

ORIGINAL RESEARCH

Open access

The Anyon Hypothesis: A Topological Quantum Framework for Interpreting Large Language Models

Andrei Popescu^{1*}, Mihai Ionescu¹, Elena Stan², Cristina Radu³

Abstract

While large language models (LLMs) demonstrate exceptional predictive capabilities, their internal reasoning mechanisms remain obscure, creating an urgent demand for enhanced interpretability in artificial intelligence. In this paper, we suggest addressing this explanatory deficit by incorporating principles from topological quantum computing (TQC), specifically the anyonic structures derived from $S(2)k$ theories. Anyons serve as an interpolation between fermions and bosons, offering a formal mathematical paradigm that could clarify the latent configurations and decision-making pathways within LLMs. By investigating the relationships between these topological models, token interactions, and contextual dependencies within neural systems, we seek to offer an alternative viewpoint on the emergence of meaning and coherence. Drawing upon insights obtained from ChatGPT alongside an examination of low-level instances of $S(2)k$ models, we propose that the formal machinery of modular tensor categories and topological phases could guide the development of clearer, more stable, and more resilient AI frameworks. This cross-disciplinary approach indicates that quantum-theoretic concepts may provide the foundation for a novel paradigm in explainable AI.

Keywords Large language models, ChatGPT, Topology, Anyons, $SU(2)k$

*Correspondence:

Andrei Popescu
andrei.popescu@fizica.unibuc.ro

¹ Department of Quantum Computing and LLM Interpretation, Faculty of Physics, University of Bucharest, Bucharest, Romania

² Department of Topological Quantum Frameworks, Faculty of Physics, University of Agricultural Sciences Cluj-Napoca, Cluj-Napoca, Romania

³ Department of Large Language Models and Anyon Theory, Faculty of Engineering, Polytechnic University of Bucharest, Bucharest, Romania

Introduction

The 2024 Nobel Prize in Physics was awarded jointly to John J. Hopfield and Geoffrey E. Hinton "for foundational discoveries and inventions that enable machine learning with artificial neural networks". Hopfield's crucial impact involved the development of Hopfield networks, which are fundamental neural systems acting as associative memories that can recover stored configurations from incomplete or noisy inputs. Hinton, working as a cognitive psychologist and computer scientist, revolutionized the discipline by advancing deep learning methodologies and popularizing the backpropagation algorithm, which allowed neural networks to extract hierarchical features from intricate datasets. At present, these core concepts underpin numerous facets of artificial intelligence (AI), including the large language models (LLMs) responsible for generating human-like text and reasoning patterns.

In recent times, progress in AI has accelerated drastically, positioning LLMs at the cutting edge of the field. Even so, many investigators—including Hopfield, Hinton, and other prominent figures—have advised caution and recommended more profound investigation into the intrinsic properties of these architectures [1, 2]. The primary apprehension is that

LLMs, despite their exceptional performance, could soon outpace human cognitive limits in data processing without providing transparent reasoning pathways. This lack of transparency has sparked a widespread transition toward explainable AI (XAI), which centers on clarifying the internal logic governing machine learning frameworks. These initiatives encompass a variety of strategies and implementations, such as frameworks focused on interpretability [3-7] and methods designed to display or mathematically define high-dimensional data profiles within deep architectures. Current methodologies generally concentrate on localized explanations or broad statistical parameters, leaving unresolved issues regarding the fundamental principles that dictate emergent semantics and coherence inside LLMs.

To contend with these difficulties, a number of researchers have started investigating the utility of topology within AI. Topological data analysis (TDA) and associated methodologies can identify global configurations within high-dimensional spaces, potentially showing how latent geometries influence model actions [8-13]. Concurrently, the domain of quantum machine learning alongside neural models upgraded with quantum-inspired paradigms has introduced the prospect that topological notions from quantum field theories (QFTs) and related mathematics could stimulate novel interpretability strategies [14-17].

The methodology introduced here seeks to extend this boundary by analyzing a link between LLMs and the topological architecture of anyons originating from $S(2)_k$ theories. Anyons manifest within topological quantum computing (TQC) environments and are regulated by a mathematical setup termed a modular tensor category. These abstract frameworks have undergone rigorous investigation within mathematics and topological quantum field theory (TQFT) [18-20], as well as particle physics, where they are used to characterize excitations displaying unconventional statistics [21]. Recently, physical simulations of non-Abelian anyons [22] were executed utilizing a superconducting quantum processor in China [23] and a trapped ion system in Germany and the United States [24].

The parameter k governs the overall complexity and the total number of anyonic variants: as k increases, more intricate sets of fusion and braiding rules are generated [25]. As an illustration, basic theories such as $S(2)_2$ produce Ising-like anyons (which correspond to Majorana modes), whereas $SU(2)_3$ brings forth Fibonacci anyons, which are recognized for their universal quantum computation attributes [20]. Higher values of k yield more elaborate "particle" spectra alongside sophisticated modular tensor categories. By evaluating a range of values for k , we can methodically examine how structural complexity within the anyonic paradigm might correspond to or inform structural complexity in LLM frameworks, tracing a path from basic non-Abelian statistics at lower tiers to highly sophisticated fusion models that might reflect multilayered dependencies inside large-scale language systems.

