On the left, a Euclidean picture of the quantum logic of subspaces, and on the right, a Venn diagram picture of Boolean logic of sets.

(Nielsen & Chuang, 2002, §2.1.6). The ‘similarity’ between a vector $|a\rangle$ and the subspace P can then be defined as the magnitude of $\pi_P(|a\rangle)$, and when P is one-dimensional and $|a\rangle$ has length one, this recovers the familiar definition of cosine similarity between two vectors.

A sketch highlighting some of the differences between Boolean and quantum logic is given in Figure 3. In the realm of quantum logic (depicted left), we can see that if a point p can be written as $\lambda A + \mu B$ for some λ and μ , it is considered part of the disjunction $A \vee B$ even if p is itself in neither A or B . The negation C is not the entire rest of the space, but is strictly the *orthogonal* complement C^\perp . This is a big contrast with Boolean logic — in the familiar Venn diagram on the right hand side of Figure 3, if the point p is in neither A nor B , then by definition it is part of $C = \neg(A \vee B)$. In quantum logic, if $p = \lambda A + \mu B + \nu C$ and none of these coefficients are zero, then we cannot say either that $p \in A \vee B$ or that $p \in C = \neg(A \vee B)$, but instead p has some of each. The conjunction operator behaves the same way in both quantum and classical logic (because the intersection of two subspaces is itself a subspace, for example, the intersection of two planes is typically a line), but disjunction and negation behave very differently. Finally, in the quantum scenario we find that if the set A, B, C forms an orthonormal basis, then the coefficients λ , μ , and ν correspond precisely to the *probability amplitudes* for observing p in one of these basis states in quantum mechanics (Isham, 1995, Ch 2).

3.1.1 SUBSPACE DISJUNCTION

In quantum logic, the disjunction of two one-dimensional subspaces (lines through the origin) is the plane that spans them both. In Figure 3 (left), the example is that the disjunction of the lines OA and OB is the plane OAB , and the notation is usually simplified by leaving out the origin O and saying that the disjunction of the lines A and B is the plane $A \vee B$. Note that this includes points that are not in either of the initial lines (these points correspond to a superposition in quantum mechanics.) In logical terms, we say that quantum logic is *non-distributive*, because it fails to satisfy the distributive law $P \wedge (Q \vee R) = (P \wedge Q) \vee (P \wedge R)$.

More intuitively, the quantum disjunction interpolates (and sometimes extrapolates) from its inputs. A conversational example that contrasts such inclusion might be:

It’s 70 or 80 miles from Oxford to Birmingham and buses leave at 9am or 11am.

Without even thinking about it, an adult reader will assume that the actual distance (about 74 miles) is expressed perfectly well by the phrase “70 or 80 miles”, but will not assume that there is also a 10am bus departure. In the first case, the combination is continuous, whereas in the second case, it is discrete, and while Boolean logic gives a better model for the latter, the former is more effectively modelled by quantum disjunction. (This is just an illustrative example, and several other approaches such as fuzzy logic can be motivated in a similar way.) The notion that an object can be modelled in a space that combines A and B without it being part of either A or part of B is a feature of quantum systems and their non-separability that we will encounter again later (and more strikingly) with tensor product representations and quantum entanglement. With subspaces, this affects the way that categories or natural kinds might be modelled in an IR system, a point discussed more fully by Van Rijsbergen (2004, Ch. 2) and Melucci (2015, §3.3).

3.1.2 SUBSPACE NEGATION

The negation of a subspace in quantum logic is not just its set complement, but its orthogonal complement: for example, for a plane in 3-dimensions, the orthogonal complement is the linear subspace normal to the plane. For example, in Figure 3, the orthogonal complement to the plane $A \vee B$ is the line C that is perpendicular to both A and B and passes through the origin. Negating a concept therefore involves projecting a vector onto this normal subspace, which can move it over large parts of the space, so in this sense quantum negation is a *non-local* operation.

3.1.3 SUBSPACE CONDITIONALS FOLLOWING VAN RIJSBERGEN

The material conditional $A \rightarrow B$ of classical logic is modelled by the inclusion $A \subseteq B$ in set theory, and the direct quantum analogy of this is the subspace inclusion $A \leq B$, which means that A is contained in B and A and B are both subspaces. In terms of the corresponding projection operators, $A \leq B$ if and only if $\pi_B \pi_A = \pi_A$, which is to say that if we have already projected onto A , projecting onto B does nothing.

In information retrieval, however, Van Rijsbergen (2004, Ch 5) is trying to describe the relationship between a query and a relevant document, modelled by the document ‘implying’ the query — in the sense that the relevant document contains *at least* the information in the query and more. Arguing that typically, a query is not *exactly* implied by any document, the notion of relevance-as-implication is qualified by the *Logical Uncertainty Principle* of Van Rijsbergen (1986):

Given any two sentences x and y ; a measure of the uncertainty of $y \rightarrow x$ relative to a given data set, is determined by the minimal extent to which we have to add information to the data set, to establish the truth of $y \rightarrow x$.

Van Rijsbergen (2004, Ch 5) argues that a conditional operator that meets this criterion is (the projection onto) the subspace $A^\perp \vee (A \wedge B)$. This corresponds to the notion that an object supports the claim “birds fly” if it is a bird that flies or if it is not a bird at all. (The fact that pigs don’t fly doesn’t provide positive evidence that birds fly, but at least it does not refute the claim!) Van Rijsbergen’s exploration is logical rather than empirical: the conditional is shown to satisfy logical compatibility properties explored by earlier authors, and in the 2004 work, applications in information retrieval such as relevance feedback and clustering were proposed but not demonstrated. Nonetheless, this work sets the stage for later developments, including the use of error terms in the classification experiments of Bankova, Coecke, Lewis, and Marsden (2019) and Lewis (2019) discussed below in Section 4.3.

3.2 Application 1: Orthogonal Negation

It turns out that these relatively simple observations can be applied directly to vector models for language. For example, suppose a corpus contains several uses of the word *suit* in the sense of

lawsuit, and several in the sense of *garment*. It follows that the vector learned from training for *suit* will involve a superposition of these senses, so

$$suit = suit_{lawsuit} + suit_{garment}. \quad (1)$$

Now, suppose a user is searching for documents about *suit* in the sense of garment. An effective way to meet this need is to project the query vector from the term *suit* onto the subspace orthogonal to *lawsuit*. This technique was explored in detail with various term-retrieval measures by Widdows and Peters (2003), which demonstrated that the immediate advantage of projection over Boolean negation in IR experiments was that it removed not only documents containing the unwanted term *lawsuit*, but other documents that use terms related to *lawsuit*. This form of vector negation reduced the occurrence of synonyms and neighbours of the negated terms by as much as 76% compared with standard Boolean methods — a practical outcome of the non-local nature of the negation operator. Its effectiveness in improving document retrieval was demonstrated by Basile, Caputo, and Semeraro (2011), with relative increases of up to 10.8% in Mean Average Precision (MAP) over a standard baseline obtained when using a re-ranking strategy involving rendering the vector average of 5 relevant documents orthogonal to that of a single irrelevant counterpart, and much greater increases when incorporating more relevance judgments (e.g. relative increase of 66.46% with feedback from 40 relevant and 40 irrelevant documents). The application of a projection operator to remove or emphasize aspects of a mixed semantic vector representation follows naturally from the fact that such vectors are generated as superpositions of vectors that represent occurrence in a (usually sense-specific) context. For a more thorough exposition of these points, see Widdows (2004, Ch 7) and Melucci (2015).

Although negation is something of a niche operation in search engines, this example helps us to start motivating an answer to the question: “What can quantum models do for artificial intelligence?” In the case of modelling negation, quantum logic from the 1930’s has provided a direct and computationally simple technique for performing logical operations in a vector model search engine. Building more thoroughly on the same foundation, Garg, Ikbal, Srivastava, Vishwakarma, Karanam, and Subramaniam (2019) have recently used the rules of quantum logic (including negation) to define objective functions for optimizing the representation of ontologies and relationships between concepts as vector space embeddings. So from this point of view, quantum theory has motivated and guided these successes. However, it is partly a historical accident that the logic of projections and subspaces is called quantum logic at all: the meet, join, and orthogonal complement operations on a vector space were all discussed back in the 1860’s by Grassmann (1862), and the logical structure could easily have been explored independently of its application to quantum mechanics by Birkhoff and von Neumann in the 1930’s. So perhaps the right question is not “Can quantum logic be applied to IR?” (it can), but “Would it be better to refer to ‘vector logic’ rather than ‘quantum logic’ throughout the literature?”

The interpretation of products between vectors as subspaces is also a key part of *geometric algebra*, which can be used a fundamental model for modelling several physical and computational operators (Dorst, Fontijne, & Mann, 2010). One further complication here is that there are other logical systems with vectors, including those of Mizraji et al. (1994), Westphal and Hardy (2005), and Clarke (2012). The question of which vector logic works most effectively for information retrieval or other word embedding systems and applications is a research area that is likely to yield many interesting new results, but which has not to date attracted as much attention as it deserves.

3.3 Application 2: Modeling Term Dependencies for Information Retrieval

Another example of the utility of quantum models for information retrieval was provided by Sordani, Nie, and Bengio (2013), who investigated whether an extension of the standard unigram / bag-of-words language model to accommodate dependencies between terms in the context of frequently occurring multiword expressions such as *climate change* can improve retrieval performance. The

issue of term dependencies is a longstanding problem in information retrieval. When constructing a probabilistic or geometric representation of a document, the question arises as to how to address the compositional nature of frequently occurring short phrases. A simple approach would involve treating the phrase *climate change* as a unit — as though it were a single word. However, information about the occurrence of the component terms *climate* and *change* would then be lost. Furthermore, this approach provides no clarity about how best to weight such composite terms, which occur with less frequency than their unigram components.

To address these issues, Sordoni and colleagues developed a Quantum Language Model (QLM), a generalization of prior language modeling approaches that provides the means to model term dependencies without severing the connection between the probability of observing a multiword expression and the probabilities of observing its component terms. Here, we will provide a brief account of the model and the accompanying empirical results. We refer the interested reader to Sordoni et al. (2013) for further implementation details such as the selection of phrases to model, smoothing and the estimation of model parameters.

Language modeling approaches to information retrieval can be represented geometrically by mapping each term in the vocabulary to an independent basis vector (familiar to readers with a machine learning background as a ‘one hot encoding’ (Géron, 2019, Ch 2)). Projection operators for each term can then be defined to recover the probability of observing a particular term in the context of a set of language model parameters learned from a document. Projection operators for single words are diagonal matrices with a single non-zero value corresponding to the basis vector of the term concerned.

With the QLM, multiword expressions are represented as a weighted superposition of the basis vectors corresponding to their component terms:

$$\overrightarrow{\text{climate_change}} = \alpha \times \overrightarrow{\text{climate}} + \beta \times \overrightarrow{\text{change}}$$

which are normalized to length one. The projections onto these vectors lead to off-diagonal entries in the corresponding matrices, so the term dependency cannot be represented as just a weighted sum of the ‘one-hot’ diagonal elements. The coefficients chosen for this weighting encode the extent to which observation of the composite phrase indicates the presence of each component term. This is intuitively appealing, and provides the means to heuristically weight the relative importance of the terms within a phrase such that, for example, documents containing the term ‘climate’ in isolation are more likely to be retrieved in response to the query ‘climate change’ than those containing the term ‘change’ alone. Given a unit of text, model parameters are learned using an approximation algorithm that attempts to find the parameters that maximize the probability of the observed document or query. Relevance ranking is performed on the basis of the divergence between parameters learned from queries and documents. Evaluation of the QLM was conducted using selection of 450 queries drawn from across four information retrieval evaluation sets. The best performance for each of the four sets was obtained by a variant of the QLM, outperforming a unigram language model baseline, with statistically significant improvements over then state-of-the-art approaches using Markov Random Fields (MRF) (Metzler & Croft, 2005) for the two larger web-based sets with relative increases in MAP for the best-performing QLM as compared with the best-performing MRF of 5.5% and 5.2%.

In summary, using quantum theoretic ideas directly in the traditional vector model for information retrieval has provided several concrete opportunities, including an account of conditionals and implications, negation and disjunction, and term dependencies, some of which have improved performance in retrieval experiments.

4. Categories, Hypernyms and Implication: More Advanced Structures in Vector Spaces

Having presented examples of how vector spaces can provide a rich set of tools for problems relevant to AI and ML, it is now time to move onto more advanced vector space structures — what insights do they provide? Representing categories and hypernym relationships will help us to motivate the use of structures in vector spaces that are more complex than the individual points referenced by vectors that we explored in the previous section.

4.1 Negation and ‘Parts’ of Vectors

The example above that uses negation to remove unwanted meanings of ambiguous words carries an important conceptual lesson: a single vector can represent multiple meanings in a high dimensional space, and specific ones can sometimes be recovered explicitly using a method such as orthogonal projection to a subspace.

This runs counter to a common belief expressed in the literature, that one point can represent only one thing. As an example of this line of thinking, consider Neelakantan, Shankar, Passos, and McCallum (2014) who start their widely cited paper with this statement:

There is rising interest in vector-space word embeddings and their use in NLP, especially given recent methods for their fast estimation at very large scale. Nearly all this work, however, assumes a single vector per word type — ignoring polysemy and thus jeopardizing their usefulness for downstream tasks.

The effectiveness of vector negation for uncovering different word senses in a single vector demonstrates a flaw in the claim that a single vector per word ignores polysemy, even though the clustering results described by Neelakantan et al. (2014) are still valuable. Mathematically, this claim would be equivalent to the claim that a single vector per particle cannot represent a superposition of pure states, a statement which quantum mechanics has shown to be flawed in the quantum realm. Nonetheless, other authors have echoed this claim, including for example Camacho-Collados and Pilehvar (2018):

The prevailing objective of representing each word type as a single point in the semantic space has a major limitation: it ignores the fact that words can have multiple meanings and conflates all these meanings into a single representation.

The assumption that a vector is equivalent to a point and therefore cannot represent something with internal structure or ingredients goes back to the beginning: “A point is that which has no parts” is Book I, Definition I of Euclid’s *Elements*. However, even in Euclidean space, a vector can also be thought of as an arrow or straight path from the origin to a particular point. This straight path certainly can have ‘parts’; in particular, each coordinate represents a length along a particular axis (in Fourier analysis, a basic harmonic function; in quantum mechanics, a basis state, etc.).

The notion that vectors have parts is much more obvious in quantum theory, because as introduced in Section 2.3, the vectors in quantum mechanics are functions: and of course, a function can have many different parts. Fourier analysis itself is a particularly canonical way of breaking a function down into a sum of different parts. From this point of view, the mistaken belief that a vector can only represent one thing would be unlikely ever to have arisen. Vectors can represent many types of things in many ways, and high-dimensional vectors in particular can accurately represent many ingredients in such a way that the ingredients can be clearly recognized and sometimes recovered from the combined representation. The mathematical foundations behind this claim are analyzed more thoroughly by Kanerva (2009) and Widdows and Cohen (2015).

Despite its mistaken foundations, work motivated by this understanding of vector models has provided useful results. The perspective of Neelakantan et al. (2014) that we criticised above still

created word vectors by clustering contexts of word occurrences (as pioneered by Schütze (1998)) and built a representation for each sense. Similarly, the survey of disambiguation techniques by Camacho-Collados and Pilehvar (2018), and the smoothed Gaussian representations of Vilnis and McCallum (2015) have both provided valuable contributions to the field. Still, the stated mathematical motivation for such work is not well-founded, and a quantum perspective on the mathematics naturally avoids this mistake, and provides a richer formulation that can potentially be leveraged to achieve novel results.

4.2 Subspaces, Disjunctions, and Generalization

Even though a single vector can represent several senses of a polysemous word, it is important to realise that individual vectors cannot mathematically represent everything that can be expressed in a vector space. The case of linear subspaces, described in Section 3, is an obvious example. A k -dimensional subspace in an n -dimensional space typically requires at least $\min(k, n - k)$ vectors to describe it. The $n - k$ representation itself provides an interesting insight: it is useful when $k > \frac{n}{2}$, which makes it more convenient to express a subspace in terms of a normal projection to the subspace. This technique is most commonly encountered in 3 dimensions where the normal to a plane is just a line (for example see again Figure 3, where the normal C could also be used to *define* the plane $A \vee B$, because $p \in A \vee B$ if and only if $p \cdot C = 0$).

The link between disjunctions and categories or other ‘natural kinds’ is that a more general category like *mammal* arises as a disjunction of examples like *cat*, *dog*, *mouse*, *elephant* (Van Rijsbergen, 2004, Ch 3). In lattice theory, a disjunction is a ‘least common ancestor’, characterized by being the most specific element available that subsumes (implies) all of the constituents (Widdows, 2004, Ch 8). Since subspaces are the natural representation of disjunctions in quantum logic, it is tempting to assume they are a good representation for categories of word vectors. This intuition has worked in some cases: for example, a combined representation built using quantum subspace disjunction with limited numbers of inputs has been shown to perform well at the task of geometrically-mediated analogical inference, generally recovering more therapeutic relationships between drugs and types of cancer than comparable superposition-based approaches (Cohen, Widdows, Vine, Schvaneveldt, & Rindflesch, 2012) and in information retrieval, subspaces have been shown to be particularly effective for modelling *negated* disjunctions — that is, removing many unwanted areas of meaning from a query vector (Widdows, 2003). However, a mathematical problem arises with positive disjunctions of many inputs: linear sums of many similar vectors tend to extrapolate and overgeneralize (Bruza, Widdows, & Woods, 2009). Specifically, a disjunction of k non-degenerate vectors *always* leads to a k -dimensional subspace, however similar or different the initial vectors are. For example, back in Figure 3, A and B would generate the exactly same plane $A \vee B$ whether they are close together or far apart, so long as they are not identical. This also means that in an n dimensional space, n very slightly different examples would nearly always generate the whole space. So while subspaces have been shown to work well for some forms of semantic generalization, we should expect this to be an incomplete and eventually over-general model.

4.3 Density Matrices, Positive Operators, and Hyponymy

Quantum mechanics is already familiar with the problem that individual vectors and even subspaces cannot represent everything we encounter in practice. A single particle in a superposition of pure states can be represented as a single vector, but an ensemble of particles cannot. Instead, ensembles are modelled using a *density matrix* or *density operator*, which takes the form

$$\rho = \sum_j P_j |\Psi_j\rangle \langle \Psi_j|.$$

In quantum mechanics, a system of particles that can be expressed as a density matrix but not a state vector is called a *mixed state*. The motivation for generalizing from states to density matrices is discussed in detail by Isham (1995, Ch 6), and Van Rijsbergen (2004, Ch 6). Mathematically, the formalism is general enough to represent all the relevant probability distributions, which an individual state vector $|\Psi\rangle$ with corresponding density operator $\rho = |\Psi\rangle\langle\Psi|$ cannot.

Density matrices are positive operators, in the sense that $\langle\Psi|\rho|\Psi\rangle \geq 0$ for all vectors Ψ . This can be used to induce an ordering on density operators, defining $A \sqsubseteq B \iff B - A$ is positive. This is called the Löwner ordering, and when the operators are restricted to projection onto subspaces, the Löwner ordering becomes equivalent to the quantum logic introduced in Section 3.1. The Löwner ordering is used by Bankova et al. (2019) and Lewis (2019) to represent graded hyponymy (where, for example, *dog* may be said to be a strict hyponym of *mammal* but a graded hyponym of *pet*, because not all dogs are pets). In cases where A and B are incomparable in this ordering, the systems proposed in these papers work by finding positive operators D and E (an error term) such that $A + D = B + E$. This avoids the problem with subspaces whereby any part of A that is not part of B adds whole new dimensions: instead, if A is nearly subsumed by B , the error term E will be correspondingly small. The smaller E , the more strict the hyponymic relationship. In further work, Lewis (2020) extends this system to produce a graded form of negation, which can be applied to more situations than the orthogonal negation of Widdows (2003).

This example is critical, and illustrates the main point of this section: while vectors are powerful representations and can represent many ingredients at once, they cannot represent (for example) all the probability distributions necessary for quantum mechanics. Exploring the ways these shortcomings have been addressed within quantum theory suggests potentially fruitful research avenues for those using vectors to model concepts in language.

5. Products of Vectors and Semantic Composition

The summary of categorization and negation in the previous section highlights one of the key places where quantum techniques are useful in AI. They provide us with a range of operators for exploring, manipulating, and generating representations of concepts using semantic vector models. Some of these are familiar, some are novel, and all open up potential new avenues for research in AI. In this section we draw attention to the problem of semantic composition. There are many ways to compose vectors — so far in this paper we have discussed the vector sum, and various operators related to subspaces and projections. More complicated structures are available, and semantic composition is often represented in vector models using the tensor product and its offshoots.

5.1 Tensor Products and Entanglement

The tensor product is one of the most significant methods used for composing vectors. After exploring addition and subtraction, the next product structure on vectors defined by Grassmann (1862, Ch 2) is the forerunner of today’s tensor product in coordinate form: if $a = \sum a_r e_r$ and $b = \sum b_s e_s$, then their product $[ab]$ is defined as $\sum a_r b_s [e_r e_s]$. By varying the rules for interpreting and identifying the basic terms $[e_r e_s]$, Grassmann showed that such product structures can be used to represent the inner (scalar) product and combinatorial (exterior) product. Note that the order matters here: unlike the sum of two vectors or the product of real or complex numbers, this product between two vectors is not commutative. This makes it more complicated, but also opens opportunities. For example, generating a document vector as a weighted sum of term-vectors is surprisingly effective for information retrieval, but a commutative sum that fails to take word-order into account is unsuitable for building a number of important applications (e.g. a conversational agent or any other dialogue system).

We have already noted that vectors can represent operators as well as states — starting with the (co)vectors in the dual space V^* that act as linear maps from V to its ground field (the *ground*

field being the number system used for coordinates, most commonly the real or complex numbers). Given this, Grassmann's product operator can be used to represent a bilinear map from $V \times V$ to the ground field — for $\alpha, \beta \in V^*$ and $u, v \in V$, we define $[\alpha\beta](u \times v) = \alpha(u)\beta(v)$, and it is easy to see that this map is linear in all of its arguments. In Dirac notation, a product $\langle \phi_i \phi_j |$ of the bra vectors $\langle \phi_1 |$ and $\langle \phi_2 |$ would map the product of two ket vectors $|\psi_i \psi_j\rangle$ to $\langle \phi_1 | \psi_i\rangle \langle \phi_2 | \psi_j\rangle$.