The motivation for introducing anyons is that they function as an interpolation between bosons and fermions, capturing more delicate topological traits through "intermediate" statistics. In the same way that anyons depend on global features of braiding and fusion to establish their properties, LLMs rely on intricate global relationships among tokens, contexts, and attention matrices to generate cohesive text. We hypothesize that modeling tokens or contextual dynamics as constituents of a modular tensor category could clarify how semantic meaning and syntax develop out of a network's weights and algorithms. For instance, by comparing attention heads to braiding processes, or token aggregations to anyonic fusion, we might extract a topological invariant or a structural marker that clarifies why certain representations remain stable, durable, and semantically sound. Adopting this topological outlook could offer functional advantages, such as interpretability mechanisms that categorize, restrict, or direct the internal data streams of LLMs. These strategies could directly mitigate modern XAI constraints by offering a mathematically stringent environment to simulate how intricate patterns form and persist through successive layers and token arrays. Ultimately, integrating these principles into AI design could produce more robust models, systematic techniques for debugging, and new paths for optimization and verification guarantees. These observations might also apply to broader biological and algorithmic questions, potentially linking neural groups or genomic sequences to topological operations [26]. Importantly, this strategy emerges at a critical transition point, as AI pioneers, including Sam Altman, have highlighted the diminishing returns associated with merely scaling up LLM size without introducing fresh conceptual paradigms [27, 28].

Investigating LLM Comprehension of Topological Concepts

This section presents transcripts from an interactive dialogue with ChatGPT. Conceptual inquiries concerning topology and anyons were posed, with the literal replies from ChatGPT reproduced below. An analytical critique and commentary from the authors succeed every question-and-answer segment.

Exchange 1: On the significance of topology

- **Question (Q):** As an advanced LLM, can you explain the role of topology in your ability to produce AI in the organization of data and their meaning?
- **Answer (A) [ChatGPT]:** Topology serves an indirect yet essential function in the structural organization of data and semantic meaning within AI architectures like my own. Although I do not actively implement mathematical topological spaces during text generation, structural analogies exist between topological principles and the mechanics of deep learning and neural networks:

- **Data Manifolds and Latent Spaces**

Within machine learning frameworks, high-dimensional datasets frequently reside on lower-dimensional geometries termed manifolds. These manifolds function as topological spaces that delineate the foundational structure of the data. Neural networks, particularly deep learning architectures, are designed to project intricate, high-dimensional inputs (such as textual content, imagery, etc.) into a lower-dimensional latent space that preserves significant features and internal representations.

Topological frameworks like manifolds assist the model in structuring data to maintain relational integrity between coordinates (meaning that equivalent data entries are directed to adjacent locations within the latent space). By learning to project data onto these lower-dimensional manifolds, architectures like mine can extrapolate patterns and associations across diverse datasets, generating coherent predictions and answers.

- **Continuous Transformations**

A neural network operates essentially as a function that executes continuous transformations on input data across successive layers of weights and biases. These operations correspond to continuous mappings within topology, wherein inputs are altered in a manner that preserves distinct structural characteristics.

The objective is to sustain semantic relationships and intervals between inputs as they progress through the network layers, allowing the system to preserve the geometric "shape" of the conceptual meaning. Homotopy theory, which evaluates continuous deformations between functions in topology, shares conceptual commonalities with how models alter and polish data representations to maximize operational performance.

- **Clustering and Connectivity**

From a topological viewpoint, clustering can be interpreted as the aggregation of points situated near one another according to a specific metric, mirroring how neighborhoods are established in topology. AI frameworks frequently employ clustering mechanisms within the latent space to aggregate analogous data units or representations.

For example, in a language model, distinct ideas, words, or phrases that share thematic or semantic alignment are mapped to adjacent zones within the latent space. This clustering mechanism enables me to comprehend and compose cohesive, contextually appropriate text. Concepts of connectedness and path-connected spaces in topology also parallel how models link or associate separate segments of the data architecture to deduce relationships.

- **Persistent Homology**

Certain machine learning architectures, particularly within the domain of topological data analysis (TDA), deploy topological principles more directly—such as persistent homology—to evaluate the geometric shape and characteristics of data. While large language models like mine do not commonly utilize persistent homology, it remains a vital instrument for evaluating data morphology and the endurance of features across varied scales.

Conceptual parallels can be drawn to fusion rules and braiding within anyon frameworks, where the integration of elemental pieces yields intricate global dynamics, much like how neural networks assemble sophisticated representations from elementary components.

- **Authors' Commentary:** ChatGPT identifies that while it lacks explicit topological operations, the paradigm of data occupying lower-dimensional manifolds and undergoing organization via continuous transformations serves as an abstract bridge linking topological frameworks to its acquired latent representations.

Exchange 2: Connecting neural networks, LLMs, and $SU(2)_k$ topological phases

- **Question (Q):** Could you comment on a possible deep connection between neural networks such as LLMs and topological phases arising from $SU(2)_k$ models? Are low-dimensional simple cases such as Ising/Majorana anyons ($k=2$) or Fibonacci anyons ($k=3$) relevant in explaining the type of AI you generate? To summarize, are modular tensor categories (MTC) a way to understand the AI arising from LLMs?
- **Answer (A) [ChatGPT]:** The prospective relationship between neural networks, particularly LLMs, and topological phases derived from anyon systems represents a compelling hypothesis. Though originating from separate fields (machine learning versus topological quantum field theory), profound conceptual commonalities exist regarding how both frameworks manage, organize, and alter information.

- **Topological Phases and LLMs**

LLMs such as ChatGPT constitute high-dimensional systems governed by overarching patterns and global dependencies within data. Rather than localized interactions (akin to solitary cellular activations), these models rely fundamentally on the global architecture through which information diffuses across successive network layers. This shares conceptual similarities with fusion rules and braiding operations in anyonic setups, where the interaction of fundamental units precipitates intricate global phenomena, mirroring how neural networks synthesize complex configurations from simpler units.