Today the space of all possible linear combinations of products of basis vectors of two Hilbert spaces is called their *tensor product*, written using the symbol \otimes , so that the tensor product of vector spaces U and V is written $U \otimes V$. It is defined more formally as an equivalence class of mappings whereby (for example) if $E(A, B) \rightarrow U$ is a bilinear map from (the Cartesian product of) A and B to U , then this corresponds to a unique linear mapping from $A \otimes B$ to U , and it turns out that all bilinear maps from $A \times B$ to U can be represented in this manner. Linear mappings from one space to another have a similar correspondence: the space $L(A, B)$ of linear transformations from A to B is naturally isomorphic to $A^* \otimes B$ (Lang, 2002, Ch. 16).

In a given coordinate basis, the tensor product of two vectors can be represented as a matrix with the same coordinates as their *outer product*. Just as the inner product or scalar product of two column vectors u, v can be written as $u^T v$, their outer product can be written uv^T . An example with vectors in \mathbb{R}^3 is the following:

$$u = \begin{pmatrix} 1 \\ 0 \\ -2 \end{pmatrix} \quad v = \begin{pmatrix} 2 \\ -1 \\ 3 \end{pmatrix} \quad u^T v = 2 + 0 - 6 = -4 \quad uv^T = \begin{pmatrix} 2 & -1 & 3 \\ 0 & 0 & 0 \\ -4 & 2 & -6 \end{pmatrix}$$

While this is intuitive and familiar, the mental identification of tensors with matrices does lead to gaps. For example, the notation $u^T v$ for the scalar product in coordinate matrix form encourages us to think of row vectors as dual to column vectors, but matrix multiplication alone does not enable us to map a $(n, 1)$ column vector to a $(1, n)$ column vector. (Remember that matrices of size (p, q) can only be multiplied on the right by matrices of size (q, r) , resulting in a matrix of size (p, r) : so a matrix of size $(n, 1)$ can only be multiplied on the right to give a matrix of size $(1, n)$ if $n = 1$, with a similar argument holding for left-multiplication.) So if we were to identify tensors with such matrix representations, we would be tricked into believing that tensors cannot be used to map vectors to covectors, which is not the case.

More generally, tensors form an algebra of multilinear maps. If V is a vector space and V^* is its dual, then the tensor space $V \otimes \dots \otimes V \otimes V^* \otimes \dots \otimes V^*$ consists of multilinear functions from $V^* \times \dots \times V^* \times V \times \dots \times V$ to the ground field. (The $*$ -symbols are deliberately switched between the first and the second expression: the elements of V^* act on elements of V , and vice versa.) Such a product with p copies of V and q copies of V^* is called a *tensor of type (p, q)* , and the number $p + q$ is sometimes called the *rank* or the *arity* of the tensor. For example, the density matrices of Section 4.3, being the product of a bra and a ket vector, are tensors of type $(1, 1)$ and rank 2.

Tensor algebra has been used extensively in many fields before computer science. For example, it has been in use in differential geometry since the early 1900's, partly because so many geometric concepts can be expressed easily using tensors, including vector fields, differential forms, metrics, volume integrals, complex structures and Hamiltonian dynamics. Because of this, some of the most readable and thorough introductions to tensor algebra are from standard texts in differential geometry (for example, Willmore (1959, Ch 5, 6), Bishop and Goldberg (1968, Ch 2)). Continuum mechanics, which models the behavior of materials and liquids as a continuous mass, also makes extensive use of tensor algebra to extend Hooke's law to high dimensions in the modelling of liquids, elasticity of materials, and other important phenomena (via the stress-strain tensor — see Spencer (2004), which includes a good introduction to matrix and tensor algebra for mechanics). Finally, as we have seen, in quantum mechanics the tensor product naturally arises when we move towards representing composite systems. More recently, in machine learning, tensors have made a sustained contribution through the introduction of systems such as TensorFlow for training neural networks

(Géron, 2019, Ch 12). Tensor products are also used to describe the very general notion of ‘processes happening simultaneously’ in physics, sometimes generalizing the notion of tensor product to mathematical settings beyond linear algebra (Coecke & Kissinger, 2017; Heunen & Vicary, 2019).

In summary, tensors have become an invaluable part of practical mathematics in many fields. However, this very multidisciplinary nature can result in confusion around terminology. For example, a matrix can be used to represent a rank-2 tensor, but then the definition of the term *rank* is different from the traditional definition in linear algebra, where the rank of a matrix is the number of linearly independent rows and columns (a key theorem in linear algebra guarantees that the row rank and the column rank are the same, so it makes sense just to talk about the rank of a matrix (Jänich, 1994, Ch 5)). This notion of rank also generalizes to tensor algebra, and the rank of a tensor is also used to mean the number of linearly-independent generators for the tensor. In some fields (including machine learning), the difference between vector spaces and their duals does not yet play nearly as important a role as it does in differential geometry, and so this difference is often glossed over, and instead of talking about tensors of type (p, q) , only the tensor rank $r = p + q$ is discussed. In particular, it is often said that the rank-2 tensors are equivalent to matrices, which we have seen is an algebraically incomplete representation (in the example above, a linear mapping from vectors to covectors can be represented as a rank-2 tensor of type $(1, 1)$, but no matrix can be found that maps column vectors to row vectors using multiplication). This teaching may change gradually, because used carefully, the distinction between vectors and covectors can be a benefit in machine learning as well. For example, Turney (2012) used precisely this approach to make a useful distinction between domain and functional similarities in distributional semantics.

5.2 Tensor Products and Entanglement in Quantum Mechanics

In quantum mechanics, superposition becomes incomplete for describing the state of systems as soon as we consider composites involving more than one particle. Consider, for example, particles with wavefunctions $|\phi\rangle$ and $|\psi\rangle$, each represented as a superposition $\sum_1^m a_i |\phi_i\rangle$ and $\sum_1^n b_j |\psi_j\rangle$, where $|\phi_i\rangle$ and $|\psi_j\rangle$ are the eigenstates for some observables A and B . In cases where the particles can be observed simultaneously, the result of measuring A and B together can be any combination of $|\phi_i\rangle, |\psi_j\rangle$, and there are mn such possibilities. Therefore the combined system of $|\phi_i\rangle$ and $|\psi_j\rangle$ can consist of any linear combination of these eigenstates, so it follows that the Hilbert space of possible states for this system has mn dimensions. This deceptively natural conclusion should be contrasted with the classical situation, where the state space for the combined system would be the Cartesian product of the state spaces for the individual systems, having dimension $m + n$. For a step-by-step example with spin states of 2- and 3-dimensional particles, see Isham (1995, §8.4). The combined state where measuring A results in $|\phi_i\rangle$ and B results in $|\psi_j\rangle$ is written as $|\phi_i\psi_j\rangle$, and the state that represents the combination of the superpositions $\sum a_i |\phi_i\rangle$ and $\sum b_j |\psi_j\rangle$ is $\sum a_i b_j |\phi_i\psi_j\rangle$, which is exactly the same as Grassmann’s definition of the product of two vectors above, just rewritten in Dirac notation.

The mathematical behavior of tensor products is directly responsible for the famous phenomenon of *quantum entanglement*. Entanglement is perhaps easiest to introduce via its opposite phenomenon; separability. A composite state is considered separable if it is possible to write it as a tensor product of two distinct states (i.e. as $|\phi\rangle \otimes |\psi\rangle$ for some $|\phi\rangle$ and $|\psi\rangle$), in which case it is generally considered to be non-interacting. In contrast, an entangled state cannot be decomposed in this manner.

How does this come about? Mathematically, the answer is most obvious if we consider the dimensions involved. If U is of dimension m and V of dimension n , their tensor product $U \otimes V$ has dimension mn , which of course is typically greater than $m + n$. Immediately it follows that not every element of $U \otimes V$ can be written as some individual product $u \otimes v$, because there are only $m + n$ degrees of freedom for choosing the u and v . In order to generate any element of $U \otimes V$, we may need to superpose many different individual products of the form $u \otimes v$, just as in the example system above where the product state must be written as a linear combination of states like $|\phi_i\rangle \otimes |\psi_j\rangle$.

A typical way of explaining the difference algebraically is that $u \otimes v$ is not a standard *element* of $U \otimes V$, but it is a standard *generator* or basis vector for $U \otimes V$. In matrix algebra, this corresponds to the result that any rank-1 matrix (in the sense of matrix rank above!) can be written as the outer product of a single row vector and a single column vector.

With superposition and quantum disjunction, we saw earlier that a quantum system might be represented as the sum of its parts without being identical with or contained in any of those parts. The tensor products allows for even richer combinations where the product is not contained in any of the ingredients, and cannot even be broken down into a combination of one simple ingredient from each part. A vector u in the subspace $A + B$ can at least be represented as the sum $a + b$ for some $a \in A, b \in B$, and this decomposition is easily obtained using projection operators. On the other hand, a tensor $u \in A \otimes B$ cannot even be decomposed like this: at its most general, it must be written as a linear combination of the form $u = \sum a_i \otimes b_i$ (and the minimum possible i is the tensor generalization of the rank of a matrix). While quantum logic can be described as non-distributive and non-local, tensor products and entanglement demonstrate *non-separability* (Coecke & Kissinger, 2017, §4.1.1).

The concept of entangled quantum systems lies at the center of some of the most interpretationally fraught aspects of quantum mechanics (Isham, 1995; Greenstein & Zajonc, 1997). For example, if a system of two particles is entangled, but then one of its constituent particles is measured and thus observed to be in an eigenstate, what does this mean for the combined system and for the other particle? While physical implications of entanglement are conceptually challenging, the concept is well-enshrined in modern physics: creating and observing entanglement between qubits (quantum bits) is one of the key necessary and challenging processes in quantum computing, a point that we will return to in Section 6. First however, it is worth highlighting some of the ways in which this theoretical phenomenon has been exploited in language models in AI.

5.3 Explicit Composition with Vectors and Tensors in AI

The use of tensor products in AI is often traced to the work of Smolensky (1990), and its uses have become much more widespread more recently. During the later decades of the 20th century, vector models were often described as “bag of words” models because of the commutative nature of the vector sum operation they applied. This choice of operator meant that the same representation could be used for a collection of words in a document, irrespective of the order in which the words appear. By contrast, discrete logical models used in formal semantics have for many decades been quite explicit about the ways words should be combined, but were often notably silent about what those words mean in themselves (see Widdows (2008), Baroni, Bernardi, Zamparelli, et al. (2014) for surveys of this methodological difference between traditions). This history of two modelling frameworks led to an unnecessary and unfortunate gap: there are many interesting product operations such as the tensor product between vectors that are well-established in linear algebra, but for years there were relatively little awareness of these alternatives in language research. Since the early 2000’s, this situation has changed markedly.

In a pioneering case for AI, tensor products were introduced at least as early as the work of Smolensky (1990), where the tensor product was used to bind a variable to its value, and the term ‘binding’ for products of vectors has been used in various language and cognitive models involving vectors since the work of Kanerva (2009). In the first decade of the current century, the use of tensor algebra for combining word vectors was explored using a quantum formalism by Aerts and Czachor (2004), whereby a sentence (w_1, \dots, w_n) would be represented as the tensor product $w_1 \otimes \dots \otimes w_n$. Taking this in explicitly linguistic directions, product operations on vectors were introduced to model grammatical and semantic composition (Clark & Pulman, 2007; Widdows, 2008). Particularly interesting examples included the work of Mitchell and Lapata (2010) on noun composition, and Baroni and Zamparelli (2010) on adjective-noun combinations. Turney (2012) used a combination of two vector spaces (a large-window space capturing domain similarities, and

a narrow-window function space capturing functional similarities) to combine relational similarity with two-word compositional behavior.

Inspired by Smolensky’s work on binding operators, the model of Clark and Pulman (2007) used the tensor product as a means to generate vector representations of phrases, noting that one desirable property of the tensor product in this context is that it does not commute, so the product $Djokovic \otimes beat \otimes Murray$ is not equal to the product $Murray \otimes beat \otimes Djokovic$. Another useful mathematical property is that the inner product of two tensors is equal to the product of the inner products of their related constituents, e.g. $(beat \otimes Murray) \cdot (defeated \otimes Andy) = (beat \cdot defeated) \times (Andy \cdot Murray)$. This means that the similarity between a pair of phrase tensors can be calculated without the need to explicitly represent them (which would require n^2 dimensions for n dimensional vectors), and, with normalized vectors, this is the product of the cosine similarities between the component vectors for terms that occupy the identical position in the sentence structure. However, the grammatical structure of related sentences may not be identical, and as such their composite representations cannot be easily compared with this method.

A solution to overcome this limitation of the tensor product approach was proposed by Coecke, Sadrzadeh, and Clark (2010). They made use of category theory (Lawvere & Schanuel, 2009) to demonstrate that vector spaces and certain types of grammar (exemplified in this work by pre-group grammars (Lambek, 2001)) fall under the same category type, compact closed categories (see Appendix 8. This realization provided a means to map between the grammatical and the vector representation, and thus demonstrated that the rules for grammatical composition can be applied to a compositional model in vector space. The mathematical roots of this work are directly related to quantum mechanics due to the work of Abramsky and Coecke (2009), which developed *Categorical Quantum Mechanics*, an axiomatic presentation of quantum theory using compact closed categories, which makes the relationship with the grammatical work of Lambek particularly clear. (An online appendix is provided for readers unfamiliar with category theory.) This point will be revisited in Section 6 because it is especially relevant to implementation on quantum computers. Since its introduction by Coecke et al. (2010), the Distributional Compositional Categorical model has become known as DisCoCat.

Grefenstette and Sadrzadeh (2011) provided an implementation and evaluation of this approach, deriving word representations from the British National Corpus. To do so, different rules of composition were assigned to words with different grammatical types in accordance with the categorical type of their pregroup, with ‘atomic’ types such as nouns represented by distributional term-by-context-term vectors, and ‘adjoint’ types generated compositionally. For example, verb representations were generated as the sum of the tensor products of the noun vectors corresponding to their subject and object across every occurrence in the corpus (intransitive verbs and adjectives can also be accommodated). Vector representations of sentences can then be compared: for example, the vector representation of the sentence “Djokovic beat Murray” is composed by pointwise multiplication of the tensor representing the verb “beat” and the tensor composed from the vectors for “Djokovic” and “Murray”: $(Djokovic \otimes Murray) \odot beat$. The model was evaluated for its ability to estimate the similarities between short phrases, and correlated with human judgment of similarity as well as, and better than, the best available models on datasets concerning verb-noun and subject-verb-object relationships. The authors report Spearman correlations of 0.17 and 0.21 for the verb-noun and subject-verb-object sets respectively with the compositional model, as compared with 0.17 for both sets with the multiplicative model of Mitchell and Lapata (2008).

In methodologically similar work, words such as verbs and adjectives that take arguments have been represented as tensor products in matrix form. For example, Baroni and Zamparelli (2010) used this approach to model the action of adjectives upon nouns, and Socher, Huval, Manning, and Ng (2012) took it to the logical destination of representing each internal node in a parse tree as a matrix operator acting upon its input arguments. This area has become known as Compositional Distributional Semantics, summarized by Baroni et al. (2014), and work in this area has continued, an example being the work of Sadrzadeh, Kartsaklis, and Balkır (2018) on sentence entailment in

this framework, where an F1-score of 0.86 is reported on recognizing semantic entailments like “robin flies \models bird moves”.

As with vector logics, the tensor product as used in quantum mechanics is only one of the composition operations that can be used with vectors, albeit an important one. One implication of using the tensor product is the inevitable explosion in dimensionality that arises: a tensor product takes two n -dimensional vectors and makes an n^2 -dimensional tensor, which is a problem for scalability on classical hardware, and creates products that are not obviously comparable with their inputs. Various algebraic structures, sometimes known as Vector Symbolic Architectures (Levy & Gayler, 2008), have been used to address this problem. They make use of circular correlation and circular convolution, to roll the tensor product coordinates back into coordinates in the initial space. The use of these structures for reasoning in continuous models has been described by Widdows and Cohen (2015).

Entanglement has also been used in compositional modelling for language concepts, motivated partly by the study of entanglement for word-association in cognitive models (Bruza, Kitto, & McEvoy, 2008; Aerts & Sozzo, 2011). For example, Kartsaklis and Sadrzadeh (2014) explore the entangled representation of transitive verbs using tensors in the categorical framework of Coecke et al. (2010). In an explicit use of entangled superpositions for reasoning and inference, predication-based semantic analysis was developed by Cohen et al. (2012), and makes use of vector binding operations to represent a concept A occurring in a relation of the form $R(A, B)$ as the sum of the products $r \otimes b$ for each relation in which it occurs. (For example, with the relational triple ‘aspirin TREATS headache’, the concept vector for *aspirin* gets incremented with the bound product vector $TREATS \otimes headache$.) In this framework, concepts become represented as superpositions of products of pairs, and these cannot be expressed as any product of single ingredients r and b . This uses some of the same mathematics that in quantum mechanics leads to entanglement, including the use of complex vectors as a ground field in some implementations (Cohen et al., 2012; Cohen & Widdows, 2015). (In an alternative approach to modelling relations using complex numbers, the work of Garg et al. (2019) explicitly uses the complexified representation $A + iB$ to form pairs, and axioms of quantum logic to formulate the compositional representation of relations.)

From the point of view of language and relationships, the notion of a general relation being derived from several varied examples makes intuitive sense. Each of the pairs (*Leto, Artemis*), (*Henry VIII, Elizabeth I*), (*Lord Byron, Ada Lovelace*), (*Darth Vader, Princess Leia*), (*Debbie Reynolds, Carrie Fisher*), is an example of a parent-child relationship, and all of these could be combined into a thorough and very varied notion of parenthood. Because of this variety, it would be surprising if there was any one ‘prototypical parent’ and ‘prototypical child’ that combines to make the relation of ‘prototypical parenthood’. Thankfully, the mathematical structures used for tensor products that give rise to entanglement enable us to represent a relation that is a rich combination of different example pairs, even for relations where no single pair of prototypical ingredients exists (Widdows & Cohen, 2015, §5).

5.4 Implicit Composition in Deep Neural Networks

Something the methods described in the previous section have in common is that they encode some *explicit* syntactic structure: a role / value binding, or a relationship in a grammatical parse tree. This raises the question: can models go beyond the bag-of-words drawbacks and encode more order-dependent language structures without using this traditional syntactic machinery? A recent and comprehensive survey of this topic is provided by Hupkes, Dankers, Mul, and Bruni (2020).

During the same years that compositional distributional semantics has been developed, neural networks have made great strides in artificial intelligence, particular the use of networks with several intervening layers, hence the term ‘deep learning’ (LeCun, Bengio, & Hinton, 2015). In some cases, work on deep learning and compositional semantics has been explicitly combined: for example, Socher et al. (2012) describe Recursive Neural Networks used for training. These should not be

confused with the the simpler and more standard Recurrent Neural Networks (RNNs), in which the output of a single neuron depends on its inputs and its previous state. RNNs have been used for many sequence-modelling tasks, as have their more sophisticated cousins, LSTMs (Long Short-Term Memory cells, where prior state may be stored for longer and updated based on learned importance) (Géron, 2019, Ch. 15).

Later, attention-based networks have been introduced, where the attention mechanism is designed to learn when pairs of inputs depend crucially on one another, a capability that has demonstrably improved machine translation by making sure that the translated output represents all of the given input even when their word-orders do not correspond exactly (Vaswani, Shazeer, Parmar, Uszkoreit, Jones, Gomez, Kaiser, & Polosukhin, 2017). This has led to rapid advances in the field, including the contextualized BERT (Devlin, Chang, Lee, & Toutanova, 2018) and ELMo (Peters, Neumann, Iyyer, Gardner, Clark, Lee, & Zettlemoyer, 2018) models. For example, the ELMo model reports good results on traditional NLP tasks including question answering, coreference resolution, semantic role labelling, and part-of-speech tagging, and the authors attribute this success to the model’s different neural-network layers implicitly representing several different kinds of linguistic structure. The survey and experiments of Hupkes et al. (2020) evaluate three such neural networks on a range of tasks related to composition, concluding that each network has strengths and weaknesses, that the results are a stepping stone rather than an endpoint, and that developing consensus around how such tasks should be designed, tested and shared is a crucial task in itself.

Even without the cost of encoding order information, Iyyer, Manjunatha, Boyd-Graber, and Daumé III (2015) demonstrated a deep *averaging* network, showing that network depth could in some cases compensate for the lack of syntactic sophistication in unordered models — with considerable computational savings. At the time of writing, this has developed into a very open research question: do neural networks need extra linguistic information as inputs to properly understand language, or can they actually *recover* such information as a byproduct of training on raw text input? Does a complete NLP system need components for tokenization, part-of-speech tagging, syntactic parsing, named entity recognition, and so on, or can some or all of these be replaced by a single vector language model? If so, how do we describe the way meaningful units are composed into larger meaningful structures in such a model?

Tensor networks are one of the possible mathematical answers to this question. Their use in libraries such as TensorFlow has become ubiquitous (see (Géron, 2019) and numerous papers and packages referenced therein), though as cautioned above, users should be wary of the equivalence between the use of matrices and multidimensional arrays and tensor algebra that is often assumed in software documentation. More explicit evidence is presented by work that continues to build upon Smolensky’s introduction of tensors to AI: for example McCoy, Linzen, Dunbar, and Smolensky (2020) present evidence that the sequence-composition effects of Recurrent Neural Networks (RNNs) can be approximated by Tensor Product Decomposition Networks, at least in cases where using this structure provides measurable benefits over bag-of-words models (see also McCoy, Linzen, Dunbar, and Smolensky (2018) for a more detailed presentation). It has also been shown that Tensor Product Networks can encode grammatical structure more effectively than LSTMs for generating image captions (Huang, Smolensky, He, Deng, & Wu, 2017), achieving for example a BLEU-4 score of 0.305 compared the CNN-LSTM’s result of 0.292. Tensor product networks have also been used to construct an attention mechanism from which grammatical structure can be recovered by unbinding role-filler tensor compositions (Huang, Deng, Wu, Liu, & He, 2019). Explicitly quantum networks for natural language processing are described by Wiebe, Bocharov, Smolensky, Troyer, and Svore (2019). The range of challenges and application opportunities in AI for compositional vector representations is by now much-studied and valued, and quantum-inspired tensor networks already successfully combine and extend many of the mathematical features that are core to AI today.