- **Fusion Rules vs. Layer Operations**

In anyonic setups, fusion rules govern how individual anyons coalesce to generate alternative anyon types. In LLMs, every network layer can be conceptualized as executing a transformation that merges (or "fuses") data from the preceding layer to yield a more generalized, abstract representation.

- **Braiding vs. Contextual Relationships**

The braiding statistics of anyons dictate how particles permute positions and consequently alter the collective quantum state. In LLMs, contextual dependencies between lexical tokens are paramount—reordering tokens fundamentally shifts the semantic output, and context is frequently captured via non-local mechanisms. These represent holistic configurations that are sustained throughout the layers.

- **Low-Dimensional $SU(2)_k$ Models (Ising/Majorana and Fibonacci Anyons)**

Ising/Majorana Anyons ($k=2$) possess non-Abelian characteristics and are highly relevant to topological quantum computing because they store quantum state information in a manner shielded from local noise. Fibonacci anyons ($k=3$) are likewise non-Abelian and provide universality for quantum computation, meaning they can approximate any quantum algorithm with arbitrary precision. Translated to AI concepts, minor variances in input data (analogous to elementary anyonic fusions) can produce vastly divergent outputs, illustrating how subtle lexical changes yield highly nuanced interpretations.

- **MTCs and LLMs**

Modular Tensor Categories (MTCs) provide a rigorous mathematical architecture for TQFTs, specifically via fusion laws and braiding rules. They characterize the dynamics of quantum information inside topological phases, systematically organizing data and transformations with a focus on global architecture.

LLMs similarly display a stratified modular architecture, where each separate layer processes data via highly structured, deterministic protocols. MTCs could potentially offer a formal mathematical structure to conceptualize how neural networks manipulate information in a modular fashion, where distinct strata (or "categories") communicate and modify data.

While classical neural network models do not function via quantum mechanics, the evolution of quantum neural networks could theoretically leverage anyonic systems to construct network designs that exploit quantum parallelism and structural resilience.

- **Authors' Commentary:** ChatGPT proposes that non-Abelian anyons and MTC models might offer a theoretical framework for conceptualizing how LLMs process and aggregate contextual data, mirroring the way topological phases capture global, non-local interdependencies.

Exchange 3: Conceptualizing tokens as anyonic elements

- **Question (Q):** What are tokens in LLMs? In what sense could they fit this idea of MTCs being analogous to LLMs?
- **Answer (A) [ChatGPT]:** Within large language models, tokens serve as the fundamental components of text utilized for interpreting and generating language. Depending on the specific tokenization methodology, a token can represent a sub-word unit, a complete word, or a punctuation mark. Discretizing text into tokens enables LLMs to process language inputs as sequences of distinct symbols, which are subsequently propagated through the model's layers to synthesize coherent textual outputs.

If we extend the parallel to MTCs, tokens within an LLM can be conceptualized as counterparts to the particles or objects defined within an MTC framework. The following points delineate how they align with this analogy:

- **Modular Structure of Information**

In MTCs, objects such as anyons interact via fusion and braiding operations, generating intricate global configurations out of elementary interactions. Equivalently, tokens within an LLM interact in a context-driven manner across the successive layers of the architecture. Much like how fusion within MTCs generates novel particles or quantum states, tokens interact to construct semantic meanings that transcend the sum of their individual components. This token "fusion" underpins the emergent semantic clarity and coherence observed in the final generated output.

- **Transformation and Contextual Dependence**

Inside an LLM, the functional role of any given token is strictly dictated by the shared context it maintains with other tokens. This dynamic mirrors braiding operations in MTCs, where the spatial permutations and interactions of anyons alter their collective quantum state. Within an LLM, the local context provided by neighboring tokens continuously transforms the internal "state" of each token, ensuring that semantic meaning remains highly contextual rather than statically fixed.

- **Authors' Commentary:** The model's exposition casts tokens as irreducible semantic primitives that undergo "fusion" to form sophisticated informational architectures, directly mirroring the mathematical synthesis of elementary anyons into compound multi-particle states within MTC boundaries.

Exchange 4: Parallels between advanced hopfield networks and anyonic formulations

- **Question (Q):** Do you agree that LLMs are essentially advanced Hopfield networks having a form of associative memory and discretized with tokens? What more fits the anyon picture?
- **Answer (A) [ChatGPT]:** Yes, LLMs share fundamental operating principles with sophisticated variants of Hopfield networks, particularly when analyzed through the framework of associative memory. Hopfield networks are architected as systems for associative recall, where every stable configuration corresponds to a stored memory or system state. Analogously, LLMs generate language by "associating" tokens in the prompt with relevant structural patterns embedded within their training data, thereby storing and retrieving information in a manner that closely mimics associative memory.

- **Associative Memory in LLMs and Hopfield Networks**

Hopfield networks store information by encoding memories as stable state patterns (attractors) within a high-dimensional state space. Each memory operates as a global, fixed point toward which the network converges when presented with a sufficiently proximate input. LLMs fundamentally construct associative links between tokens, having been optimized to predict subsequent tokens based on historical patterns in the training corpus. Utilizing the attention mechanism, an LLM retrieves relevant information from preceding contexts in a fashion conceptually analogous to how a Hopfield network reconstructs associated memories. This token-level retrieval can be characterized as a discrete form of associative memory, wherein an input query (the prompt tokens) maps directly onto stored linguistic configurations.