6. Physical Implementations on Quantum Computers

Quantum computing has regularly made the front page of scientific news from 2019 to 2021. This section gives a glimpse of opportunities this opens for quantum mathematics in AI.

With AI based on classical computation, vectors and matrices are already ubiquitous. While the algebra and geometry behind them contains many riches that overlap with quantum mechanics, studying this overlap in detail is an investment that many researchers might consider esoteric and risky — there are more mainstream state-of-the-art ways to make dramatic progress in AI and machine learning without learning quantum theory and tensor algebra. But what if there was a platform for computation where vectors could be represented in exponentially smaller memory, and instead of being an operation with quadratic cost, the tensor product was just the most suitable mathematical representation of a natural system? Quantum computers may soon provide this.

As outlined in Section 5.2, the tensor product arises naturally in quantum mechanics. Computations using quantum circuits can be used to manipulate tensor products (for example, from separable to entangled states (Bernhardt, 2019, Ch 4)), but no ‘work’ is required to compute the tensor product of a 2-qubit system in the first place: the wavefunction or state of a 2-qubit system *is* represented as the tensor product of the individual qubit systems, irrespective of any computational operation we perform on these qubits (Ying, 2010, §2).

Perhaps the biggest computational promise that follows from the tensor product is the potential for representing exponentially more information. A classical register of n physical bits can represent n bits of information (choices between 0 and 1). In quantum computing, these combinations correspond to basis vectors — if the state vector of a qubit is represented as a vector in \mathbb{C}^2 with basis vectors $\{|0\rangle, |1\rangle\}$, then a register with n qubits is represented as a state vector in $\bigotimes_1^n \mathbb{C}^2$, which has dimension 2^n . Instead of a state of the whole system, each of those 2^n combinations represents a basis vector for a coordinate, and if each of these coordinates can be accessed, the capacity of the memory grows exponentially instead of linearly in the number of qubits in the register! In AI applications, it is easy to see the appeal of this — an embedding vector of 256 dimensions using a 4-byte floating point number for each coordinate requires 1KB of RAM, which can in theory be represented by a tensor product of 8 qubits. Of course, there are huge challenges with this idea, starting with the physical fact that we could never observe such continuous coordinates directly, only their probabilistic quantization to a pure state upon measurement. Nonetheless, the promise of exponential quantum memory has led to ingenious research ideas (particularly the ‘bucket-brigade’ qRAM protocol of (Giovannetti, Lloyd, & Maccone, 2008)), and the idea is even presented in an early-reader board book for children (Ferrie, 2018). There are still practical physical challenges and theoretical caveats to any premature claims that ‘quantum algorithms give exponential improvements’ (Aaronson, 2015). Still, if the promise of exponential qRAM is obvious and compelling enough to excite the youngest readers, it is hard to imagine that no breakthroughs will happen in the next few decades.

Still, by the end of 2019, most of the papers published on quantum algorithms were, from a physical point-of-view, advanced thought-experiments or simulations: they present mathematical techniques or simulated results on classical hardware, not experimental results from quantum computers. Shor’s quantum algorithm for prime factorization was published in 1994 (Bernhardt, 2019, Ch 9), but at the time of writing, the largest integer factorized on a quantum computer appears to be 291,311 (Li, Dattani, Chen, Liu, Wang, Tanburn, Chen, Peng, & Du, 2017). There are very particular problems that can already be solved much faster on a quantum computer than a classical computer, including the problem used by (Arute, Arya, Babbush, Bacon, Bardin, Barends, Biswas, Boixo, Brandao, Buell, et al., 2019) which was to simulate and predict the output of a pseudo-random quantum circuit. However, this task was proposed *in order to* demonstrate quantum advantages, rather than for some existing use-case, a critique expounded by Zhong, Wang, Deng, Chen, Peng, Luo, Qin, Wu, Ding, Hu, et al. (2020).

There are several research developments that *are* poised to take advantage of quantum computers in AI, in areas including deep learning in particular (Wiebe, Kapoor, & Svore, 2014), machine

learning in general (Biamonte, Wittek, Pancotti, Rebentrost, Wiebe, & Lloyd, 2017), and language processing including parsing (Wiebe et al., 2019), but it remains normal for such papers to develop algorithms and sometimes results on real datasets without actually implementing them on quantum machines. Even papers that describe physical implementation in detail including those in quantum optical neural networks (Steinbrecher, Olson, Englund, & Carolan, 2019) and single photon image classification (Fischbacher & Sbaiz, 2020) report results on high-end classical hardware and proposals for quantum hardware. This is not a negative criticism, just a description of the research frontier: in celestial mechanics, the orbit of a satellite was studied and predicted for centuries before spacecraft could actually be launched.

In 2020 and 2021, this has been changing: AI experiments on quantum hardware have been successfully carried out. The single-photon classification experiment of Wang, Xiao, Yi, Ran, and Xue (2020) uses a quantum-mechanical platform with a single qubit. For comparison, the simulated work of Fischbacher and Sbaiz (2020) classifies all ten digits, whereas Wang et al. (2020) build a classifier that only distinguishes zeros and ones: this exemplifies the sort of tradeoff researchers have made in using real-but-limited quantum computing resources. In the work recently reported by Abbas, Sutter, Zoufal, Lucchi, Figalli, and Woerner (2021), an actual quantum neural network is trained and shown to learn faster and more effectively than a classical network (as measured by Fisher information and effective dimension), showing as much as a 250% improvement over classical training using the `ibmq_montreal` 27-qubit hardware.

In natural language processing on quantum computers, the most dramatic development so far is perhaps the demonstration of a working Quantum NLP system on one of IBM’s quantum devices (Meichanetzidis, De Felice, Toumi, Coecke, Gogioso, & Chiappori, 2020; Coecke, de Felice, Meichanetzidis, & Toumi, 2020). Mathematically speaking, this work has several key ingredients. A key part is the use of compact closed categories as an axiomatic model for quantum mechanics (Abramsky & Coecke, 2009), which enables the DisCoCat model from Coecke et al. (2010) to use the same mathematical language. Diagrammatic calculus, developed over some years and explained thoroughly by Coecke and Kissinger (2017), enables quantum structures for sentences that combine their syntactic structure (in terms of combinations) and semantic content (in terms of vectors) to be represented together. And the paper of Meichanetzidis et al. (2020) explains in detail how this enables a quantum compiler to break these expressions down into quantum gates, compiled into quantum circuits and run on a quantum computer. The sentences used to train the model are still simple toy examples such as ‘*Alice loves Bob*’ and ‘*Bob who is cute loves Alice who is rich*’ — as with prime factorization and image recognition, the implementation using qubits in quantum computers does not yet rival the scale of NLP on classical hardware — but the first implementation of a compositional NLP system on quantum hardware has been accomplished successfully, and given the level of investment in both quantum computing and NLP, more will follow.

7. Other Areas Related to AI and Quantum Theory

This final section refers briefly to other areas related to both AI and quantum theory that, to avoid excessive length, have not been emphasized in this paper, but which are likely to be important avenues for future work. Many more papers can be found on this stream of topics than we are able to list here. We encourage interested readers to investigate the references provided in this section for further links to other papers and more detailed avenues.

7.1 Quantum Search and Automated Problem Solving

Quantum search is a large algorithmic topic in itself. In this context, we mean ‘search’ in the sense of ‘database search’ or ‘tree search’, rather than the search engines of information retrieval: quantum search algorithms are typically designed to find a particular element that solves a given problem in computationally faster time than classical algorithms (Nielsen & Chuang, 2002, Ch 6). A canonical

example of such a speedup is Grover’s search algorithm, which locates a unique element out of n choices in \sqrt{n} time (Bernhardt, 2019, Ch 9). The details are not covered here, though the process itself is a hallmark example of a quantum oracle that manipulates bases step-by-step in such a way that the states of the non-solutions cancel one another out and the state of the solution eventually ‘sticks out’ after at most \sqrt{n} repetitions.

Grover’s search algorithm has been used as a building block for approaching other higher-level problems. These include various forms of tree search (Tarrataca & Wichert, 2011b) and challenges that can be formulated as symbol manipulation problems, such as a block puzzle (Tarrataca & Wichert, 2011a). For an overview see Wichert (2020, Ch 10–12). It is worth noting that core data structure operations such as tree search have been investigated for basic hardware components such as qRAM (Giovannetti et al., 2008) while they are also being developed for higher-level problem solving in quantum AI: for example, Arunachalam, Gheorghiu, Jochym-O’Connor, Mosca, and Srinivasan (2015) analyse the importance of reducing the number of lookup operations in the ‘bucket brigade’ qRAM protocol error-rates (demonstrating in the process some of the great difficulties in implementing a real qRAM).

This evolution is historically different from the development of classical computing, in which basic data structures were reliably and readily available decades before the heyday of machine learning. (This is not to say that there are no more developments to be made in classical data structures, but that the machine learning practitioner often has no reason to wonder whether a dictionary lookup is from a binary search tree vs. a hashmap — they ‘just work’.) By contrast, in quantum computing, these innovations are happening at the same time and facing similar challenges — quantum heuristics are sometimes needed not only to expedite computation, but to give *more* reliable results, because reducing the number of operations reduces errors. While such details might be seen as frustratingly low-level in most of machine learning today, this is an area where quantum computing and quantum AI may cross-fertilize each other’s early development.

7.2 Quantum Probability and Cognition

The geometric and algebraic structures of quantum theory affect probability as well as logic. In quantum theory, the probability of finding a system in state $|\psi\rangle$ to be in the pure basis state $|\lambda_i\rangle$ is equal to $|\langle\lambda_i|\psi\rangle|^2$, which is called the Born rule. For normalized unit vectors, the scalar product $\langle\lambda_i|\psi\rangle$ is well-known to be the cosine of the angle between the two vectors — so in quantum theory, probability is geometrically related to angles between vectors. This contrasts with so-called classical probability where probabilities are obtained from ratios between volumes (Isham, 1995, Ch 2)(Khrennikov, 2010, Ch 2). Quantum probability behaves differently from classical probability: for example, conditional probability depends on the spectrum of the observable in question (Khrennikov, 2010, p. 35), and the fact that a measurement of one of the variables was performed beforehand (Isham, 1995, p. 165).

Quantum probability has been used to explain and accurately model order-effects on attitudes, where asking people the same questions in a different order has been shown to give different answers (Busemeyer & Bruza, 2012, Ch 3), resulting in various disjunction effects in decision making. For example, preferences made with more information have been shown to be different from those made with less information, irrespective of what information has been learned in the meantime, a violation of the Sure Thing Principle (Pothos & Busemeyer, 2013). In semantics, such models have been used to model compositional behavior including nonseparability in the human lexical representations (Bruza, Kitto, Ramm, & Sitbon, 2015).

The affordances of quantum probabilities in cognitive modelling have been exploited by some authors to support more human-like models of automated reasoning and decision making. For example, Kitto and Boschetti (2013) made use of the structural nature of quantum probability in an agent based model capable of representing phenomena such as cognitive dissonance and attitude change in a social context. Similarly, a stream of work in developing quantum-like Bayesian network

models (Moreira, Tiwari, Pandey, Bruza, & Wichert, 2020) has resulted in sophisticated methods for modelling the evolution of the beliefs of a decision maker.

7.3 Quantum Probability and Language Modelling

A key part of the language modeling work of Sordoni et al. (2013) was the use of off-diagonal elements in a matrix to model the joint probability of bigrams (where the on-diagonal elements model unigram probabilities). The mathematics of this particular construction is presented intuitively and explored thoroughly in the *statistical algebra* of Bradley (2020). This leads to a surprising and clear correspondence between Formal Concept Analysis (Ganter & Wille, 1999), in which objects and attributes are represented in a lattice, and quantum probability. The basic construction is to take the cross table (the matrix showing which objects possess which attributes), and multiply it by its transpose, which gives a matrix whose off-diagonal elements capture the overlaps between objects. This conditional dependence information is reflected in the partial trace, which is a basis-dependent mapping from the product space $V^* \otimes V$ back to V . Bradley (2020) argues that the partial trace is the quantum analogy of the classical process of marginalizing a joint distribution — but whereas in classical probability, marginal probabilities retain no information about the overlaps in the original joint distribution, in the quantum case, some of this information is preserved. This richness may lead to quantum probability becoming a more common tool of choice for language and concept models.

7.4 Real and Complex Numbers in AI

One place where physics, classical logic, and machine learning use quite different fundamental mathematical building blocks is in the choice of number field. Machine learning and AI tend to use real vectors because the basic features are measurements; classical logic and computer science have used binary numbers thanks largely to the enormous influence of George Boole (1854, Ch 2); whereas physics most prevalently uses complex vectors, not just in quantum mechanics but also in electrodynamics.

So far, the reasons for *not* using complex and binary numbers in machine learning seem to be twofold (and reasonable): there are no widespread and immediately intuitive interpretations for them; and the field is progressing very quickly without them. However, there are computational benefits to using complex numbers. For example, the convolution operator, sometimes used for vector binding, is just the addition of phase angles, and is thus $O(n)$ rather than the $O(n \log n)$ of the real convolution operation optimized using Fast Fourier Transforms (see Plate (2003)). And in several experiments on automatic inference, we have ourselves found that using binary and complex vectors sometimes yields much better results than their real-valued counterparts, for reasons not yet properly understood (Widdows & Cohen, 2015). There are a few machine learning papers on using complex numbers and even quaternions and octonions in machine learning, for example Trabelsi, Bilaniuk, Zhang, Serdyuk, Subramanian, Santos, Mehri, Rostamzadeh, Bengio, and Pal (2018). As quantum computing becomes more mainstream, we should expect the use of complex numbers to become more widespread.

7.5 Complex Time Evolution

Imaginary and complex numbers are particularly important in quantum mechanics for representing the momentum operator $\hat{p} = -i\hbar \frac{\partial}{\partial x}$ and the resulting Hamiltonian in the Schrödinger wave equation. This technique has roots in harmonic analysis, particularly with the simple form of the plane wavefunction $\Psi(x, t) = Ae^{i(kx - \omega t)}$. Harmonic functions can be written in this way thanks to the identity $e^{i\theta} = \cos \theta + i \sin \theta$, and the technique for representing time evolution could potentially be applied to situations where harmonic functions are used to model word position as attention moves along a sentence (Vaswani et al., 2017).

7.6 Again, Why Quantum Theory?

Throughout this paper we have introduced and surveyed several areas where mathematics used in quantum theory is also fruitful in AI. However, we are yet to answer the questions “What is quantum mathematics?” and “Why is it important in AI?”, delaying them to this point. This has enabled us to demonstrate the many and diverse ways in which mathematics has been abstracted from quantum theory and applied to AI. But there is no clear characterization of what precisely makes mathematics ‘quantum’.

This is well-illustrated by contrast with the other great pillar of 20th century physics, general relativity. The mathematics of general relativity can be defined crisply as the differential geometry of Lorentzian manifolds, which are smooth 4-manifolds with a pseudo-Riemannian metric of signature (1, 3) — that is, there must be 4 dimensions, and the metric must treat three of them interchangeably (the spacelike dimensions) and one with the opposite sign (the timelike dimension), giving a local distance metric that takes the form $dt^2 - dx^2 - dy^2 - dz^2$ (choosing units so that the speed of light is 1).

After years of searching, we are confident that *there is no similarly clear characterization of quantum mathematics*. That disjunctions are non-distributive or products non-separable have been proposed, but can be rejected on the grounds that such mathematical structures are prominent in many other areas, and many quantum models do not explicitly use these properties.

So the question “What makes mathematics distinctly quantum?” remains a matter for discussion. In our opinion, quantum mathematics as a whole is the mathematics of self-adjoint operators in complex Hilbert space. In theory, complex numbers are needed to guarantee that roots of polynomials exist, which guarantees that eigenvalues and eigenvectors exist. The self-adjoint requirement guarantees that these eigenvalues are real numbers, which is necessary for them to represent physical properties. This mathematical argument is standard: in this case, we are following particularly the narrative of Coecke and Kissinger (2017, Ch 6), a similar presentation is given by Van Rijsbergen (2004), and indeed Dirac (1930, Ch 1, 2).

Much of this mathematical apparatus is general, and has been used in a wide array of fields beyond quantum theory. In this paper we have demonstrated a number of these overlaps with AI. However, with this point comes an important caveat: most AI applications are only using a part of quantum mathematics, since few use complex numbers. For example, the information retrieval applications discussed in Section 3 are explained explicitly in terms of quantum logic (in the case of Widdows and Peters (2003)) and quantum probability (in the case of Sordani et al. (2013)), but use only real numbers throughout. Rather than a criticism, this could become a recommendation: quantum-inspired AI systems should systematically consider and investigate the use of complex numbers.

Instead of embodying a distinct piece of mathematics, the motivation for developing quantum-inspired models in AI and related areas has sometimes been the obvious shortcomings of so-called ‘classical’ models. Quantum theory acknowledges that the future is not wholly predictable, that combined systems cannot always be separated into atomic parts, and that the order in which observations are made sometimes affects the outcome. These points are often obscured by the oft-repeated claim “classical physics is the physics of everyday life”, even though *quantum physics is more like daily life* in these respects (points discussed in more detail by Busemeyer and Bruza (2012) and Widdows and Bruza (2007)). As these areas have been explored in detail, with hindsight it is easy to understand why quantum structures were seen as promising alternatives to reductionist symbolic determinism: and once this approach is considered, mathematical overlaps between quantum mechanics and other application areas are easy to find. Without the advent of quantum computing, these overlaps may have been considered a mathematical curiosity: a good motivation perhaps for individual system designs so long as they perform well in practice. But now with quantum computing coming of age, the importance of quantum mathematics in computer science is set to increase, because it will be key to answering the question “What’s the most effective way to implement this?”

8. Conclusion

Quantum theory and artificial intelligence use much of the same mathematics, and the importance of the overlapping areas has grown with the ubiquitous use of vectors in AI. As well as simple superposition and similarity operations, quantum theory offers well-studied formulations of logical inference, concept combination, and probability, many of which have been demonstrated to give good results on various AI tasks, including examples from natural language processing, image processing, reasoning and inference. The mathematics used in these applications is often well-developed thanks to its use in other fields, and in particular, Grassmannian and tensor algebra offers more in terms of duality and symmetry than is typically used in AI. Some of these algebraic operations are costly on classical hardware, but much cheaper or even free in quantum computation, which in 2021 is rapidly becoming more practical. The authors hope that this paper helps AI researchers familiar with vectors and their uses to turn more often to quantum theory as a rich source of mathematical motivation for advances today, and to approach the advent of quantum computing with confidence and excitement.

Appendix: Category Theory for Quantum AI

Category theory is one of the less familiar branches of mathematics to many researchers in AI, and this appendix is designed as a very brief supplement. Interested readers are referred particularly to the survey article of Selinger (2010) and the graduate text of Heunen and Vicary (2019) for a proper explanation.

Much of modern mathematics can be thought of as the study of objects with some defined structure, along with maps that preserve this structures. Standard examples well-known by the middle of the 20th century include sets and maps, groups and homomorphisms, topological spaces and continuous maps, differentiable manifolds and differentiable maps, and of course, vector spaces and linear maps.

The abstract notion of a category was developed in the 20th century to reason about such structures in a common fashion. A category consists of a collection of *objects*, and for every pair of objects A and B , a collection of *morphisms* from A to B . Morphisms must be composable in sequence in an associative manner (so that $f \circ (g \circ h) = (f \circ g) \circ h$), and there must be an identity morphism id so that $\text{id} \circ f = f \circ \text{id} = f$.

In many categories, the morphisms are some kind of function. These include **Set**, the category of sets and maps, and **Vect**, the category of vector spaces and linear maps. There are also important cases where the morphisms are not functions, including the category of sets and *relations*, which is often called **Rel**. A relation between two sets A and B is a subset of their Cartesian product $A \times B$, intuitively a ‘many to many correspondence’, and relations can also be composed in a way that satisfies the definition of a category (Heunen & Vicary, 2019, §0.1.3). Because of the potential for one-to-many correspondences, if we think of ‘applying a morphism’ as ‘evolving in time’, the category **Rel** can be used as a model for ‘nondeterministic classical physics’, because a single state representing the present can evolve into more than one possible state representing the future.

Following the pattern of abstract algebra, further definitions are introduced and their properties and consequences explored. Some of these (here described intuitively but not precisely) are as follows. A *functor* is itself a mapping between two categories that maps objects to objects and morphisms to morphisms in a regular composable fashion. The *opposite* or *dual* category C^{op} has the same objects as C but ‘reversed’ morphisms: if we define the set of morphisms from A to B in the category C as $C(A, B)$, then the set of morphisms from B to A in C^{op} is defined by $C^{op}(B, A) = C(A, B)$. This is simple but initially confusing, especially with the most obvious category **Set** of sets and functions — how can a function from A to B be considered a morphism from B to A ? In this surprising case, the answer is just ‘the definition satisfies identity and associativity, so it forms a category’. However, with the categories of **Rel** of sets and relations, and **FVect** of (finite dimensional) vector

spaces and linear maps, the notion of dual categories becomes much more concrete, as the objects and morphisms themselves have duals, not only the category as a whole. Much of the power in categorical quantum mechanics arises from the use of category theory to explore the properties and consequences of duality in these settings.

A *closed* category is one where the morphisms themselves have the structure of objects in the category. For example, the mappings from a finite set A to a finite set B themselves form a finite set (an element of B is selected for each element of A , so the number of such possibilities is $|B|^{|A|}$), so finite sets are a closed category. A *monoidal* category is one where objects from the same category can be composed with one another ‘in parallel’ using an operator written \otimes (as well as morphisms being composed ‘in sequence’), and there needs to be an identity object I that satisfies basic compatibility conditions with the identity morphism. This gives the objects of the category the structure of a monoid, which is a group without the requirement of each element having an inverse. In a monoidal category, an object A may have a *dual* object A^* , with canonical morphisms from I to $A^* \otimes A$ and vice versa. If these morphisms contract to the identity in given ways, the category is called *compact*.

A *dagger* category is one where each morphism $f : A \rightarrow B$ has a corresponding dagger morphism $f^\dagger : B \rightarrow A$, with algebraic properties generalized from those of the adjoint operator in linear algebra. Dagger categories were used by Abramsky and Coecke (2004) for describing quantum information protocols, using the term *strongly compact closed categories*. The term ‘dagger category’ was introduced as part of a more general exploration of their mathematical properties by Selinger (2007).