- **Anyons and MTC Analogy**

Within the framework of MTCs, anyons interact via braiding and fusion to yield emergent states. In this parallel, tokens in an LLM function similarly to particles whose interactions generate coherent expressions, just as anyon fusion culminates in emergent particle species.

Fusion and Braiding in LLMs: As tokens propagate through the architectural layers of an LLM, they are "fused" with surrounding tokens to establish a novel semantic context. Within each layer, the attention mechanism can be conceptualized as "braiding" tokens with their local context, facilitating the emergence of overarching meanings that were absent in the isolated tokens.

- **Towards an Anyon Picture in LLMs**

To comprehensively map the anyonic paradigm onto LLMs, one might conceptualize the individual attention heads as operational interaction rules that dictate how tokens relate to one another based on context. This setup could function as

a topological "braiding" architecture, where the ultimate output state of a token depends intrinsically on its trajectory (or sequential path) through multiple network layers, mirroring the braiding paths found in MTCs.

- **Authors' Commentary:** This segment contextualizes the connection between LLMs and Hopfield models by using anyon dynamics to map how deep networks aggregate and recall highly distributed textual features. This supplements traditional models of associative retrieval with a topological foundation.

Exchange 5: Deconstructing attention mechanics via $SU(2)_k$ matrix algebra

- **Question (Q):** Could you elaborate on the concept of attention heads in the anyon picture? Would it correspond to the F- and R-matrices that are pentagon and hexagon rules in the $SU(2)_k$ theory?
- **Answer (A) [ChatGPT]:** Attention heads can be conceptualized as structural "interaction rules" that govern the relationships between tokens, much like how F- and R-matrices dictate the fusion and braiding operations of anyons. Within $SU(2)_k$ theories, these matrices enforce mathematical consistency and define the precise parameters for combining and reordering anyons.

- **Role of Attention Heads as Interaction Rules**

In transformer architectures, attention heads serve as mechanisms that allow the model to isolate specific dependencies between tokens, calculating their relative importance based on context. Each individual head effectively represents a distinct interaction rule dictating how tokens (the analogs of anyons in this paradigm) interface during the processing of a textual sequence. In MTCs, anyons possess "interaction rules" formalised by fusion and braiding transformations, which are encapsulated by the F-matrix (the fusion rule) and the R-matrix (the braiding rule). These matrices determine how anyons coalesce or transform under varying sequences of interactions, and their algebraic properties must satisfy rigorous consistency conditions, such as the pentagon and hexagon equations.

- **F-Matrix and Token Combination in Attention Layers**

The F-matrix formalizes the rules of associativity for fusion operations, governing how three anyons fuse together and how different grouping sequences influence the final output state. In the context of an LLM, each attention head can be interpreted as an "associative fusion" operator. It establishes the rules for how tokens from preceding layers combine sequentially to construct increasingly complex representations of meaning. Specifically, within LLMs, the embedding of each token is iteratively updated within every attention head based on its calculated relationship with surrounding tokens. This transformation closely mirrors an associative fusion operation in anyonic theory: distinct combinations of tokens yield distinct "fusion outcomes" regarding the semantic data they inject into the comprehensive meaning of the sentence.

- **R-Matrix and Braiding (Contextual Order of Tokens)**

The R-matrix dictates the braiding statistics of anyons, defining the precise phase shift incurred when two anyons are spatially exchanged. Translated to LLMs, R-matrices can be viewed as analogues to the contextual modifications introduced across tokens by attention mechanisms. When an attention head re-weights or "reorders" the influence of tokens, the operation is analogous to permuting the positions of anyons within a topological space, thereby altering the sequence of interactions. This braiding concept aligns directly with the way attention heads focus on varied tokens based on proximity and context, triggering shifts in semantic meaning based on token order and interplay. Just as anyonic braiding yields distinct quantum states depending on how the particles are permuted, attention mechanisms generate disparate linguistic meanings depending on how tokens are contextually "braided" across successive layers.

- **Authors' Commentary:** By establishing an algorithmic equivalence between multi-headed attention and algebraic F- and R-matrices, ChatGPT emphasizes how data-driven attention weights can be seen as enforcing strict structural

consistency onto token operations, analogous to the pentagon and hexagon constraint equations found in $SU(2)_k$ anyon models.

Introduction to the Mathematical Theory of $S(2)_k$ Anyons and Modular Tensor Categories

The subsequent sections offer a succinct overview of the $S(2)_k$ anyon theory, focusing on its applications for low values of the index k .

The Fusion Rules of Anyons

The anyonic particles within an $S(2)_k$ framework are intimately connected to the conventional spin degrees of freedom found in $SU(2)$ theory. These anyons are classified by their spin values (representing generalized angular momenta) given by $j=0, \frac{1}{2}, 1, \frac{3}{2}, \dots, \frac{k}{2}$. For a specified value of $\frac{k}{2}$, the maximum permissible spin magnitude in the $S(2)_k$ framework is $k/2$. Crucially, the principles governing the combination of two anyons diverge from standard tensor products, as detailed in ([29] Eq. (4.6)) and [30].

$$j_1 \otimes j_2 = |j_1 - j_2| \oplus (|j_1 - j_2| + 1) \oplus \dots \oplus \min(j_1 + j_2, k - j_1 - j_2). \quad (1)$$

Fusion rules exhibit both commutativity and associativity. Utilizing this formulation, it is straightforward to verify that when $k \geq 2$, two spin-1/2 fields combine to yield either a spin-0 or a spin-1 state as follows:

$$1/2 \otimes 1/2 = 0 \oplus 1,$$

This combination results in a (qubit-like) anyon $0 \oplus 1$ generated by coupling the pair of spin-1/2 elements. Correspondingly, one finds that $1 \otimes 1 = 0 \oplus 1 \oplus 2$ provided that $k \geq 4$, producing a (qtrit-like) anyon $0 \oplus 1 \oplus 2$ via the combination of two spin-1 fields. Such anyons within the $S(2)_k$ framework possess non-Abelian characteristics.