The connection of notation and terminology with linear algebra is quite deliberate: finite dimensional complex vector spaces form a monoidal category whose monoidal product is the tensor product, and whose dual objects are the dual vector spaces. Hence the notation A^* for the dual of A , and $A \otimes B$ for the monoidal category product, are just the same definitions as given in linear algebra.

The relationship with quantum mechanics, called *categorical quantum mechanics*, is due particularly to Abramsky and Coecke (2009), and has become part of the mathematical backbone of work on understanding quantum processes by Coecke and Kissinger (2017) and in related works. Categorical quantum mechanics seeks to describe quantum mechanics and quantum processes in terms of mathematical properties like those outlined above.

This contributes to quantum mathematics in two ways. First, the results of quantum mechanics can be linked even more directly to key mathematical concepts. For example, Heunen and Vicary (2019, Ch. 4) uses the contrasting monoidal structures of the categories **Set** and **FHilb** (finite dimensional Hilbert Spaces and bounded linear maps) to explain why we should expect that state can be uniformly copied in classical computing but not in quantum computing; and they relate categorical properties to aspects of entanglement and quantum teleportation throughout. Second, understanding such phenomena at a more abstract mathematical level can lead to applications in other areas and cross-fertilization. For example, the category of sets and relations (more general than the category of sets and maps, because a relation can be many-to-many) sometimes behaves much more like the category of vector spaces than that of sets, particularly in relation to entanglement, which the statistical algebra of Bradley (2020) uses to describe concepts in a formal concept lattice. And of particular importance, the use of category theory links the *pregroup grammars* of Lambek calculus to the composed vector representations of Coecke et al. (2010), the algebraic representation of pronouns (Sadrzadeh, Clark, & Coecke, 2013, 2014), and the gate-based quantum circuit implementation of Meichanetzidis et al. (2020). Here, for example, the notion of monoidal products and duals is used to give a categorical model for a linguistic notion such as “a transitive verb operates on two nounphrases” and a vector notion such as “a type $(0, 2)$ tensor operates on two type $(1, 0)$ vectors”, in such a way that the verb and noun representations can be implemented on a quantum computer.

While category theory can sound very abstract and technical, it gives clearer insights into why apparently similar systems sometimes behave differently, and why some apparently very different systems share common structures. Today especially, it provides an ideal opportunity for mathe-

maticians with a background in algebra to contribute to technologies for computing, security, and communication.

References

- Aaronson, S. (2013). *Quantum computing since Democritus*. Cambridge University Press.
- Aaronson, S. (2015). Read the fine print. *Nature Physics*, 11(4), 291–293.
- Abbas, A., Sutter, D., Zoufal, C., Lucchi, A., Figalli, A., & Woerner, S. (2021). The power of quantum neural networks. *Nature Computational Science*, 1(6), 403–409.
- Abramsky, S., & Coecke, B. (2004). A categorical semantics of quantum protocols. In *Proceedings of the 19th Annual IEEE Symposium on Logic in Computer Science, 2004.*, pp. 415–425. IEEE.
- Abramsky, S., & Coecke, B. (2009). Categorical quantum mechanics. *Handbook of quantum logic and quantum structures*, 2, 261–325.
- Aerts, D., & Czachor, M. (2004). Quantum aspects of semantic analysis and symbolic artificial intelligence. *J. Phys. A: Math. Gen.*, 37, L123–L132.
- Aerts, D., Durt, T., Grib, A., Van Bogaert, B., & Zapatrin, R. (1993). Quantum structures in macroscopic reality. *International Journal of Theoretical Physics*, 32(3), pp489.
- Aerts, D., & Sozzo, S. (2011). Quantum structure in cognition: Why and how concepts are entangled. In *International Symposium on Quantum Interaction*, pp. 116–127. Springer.
- Arunachalam, S., Gheorghiu, V., Jochym-O’Connor, T., Mosca, M., & Srinivasan, P. V. (2015). On the robustness of bucket brigade quantum ram. *New Journal of Physics*, 17(12), 123010.
- Arute, F., Arya, K., Babbush, R., Bacon, D., Bardin, J. C., Barends, R., Biswas, R., Boixo, S., Brandao, F. G., Buell, D. A., et al. (2019). Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779), 505–510.
- Bankova, D., Coecke, B., Lewis, M., & Marsden, D. (2019). Graded hyponymy for compositional distributional semantics. *Journal of Language Modelling*, 6(2), 225–260.
- Baroni, M., Bernardi, R., Zamparelli, R., et al. (2014). Frege in space: A program for compositional distributional semantics. *Linguistic Issues in language technology*, 9(6), 5–110.
- Baroni, M., & Zamparelli, R. (2010). Nouns are vectors, adjectives are matrices: Representing adjective-noun constructions in semantic space. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Basile, P., Caputo, A., & Semeraro, G. (2011). *Negation for Document Re-ranking in Ad-hoc Retrieval*, pp. 285–296. Springer, Berlin, Heidelberg.
- Bernhardt, C. (2019). *Quantum Computing for Everyone*. Mit Press.
- Biamonte, J., Wittek, P., Pancotti, N., Rebentrost, P., Wiebe, N., & Lloyd, S. (2017). Quantum machine learning. *Nature*, 549(7671), 195–202.
- Birkhoff, G., & von Neumann, J. (1936). The logic of quantum mechanics. *Annals of Mathematics*, 37, 823–843.
- Bishop, R. L., & Goldberg, S. I. (1968). *Tensor analysis on manifolds*. Macmillan, republished by Dover (2012).
- Bollobás, B. (1998). *Modern Graph Theory*. No. 184 in Graduate Texts in Mathematics. Springer-Verlag.
- Boole, G. (1854). *An Investigation of the Laws of Thought*. Macmillan. Dover edition, 1958.
- Bradley, T.-D. (2020). *At the Interface of Algebra and Statistics*. Ph.D. thesis, City University of New York.

- Bruza, P., Kitto, K., & McEvoy, D. (2008). Entangling words and meaning. In *Quantum Interaction: Proceedings of the Second Quantum Interaction Symposium (QI-2008)*, pp. 118–124. College Publications.
- Bruza, P. D., Kitto, K., Ramm, B. J., & Sitbon, L. (2015). A probabilistic framework for analysing the compositionality of conceptual combinations. *Journal of Mathematical Psychology*, *67*, 26–38.
- Bruza, P. D., Widdows, D., & Woods, J. (2009). A quantum logic of down below. In Engesser, K., Gabbay, D. M., & Lehmann, D. (Eds.), *Handbook of quantum logic and quantum structures*. Elsevier.
- Busemeyer, J., & Bruza, P. (2012). *Quantum Models of Cognition and Decision*. Cambridge University Press.
- Camacho-Collados, J., & Pilehvar, M. T. (2018). From word to sense embeddings: A survey on vector representations of meaning. *Journal of Artificial Intelligence Research*, *63*, 743–788.
- Clark, S., & Pulman, S. (2007). Combining symbolic and distributional models of meaning.. In *AAAI Spring Symposium: Quantum Interaction*, pp. 52–55.
- Clarke, D. (2012). A context-theoretic framework for compositionality in distributional semantics. *Computational Linguistics*, *38*(1), 41–71.
- Coecke, B., de Felice, G., Meichanetzidis, K., & Toumi, A. (2020). Foundations for near-term quantum natural language processing. *arXiv preprint, arXiv:2012.03755*.
- Coecke, B., & Kissinger, A. (2017). *Picturing quantum processes*. Cambridge University Press.
- Coecke, B., Sadrzadeh, M., & Clark, S. (2010). Mathematical foundations for a compositional distributional model of meaning. *arXiv preprint, arXiv:1003.4394*.
- Cohen, T., & Widdows, D. (2015). Embedding probabilities in predication space with hermitian holographic reduced representations. In *Ninth International Symposium on Quantum Interaction*, pp. 245–257. Springer.
- Cohen, T., Widdows, D., Vine, L. D., Schvaneveldt, R., & Rindfleisch, T. C. (2012). Many paths lead to discovery: Analogical retrieval of cancer therapies. In *Sixth International Symposium on Quantum Interaction*.
- Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2018). BERT: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint, arXiv:1810.04805*.
- Dirac, P. (1930). *The principles of quantum mechanics* (4th edition, 1958, reprinted 1982 edition). Clarendon Press, Oxford.
- Dorst, L., Fontijne, D., & Mann, S. (2010). *Geometric algebra for computer science: an object-oriented approach to geometry*. Elsevier.
- Fearnley-Sander, D. (1979). Hermann Grassmann and the creation of linear algebra. *The American Mathematical Monthly*, *86*(10), 809–817.
- Ferrie, C. (2018). *Quantum Information for Babies*. Sourcebooks Explore.
- Fischbacher, T., & Sbaiz, L. (2020). Single-photon image classification. *arXiv preprint, arXiv:2008.05859*.
- Ganter, B., & Wille, R. (1999). *Formal Concept Analysis: Mathematical Foundations*. Springer.
- Garg, D., Ikbāl, S., Srivastava, S. K., Vishwakarma, H., Karanam, H., & Subramaniam, L. V. (2019). Quantum embedding of knowledge for reasoning. In *Advances in Neural Information Processing Systems*, pp. 5594–5604.
- Géron, A. (2019). *Hands-On Machine Learning with Scikit-Learn, Keras, and TensorFlow: Concepts, Tools, and Techniques to Build Intelligent Systems*. O’Reilly Media.

- Giovannetti, V., Lloyd, S., & Maccone, L. (2008). Quantum random access memory. *Physical review letters*, *100*(16), 160501.
- Grassmann, H. (1862). *Extension Theory*. History of Mathematics Sources. American Mathematical Society, London Mathematical Society. Translated by Lloyd C. Kannenberg (2000).
- Greenstein, G., & Zajonc, A. G. (1997). *The Quantum Challenge: modern research on the Foundations of Quantum Mechanics*. Jones and Bartlett Publishers, Sudbury, Massachusetts.
- Grefenstette, E., & Sadrzadeh, M. (2011). Experimental support for a categorical compositional distributional model of meaning. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Hastie, T., Tibshirani, R., & Friedman, J. H. (2001). *The Elements of Statistical Learning*. Springer Series in Statistics.
- Heunen, C., & Vicary, J. (2019). *Categories for Quantum Theory: an introduction*. Oxford University Press.
- Hinton, G. E. (1990). Preface to the special issue on connectionist symbol processing. *Artificial Intelligence*, *46*(1-2), 1–4.
- Huang, Q., Deng, L., Wu, D., Liu, C., & He, X. (2019). Attentive tensor product learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33, pp. 1344–1351.
- Huang, Q., Smolensky, P., He, X., Deng, L., & Wu, D. (2017). Tensor product generation networks for deep NLP modeling. *arXiv preprint, arXiv:1709.09118*.
- Hupkes, D., Dankers, V., Mul, M., & Bruni, E. (2020). Compositionality decomposed: how do neural networks generalise?. *Journal of Artificial Intelligence Research*, *67*, 757–795.
- Isham, C. J. (1995). *Lectures on Quantum Theory*. Imperial College Press, London.
- Iyyer, M., Manjunatha, V., Boyd-Graber, J., & Daumé III, H. (2015). Deep unordered composition rivals syntactic methods for text classification. In *Proceedings of the 53rd annual meeting of the Association for Computational Linguistics (ACL)*, pp. 1681–1691.
- Jänich, K. (1994). *Linear algebra*. Undergraduate Texts in Mathematics. Springer-Verlag.
- Kanerva, P. (2009). Hyperdimensional computing: An introduction to computing in distributed representation with high-dimensional random vectors. *Cognitive Computation*, *1*(2), 139–159.
- Kartsaklis, D., & Sadrzadeh, M. (2014). A study of entanglement in a categorical framework of natural language. *arXiv preprint, arXiv:1405.2874*.
- Khrennikov, A. (2010). *Ubiquitous Quantum Structure: From Psychology to Finance*. Springer.
- Kitto, K., & Boschetti, F. (2013). Attitudes, ideologies and self-organization: information load minimization in multi-agent decision making. *Advances in Complex Systems*, *16*(02n03), 1350029.
- Lambek, J. (2001). Type grammars as pregroups. *Grammars*, *4*(1), 21–39.
- Landauer, T., & Dumais, S. (1997). A solution to Plato’s problem: The latent semantic analysis theory of acquisition. *Psychological Review*, *104*(2), 211–240.
- Lang, S. (2002). *Algebra*. No. 211 in Graduate Texts in Mathematics. Springer-Verlag.
- Lawvere, F. W., & Schanuel, S. H. (2009). *Conceptual mathematics: a first introduction to categories*. Cambridge University Press.
- LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, *521*(7553), 436–444.
- Levy, S. D., & Gayler, R. (2008). Vector symbolic architectures: A new building material for artificial general intelligence. In *Proceedings of the 2008 conference on Artificial General Intelligence 2008*, pp. 414–418. IOS Press.