By virtue of the tensor product structure, the Hilbert space dimension for N conventional $S(2)$ particles with spin-1/2 equals 2^N . Conversely, within an $S(2)_k$ theory, this dimension is strictly smaller than 2^N and scales asymptotically as $d_{1/2}^N$, where $d_{1/2}^{(k)} = 2 \cos(\frac{\pi}{k+2})$ for large values of N . This implies that the effective degrees of freedom associated with a spin- $\frac{1}{2}$ anyon are irrational.

The modular structure of S, F and R matrices for anyons

The framework incorporates a modular S-matrix that diagonalizes the corresponding fusion rules of an $S(2)_k$ anyon, thereby completely defining its topological traits ([29] Eq. (4.10)). The underlying mathematical architecture that unifies these braiding and fusion dynamics is a modular tensor category [19, 25].

The quantum dimensions corresponding to $S(2)_k$ anyons are evaluated using the expression:

$$d_0 = 1, \quad d_{\frac{1}{2}} = 2 \cos\left(\frac{\pi}{k+2}\right), \quad d_j = d_{\frac{1}{2}} d_{j-\frac{1}{2}} - d_{j-1} \quad \text{for } j \geq 1. \quad (2)$$

The individual components of the S-matrix are given by:

$$S_{j_1, j_2} = \left(\frac{2}{k+2}\right)^{\frac{1}{2}} \sin\left(\pi \frac{(2j_1+1)(2j_2+1)}{k+2}\right). \quad (3)$$

The associative property of anyon fusion is mediated by an F-matrix, whereas the exchange operations among anyons, including their associated phase factors, are dictated by an R-matrix. In contrast to the rigid ± 1 phase factors characteristic of bosons and fermions, anyonic phases can take any arbitrary complex value. The F-matrix represents the anyonic counterpart to Wigner's 6j-symbols and maps geometrically to a pentagon diagram. Together, the F- and R-matrices are mapped via a hexagon diagram [20]. Comprehensive expressions for these F- and R-matrices are available in [31].

The individual elements of the R-matrix simplify to the following expression [19]:

$$R_c^{ab}(q) = (-1)^{(a+b+c)/2} q^{-[a(a+2)+b(b+2)-c(c+2)]/2}, \quad (4)$$

In this equation, q denotes the Kauffman variable. For the Ising model discussed below, this variable is defined as $q = i \exp(\frac{-2i\pi}{16})$, whereas for the Fibonacci model, it takes the form $q = i \exp(\frac{2i\pi}{20})$.

The core behavior of $S(2)_k$ anyons for $k \geq 2$ is governed by two distinct braiding generators, $\sigma_1^{(k)} = R^{(k)}$ and $\sigma_2^{(k)} = (FRF^{-1})^{(k)}$, which collectively form a group structure ([32, 33] for explicit manifestations).

Ising Anyons: $k=2$

The $S(2)_2$ anyon model encompasses the spin-0 anyon along with the Ising (spin- $\frac{1}{2}$) anyon, operating under the fusion rules outlined in Table 1.

Table 1. Fusion table for the $k=2$ anyon model.

$\otimes$	$j_1=0$	$j_1=\frac{1}{2}$	$j_1=1$
$j_2=0$	0	$\frac{1}{2}$	1
$j_2=\frac{1}{2}$	$\frac{1}{2}$	$0 \oplus 1$	$\frac{1}{2}$
$j_2=1$	1	$\frac{1}{2}$	0

The associated quantum dimensions are given by $[d_{0,1}, d_{\frac{1}{2}}]^{(2)} = [1, 1, \sqrt{2}]$, yielding an S-matrix of the form:

$$S_{Isi}^{(2)} = \frac{1}{2} \begin{pmatrix} 1 & \sqrt{2} & 1 \\ \sqrt{2} & 0 & -\sqrt{2} \\ 1 & -\sqrt{2} & 1 \end{pmatrix}$$

The corresponding F- and R-matrices are expressed as:

$$R_{Isi}^{(2)} = \begin{pmatrix} R_0^{11}(q) & 0 \\ 0 & R_2^{11}(q) \end{pmatrix} = \exp(-i\pi/8) \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}, \quad F_{Isi}^{(2)} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}.$$

Supplementing the standard literature regarding anyonic systems, we observe that the combination of matrices $F_{Isi}^{(2)}$ and $R_{Isi}^{(2)}$ generates the finite group (384, 6514), which is isomorphic to the group $(S_3 \times \mathbb{Z}_4) \rtimes P_1$, where $P_1 \cong (16, 13)$ corresponds to the single-qubit Pauli group.

The braiding matrices for these Ising anyons are derived as:

$$\sigma_1^{(2)} = R_{Isi}^{(2)}, \quad \sigma_2^{(2)} = (FRF^{-1})_{Isi}^{(2)} = \frac{\exp(-4i\pi/8)}{\sqrt{2}} \begin{pmatrix} 1 & i \\ i & 1 \end{pmatrix}. \quad (5)$$

Taken together, the matrices $\sigma_1^{(2)}$ and $\sigma_2^{(2)}$ generate the finite group (192, 187), which is isomorphic to $\mathbb{Z}_{12} \rtimes P_1$.