- Lewis, M. (2019). Compositional hyponymy with positive operators. In *Proceedings of the International Conference on Recent Advances in Natural Language Processing (RANLP 2019)*, pp. 638–647, Varna, Bulgaria.
- Lewis, M. (2020). Towards logical negation for compositional distributional semantics. *arXiv preprint, arXiv:2005.04929*.
- Li, Z., Dattani, N. S., Chen, X., Liu, X., Wang, H., Tanburn, R., Chen, H., Peng, X., & Du, J. (2017). High-fidelity adiabatic quantum computation using the intrinsic hamiltonian of a spin system: Application to the experimental factorization of 291311. *arXiv preprint, arXiv:1706.08061*.
- Manning, C. D., Raghavan, P., & Schütze, H. (2008). *Introduction to information retrieval*. Cambridge university press.
- McCoy, R. T., Linzen, T., Dunbar, E., & Smolensky, P. (2018). RNNs implicitly implement tensor product representations. *arXiv preprint, arXiv:1812.08718*.
- McCoy, R. T., Linzen, T., Dunbar, E., & Smolensky, P. (2020). Tensor product decomposition networks: Uncovering representations of structure learned by neural networks. *Proceedings of the Society for Computation in Linguistics*, 3(1), 474–475.
- Meichanetzidis, K., De Felice, G., Toumi, A., Coecke, B., Gogioso, S., & Chiappori, N. (2020). Quantum natural language processing on near-term quantum computers. In *Semantic Spaces at the Intersection of NLP, Physics, and Cognitive Science (SemSpace2020)*. arXiv preprint arXiv:2005.04147.
- Melucci, M. (2015). *Introduction to Information Retrieval and Quantum Mechanics*. Springer.
- Metzler, D., & Croft, W. B. (2005). A Markov random field model for term dependencies. In *Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval*, pp. 472–479. ACM.
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv preprint, arXiv:1301.3781*.
- Mitchell, J., & Lapata, M. (2008). Vector-based models of semantic composition.. In *ACL*, pp. 236–244.
- Mitchell, J., & Lapata, M. (2010). Composition in distributional models of semantics. *Cognitive science*, 34(8), 1388–1429.
- Mizraji, E., et al. (1994). Modalities in vector logic.. *Notre Dame Journal of Formal Logic*, 35(2), 272–283.
- Moreira, C., Tiwari, P., Pandey, H. M., Bruza, P., & Wichert, A. (2020). Quantum-like influence diagrams for decision-making. *Neural Networks*, 132, 190–210.
- Neelakantan, A., Shankar, J., Passos, A., & McCallum, A. (2014). Efficient non-parametric estimation of multiple embeddings per word in vector space. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Doha, Qatar. Association for Computational Linguistics.
- Nielsen, M. A., & Chuang, I. (2002). *Quantum computation and quantum information*..
- Penrose, R. (1999). *The emperor’s new mind: Concerning computers, minds, and the laws of physics*. Oxford University Press.
- Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., & Zettlemoyer, L. (2018). Deep contextualized word representations. *arXiv preprint, arXiv:1802.05365*.
- Plate, T. A. (2003). *Holographic Reduced Representations: Distributed Representation for Cognitive Structures*. CSLI Publications.

- Pothos, E. M., & Busemeyer, J. R. (2013). Can quantum probability provide a new direction for cognitive modeling?. *Behavioral and brain sciences*, *36*(3), 255–274.
- Rieffel, E. G., & Polak, W. H. (2014). *Quantum Computing: A Gentle Introduction (Scientific and Engineering Computation)*. The MIT Press.
- Sadrzadeh, M., Clark, S., & Coecke, B. (2013). The Frobenius anatomy of word meanings i: subject and object relative pronouns. *Journal of Logic and Computation*, *23*(6), 1293–1317.
- Sadrzadeh, M., Clark, S., & Coecke, B. (2014). The Frobenius anatomy of word meanings ii: possessive relative pronouns. *Journal of Logic and Computation*, *26*(2), 785–815.
- Sadrzadeh, M., Kartsaklis, D., & Balkır, E. (2018). Sentence entailment in compositional distributional semantics. *Annals of Mathematics and Artificial Intelligence*, *82*(4), 189–218.
- Salton, G., & McGill, M. (1983). *Introduction to modern information retrieval*. McGraw-Hill, New York, NY.
- Schütze, H. (1998). Automatic word sense discrimination. *Computational Linguistics*, *24*(1), 97–124.
- Selinger, P. (2007). Dagger compact closed categories and completely positive maps. *Electronic Notes in Theoretical computer science*, *170*, 139–163.
- Selinger, P. (2010). A survey of graphical languages for monoidal categories. In *New structures for physics*, pp. 289–355. Springer.
- Sloan, P.-P., Kautz, J., & Snyder, J. (2002). Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments. In *Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '02*, p. 527–536, New York, NY, USA. Association for Computing Machinery.
- Smolensky, P. (1986). Information processing in dynamical systems: Foundations of harmony theory. Tech. rep., Colorado Univ at Boulder Dept of Computer Science.
- Smolensky, P. (1990). Tensor product variable binding and the representation of symbolic structures in connectionist systems. *Artificial intelligence*, *46*(1), 159–216.
- Socher, R., Huval, B., Manning, C. D., & Ng, A. Y. (2012). Semantic compositionality through recursive matrix-vector spaces. In *EMNLP 2012, EMNLP-CoNLL '12*, pp. 1201–1211. Association for Computational Linguistics.
- Sordani, A., Nie, J.-Y., & Bengio, Y. (2013). Modeling Term Dependencies with Quantum Language Models for IR. In *Proceedings of the 36th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '13*, pp. 653–662.
- Spencer, A. J. M. (2004). *Continuum mechanics*. Courier Corporation.
- Steinbrecher, G. R., Olson, J. P., Englund, D., & Carolan, J. (2019). Quantum optical neural networks. *NPJ Quantum Information*, *5*(1), 1–9.
- Strang, G. (1993). *Introduction to linear algebra*, Vol. 3. Wellesley-Cambridge Press Wellesley, MA.
- Switzer, P. (1965). Vector images in document retrieval. In *Statistical association methods for mechanized documentation*, pp. 163–171.
- Tarrataca, L., & Wichert, A. (2011a). Problem-solving and quantum computation. *Cognitive computation*, *3*(4), 510–524.
- Tarrataca, L., & Wichert, A. (2011b). Tree search and quantum computation. *Quantum Information Processing*, *10*(4), 475–500.
- Trabelsi, C., Bilaniuk, O., Zhang, Y., Serdyuk, D., Subramanian, S., Santos, J. F., Mehri, S., Rostamzadeh, N., Bengio, Y., & Pal, C. J. (2018). Deep complex networks. In *ICLR 2018*.
- Turney, P. D. (2012). Domain and function: A dual-space model of semantic relations and compositions.. *J. Artif. Intell. Res.(JAIR)*, *44*, 533–585.

- Turney, P. D., Pantel, P., et al. (2010). From frequency to meaning: Vector space models of semantics. *Journal of artificial intelligence research*, 37(1), 141–188.
- Van Rijsbergen, C. J. (1986). A non-classical logic for information retrieval. *The Computer Journal*, 29(6), 481–485.
- Van Rijsbergen, C. J. (2004). *The Geometry of Information Retrieval*. Cambridge University Press.
- Varadarajan, V. S. (1985). *Geometry of Quantum Theory*. Springer-Verlag.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008.
- Vilnis, L., & McCallum, A. (2015). Word representations via gaussian embedding. In *International Conference on Learning Representations*. arXiv preprint arXiv:1412.6623.
- Wang, K., Xiao, L., Yi, W., Ran, S.-J., & Xue, P. (2020). Quantum image classifier with single photons. *arXiv preprint, arXiv:2003.08551*.
- Westphal, J., & Hardy, J. (2005). Logic as a vector system.. *Journal of Logic & Computation*, 15(5).
- Wichert, A. (2020). *Principles of quantum artificial intelligence: Quantum Problem Solving and Machine Learning (Second Edition)*. World scientific.
- Widdows, D. (2003). Orthogonal negation in vector spaces for modelling word-meanings and document retrieval. In *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics (ACL)*, Sapporo, Japan.
- Widdows, D. (2004). *Geometry and Meaning*. CSLI Publications.
- Widdows, D. (2008). Semantic vector products: Some initial investigations. In *Proceedings of the Second International Symposium on Quantum Interaction*.
- Widdows, D., & Bruza, P. (2007). Quantum information dynamics and open world science. In *Quantum Interaction: Papers from the 2007 AAAI Spring Symposium (Technical Report SS-07-08)*, pp. 126–133. AAAI Press.
- Widdows, D., & Cohen, T. (2015). Reasoning with vectors: a continuous model for fast robust inference. *Logic Journal of IGPL*, 23(2), 141–173.
- Widdows, D., & Peters, S. (2003). Word vectors and quantum logic. In *Proceedings of the Eighth Mathematics of Language Conference*, Bloomington, Indiana.
- Wiebe, N., Bocharov, A., Smolensky, P., Troyer, M., & Svore, K. M. (2019). Quantum language processing. *arXiv preprint, arXiv:1902.05162*.
- Wiebe, N., Kapoor, A., & Svore, K. M. (2014). Quantum deep learning. *arXiv preprint, arXiv:1412.3489*.
- Willmore, T. J. (1959). *An introduction to differential geometry*. Oxford University Press, republished Dover (2012).
- Wittek, P. (2014). *Quantum machine learning: what quantum computing means to data mining*. Elsevier.
- Ying, M. (2010). Quantum computation, quantum theory and AI. *Artificial Intelligence*, 174(2), 162–176.
- Zhong, H.-S., Wang, H., Deng, Y.-H., Chen, M.-C., Peng, L.-C., Luo, Y.-H., Qin, J., Wu, D., Ding, X., Hu, Y., et al. (2020). Quantum computational advantage using photons. *Science*, 370(6523), 1460–1463.