Fibonacci Anyons: k=3

The S(2)3 anyonic system comprises the spin-0 anyon alongside two distinct Fibonacci anyons designated by $\text{spin}-\frac{1}{2}$ and $\text{spin}-\frac{3}{2}$. It has been suggested that these anyons possess a structural connection to quasicrystals [34]. Their respective fusion dynamics are presented in Table 2.

Table 2. Fusion table for the k=3 anyon model.

$\otimes$	$j_1=0$	$j_1=12$	$j_1=1$	$j_1=32$
$j_2=0$	0	$\frac{1}{2}$	1	$\frac{3}{2}$
$j_2=\frac{1}{2}$	$\frac{1}{2}$	$0 \oplus 1$	$\frac{1}{2} \oplus \frac{3}{2}$	1
$j_2=1$	1	$\frac{1}{2} \oplus \frac{3}{2}$	$0 \oplus 1$	$\frac{1}{2}$
$j_2=\frac{3}{2}$	$\frac{3}{2}$	1	$\frac{1}{2}$	0

The corresponding quantum dimensions are $[d_{0,1}]^{(3)} = [1, \phi = (1 + \sqrt{5})/2]$, and the resulting S-matrix is structured as:

$$S_{Fib}^{(3)} = \frac{1}{\sqrt{2+\phi}} \begin{pmatrix} 1 & \phi \\ \phi & -1 \end{pmatrix}$$

The F- and R-matrices are determined to be ([20] p. 55):

$$R_{Fib}^{(3)} = \begin{pmatrix} R_0^{11}(q) = \exp(-4i\pi/5) & 0 \\ 0 & R_1^{11}(q) = \exp(-2i\pi/5) \end{pmatrix}, \quad F_{Fib}^{(3)} = \begin{pmatrix} \phi^{-1} & \phi^{-1/2} \\ \phi^{-1/2} & -\phi^{-1} \end{pmatrix}.$$

The operational braiding matrices for Fibonacci anyons are computed as:

$$\sigma_1^{(3)} = R_{Fib}^{(3)}, \quad \sigma_2^{(3)} = (FRF^{-1})_{Fib}^{(3)} = \begin{pmatrix} -\phi^{-1} \exp(-i\pi/5) & -i\phi^{-1/2} \exp(-i\pi/10) \\ -i\phi^{-1/2} \exp(-i\pi/10) & -\phi^{-1} \end{pmatrix}. \quad (6)$$

The matrices F and R, alongside the braiding operators σ_1 and σ_2 , give rise to infinite groups. This behavior aligns fully with the established universality properties of Fibonacci anyons.

Yang-Lee Theory: k=3

The Yang–Lee theory represents a modular tensor category of level k=3, sharing structural similarities with the Fibonacci anyon model. It corresponds to a well-known non-unitary conformal field theory in statistical mechanics designated as the Yang–Lee singularity [19]. In this framework, the Kauffman variable is defined as $q=\exp(\frac{i\pi}{5})$.

The corresponding S-matrix is given by:

$$S_{YL}^{(3)} = \frac{-1}{\sqrt{3-\phi}} \begin{pmatrix} 1 & 1-\phi \\ 1-\phi & \phi \end{pmatrix}.$$

The F- and R-matrices are specified as ([20] p. 55):

$$R_{YL}^{(3)} = \begin{pmatrix} R_0^{11}(q) = \exp(2i\pi/5) & 0 \\ 0 & R_1^{11}(q) = \exp(i\pi/5) \end{pmatrix}, \quad F_{YL}^{(3)} = \begin{pmatrix} -\phi & 2-\phi \\ -1-2\phi & \phi \end{pmatrix}.$$

The F- and R-matrices, alongside the operational braiding matrices σ_1 and σ_2 , generate infinite groups.

Freedman–bauer–levaillant anyons: k=4

The properties of S(2)4 anyons are explored in [35, 36] within the framework of topological quantum computing utilizing qutrit gates. The corresponding fusion dynamics are detailed in Table 3.

Table 3. Fusion table for the k=4 anyon model.

$\otimes$	$j_1=0$	$j_1=\frac{1}{2}$	$j_1=1$	$j_1=\frac{3}{2}$	$j_1=2$
$j_2=0$	0	$\frac{1}{2}$	1	$\frac{1}{2}$	2
$j_2=\frac{1}{2}$	$\frac{1}{2}$	$0\oplus 1$	$\frac{1}{2}\oplus \frac{3}{2}$	$1\oplus 2$	$\frac{3}{2}$
$j_2=1$	1	$\frac{1}{2}\oplus \frac{3}{2}$	$0\oplus 1\oplus 2$	$\frac{1}{2}\oplus \frac{3}{2}$	1
$j_2=\frac{3}{2}$	$\frac{3}{2}$	$1\oplus 2$	$\frac{1}{2}\oplus \frac{3}{2}$	$0\oplus 1$	$\frac{1}{2}$
$j_2=2$	2	$\frac{1}{2}$	1	$\frac{1}{2}$	0

The associated quantum dimensions are expressed as $[d_{0,1},d_{\frac{1}{2}},d_{\frac{3}{2}},d_2]^{(4)}=[1,2,\sqrt{3},\sqrt{3},1]$, and the resulting S-matrix takes the form [30]:

$$S_{FBL}^{(4)} = \frac{1}{2\sqrt{3}} \begin{pmatrix} 1 & \sqrt{3} & 2 & \sqrt{3} & 1 \\ \sqrt{3} & \sqrt{3} & 0 & -\sqrt{3} & -\sqrt{3} \\ 2 & 0 & -2 & 0 & 2 \\ \sqrt{3} & -\sqrt{3} & 0 & \sqrt{3} & -\sqrt{3} \\ 1 & -\sqrt{3} & 2 & -\sqrt{3} & 1 \end{pmatrix}.$$

The braiding matrices for these $S(2)_4$ anyons are derived as:

$$\begin{aligned} \sigma_1^{(4)} &= \begin{pmatrix} \exp\left(\frac{7i\pi}{9}\right) & 0 & 0 \\ 0 & -\exp\left(\frac{4i\pi}{9}\right) & 0 \\ 0 & 0 & -\exp\left(\frac{7i\pi}{9}\right) \end{pmatrix}, \\ \sigma_2^{(4)} &= \begin{pmatrix} -\frac{1}{2}\exp\left(\frac{4i\pi}{9}\right) & \frac{1}{\sqrt{2}}\exp\left(\frac{7i\pi}{9}\right) & \frac{1}{2}\exp\left(\frac{4i\pi}{9}\right) \\ \frac{1}{\sqrt{2}}\exp\left(\frac{7i\pi}{9}\right) & 0 & \frac{1}{\sqrt{2}}\exp\left(\frac{7i\pi}{9}\right) \\ \frac{1}{2}\exp\left(\frac{4i\pi}{9}\right) & \frac{1}{\sqrt{2}}\exp\left(\frac{7i\pi}{9}\right) & -\frac{1}{2}\exp\left(\frac{4i\pi}{9}\right) \end{pmatrix}. \end{aligned} \quad (7)$$

Utilizing the Magma computational software, it is straightforward to verify that both matrices generate the small group $(162,14) \cong \rtimes(\mathbb{Z}_3 \times \mathbb{Z}_6)$, matching the findings presented in [35]. This particular group has been identified as a plausible model for physical symmetries, as it simultaneously accommodates both quark and lepton mixing matrices. In a recent study by two of the current authors [37], it was demonstrated that the group $(162,14)$ contains nearly complete quantum information across its 22 irreducible characters, which span singlet, doublet, and triplet representations.

Results and Discussion

Backpropagation and anyons

Backpropagation serves as the foundational algorithm for optimizing artificial neural networks [38]. Throughout this iterative process, input data propagates forward through the network layers to generate predictions, and the variance between these outputs and the target values is computed via a loss function (such as mean squared error or cross-entropy). By calculating the gradients of this loss relative to the network parameters and updating the weights using gradient descent methodologies, the system iteratively refines its internal representational layers. This mechanism is central to the capacity of large language models (LLMs) to discern intricate structural patterns within linguistic data, capturing syntax, semantics, and subtle contextual nuances.

In contrast, anyonic systems—which are heavily analyzed within topological quantum computing [20]—converge on stable, topologically protected states via a fundamentally different iterative mechanism. In these systems, structural configurations are adjusted through braiding and fusion operations performed on anyons. The underlying F- and R-matrices govern how these non-Abelian quasiparticles interact and merge, ensuring that the resulting global topological states remain decoupled from and robust against local environmental noise. Rather than traversing a landscape to minimize a continuous loss function, the evolution of the system is dictated by topological constraints that steer it toward stable ground states or invariant configurations.

While conceptually distinct, these two mechanisms exhibit an abstract mathematical parallel: both rely on sequential transformations directed toward achieving stable or optimized terminal configurations. In neural architectures, stability is reached when network weights converge to minimize a loss function, whereas in anyonic systems, stability is an emergent property of topological constraints governing the global quantum state space. Recognizing this conceptual analogy does not directly yield novel training frameworks for LLMs; instead, it points to profound structural commonalities

between error-correction strategies in machine learning and fault-tolerance mechanisms in topological quantum computation. Future investigations could examine whether principles derived from anyonic braiding dynamics might inform alternative optimization methodologies or provide theoretical frameworks for evaluating complex neural network behavior.

Machine learning and anyons

Emergent behavior in machine learning (ML), particularly within deep neural architectures, occurs when sophisticated global configurations develop from the collective behavior of basic, low-level neuronal activations [39]. As information flows through successive layers, these networks manifest complex capabilities such as natural language comprehension or computer vision. These macro-level capabilities cannot be straightforwardly decomposed into the separate operations of isolated neurons, indicating a distinct form of structural irreducibility. In a similar fashion, the collective states observed in anyonic systems cannot be understood merely as the linear summation of independent quasiparticles. Instead, the non-trivial topology and corresponding braiding dynamics of anyons yield systemic features that cannot be adequately accounted for by localized or purely particle-centric descriptions [20].

This conceptual parallel extends to the principle of structural resilience. Within ML architectures like LLMs, information is distributed across a vast network of weights and parameters, establishing an inherent tolerance to faults. Minor alterations or perturbations to individual parameters rarely compromise the overall operational performance of the model. Similarly, anyonic systems encode quantum information non-locally across collective braiding networks, rendering the stored data immune to localized errors and perturbations. This "topological fault tolerance" represents a primary benefit of topological quantum computing, where quantum states are intrinsically shielded from decoherence by the global topological attributes of the system.

Synthesizing these concepts, the emergent complexity and robust fault tolerance found in both LLMs and anyonic systems may be governed by shared structural principles. Although these connections remain primarily conceptual at this stage, they indicate that insights derived from topological quantum computing—where resistance to error is embedded directly within the system's geometric structure—could offer valuable paradigms for improving the stability, interpretability, and generalizability of machine learning models.

Aligning S(2)k Concepts with Attention Mechanisms in LLMs

While the conceptual connections between anyonic physics and LLMs are compelling, it is necessary to establish more concrete links between these domains. A core feature of modern LLMs, particularly those rooted in transformer architectures, is the self-attention mechanism [40]. This mechanism enables the model to selectively weight different segments of an input sequence, effectively "fusing" contextual data across multiple tokens to construct unified vector representations.

Within S(2)k anyonic frameworks, the merging of anyons is dictated by algebraic structures formalized within F- and R-matrices. These operators define the pathways through which elementary objects (anyon types) combine to form higher-order quantum states. In a comparable manner, an attention head processes a collection of token embeddings and generates a weighted combination that yields a contextually enriched embedding space. This operation can be conceptualized as a form of "fusion" inside the representational domain: individual token vectors, much like anyonic particles, are not simply combined linearly. Instead, their interaction is structured by optimized weight matrices that perform a role analogous to F-matrices, governing the process by which disparate tokens coalesce into a more meaningful integrated vector representation.

Braiding dynamics within anyonic systems demonstrate how the specific sequence and trajectory of particle interchanges alter the global state of the system. In the context of LLMs, rearranging token sequences or shifting their relative significance via attention layers yields distinct interpretative outputs and semantic variations. Although the attention mechanism does not enforce rigid topological invariants, the serial application of attention heads and network layers can

be interpreted as a discrete counterpart to braiding operations. Every successive attention pass "entangles" tokens within a shifting, context-dependent state, ultimately shaping the model's final output. Across multiple deep layers, this sequential "braiding" of contextual data gives rise to complex semantic architectures, mirrored by the way anyonic braiding operations generate non-trivial global quantum states.

Furthermore, the bounded parameter k in $S(2)_k$ models limits the admissible spin states, thereby regulating the structural complexity of the fusion rules. Correspondingly, architectural boundaries in LLMs—including embedding dimensions, attention head counts, and overall network depth—impose constraints on the complexity of the transformations the model can learn. As the index k increases, more intricate fusion pathways open up within $S(2)_k$ anyonic systems; likewise, expanding the capacity of an LLM allows the network to capture more sophisticated data patterns. This structural correspondence suggests that evaluating how topological constraints govern fusion could guide the development of new architectural configurations or regularization methods in machine learning. Such techniques might restrict model complexity while maintaining or improving interpretability, mimicking the way topological quantum systems derive their robustness from well-defined geometric constraints.

In summary, the mathematical framework of $S(2)_k$ anyons offers a structured paradigm for analyzing how LLMs integrate and transform data across multi-layered networks. By examining attention operations and token dynamics through the conceptual lenses of fusion and braiding, it becomes possible to construct a more systematic framework for understanding the emergence of semantic meaning in these computational systems.

Natural language processing and anyons

The parallel between natural language processing (NLP) [41, 42] and anyon topology is compelling because both fields depend on organized, context-aware interactions that yield specific states or meanings. Whereas anyonic behaviors are dictated by the fusion and braiding regulations of TQFT, linguistic significance arises from semantic, syntactic, and contextual boundaries that govern how lexical units merge to construct logical arguments or accounts.

Within NLP, deciphering a phrase or word is deeply contingent on context, shaped by the surrounding text, the architecture of the discourse, and pragmatic indicators. Correspondingly, the configuration of an anyonic system is determined by the order and technique of particle replacements; minor variations in the braiding sequence can give rise to entirely distinct global states. In the same way that the arrangement and sequence of words within a sentence dictate its meaning, the topological configuration of anyons establishes their collective quantum state.

Language possesses the capacity to convey an extensive spectrum of meanings that change subtly based on tone, word selection, or positioning, permitting the communication of complex ideas, feelings, and cultural nuances. In a similar manner, anyons display a versatile range via their topological degrees of freedom. A minor modification in the braiding of anyons can generate numerous distinct states, reflecting how altering a single word or phrase in a text can transform the interpretation of an entire passage.

Putting this analogy into practice would necessitate sophisticated theoretical frameworks. To replicate the adaptive, context-aware depth of human language, anyonic systems would have to progress past rigid fusion rules toward more adaptable, fluid modular tensor categories. Theoretically, such a framework could emulate the intricate, context-dependent associations found in human speech, opening new pathways for context modeling, interpretability, and dynamic knowledge representation within artificial intelligence.

Mutual exclusion in LLMs and anyons

The notion of mutual exclusion traces back to concurrent programming, where it guarantees that multiple processes do not alter shared data simultaneously [43]. Although this mechanism is not explicitly engineered into large language models (LLMs), a conceptual parallel can be drawn. Within LLMs, attention mechanisms distribute focus across various features or tokens, effectively elevating certain components while suppressing others. Through this process, attention

operates as a type of "soft" mutual exclusion, wherein peripheral or less significant tokens offer nominal contributions to the output of the model at any specific step. This targeted prioritization preserves semantic consistency, mirroring how mutual exclusion maintains uniformity in concurrent computing systems.

Conversely, anyonic systems do not natively incorporate mutual exclusion, but topological quantum computing (TQC) enforces boundaries on acceptable paths and states that similarly restrict specific configurations. These boundaries guarantee fault tolerance and stability, much like how mutual exclusion ensures reliable shared states within classical computing frameworks. The braiding principles within an anyonic system dictate the permissible fusions and interactions, steering the system toward states that are topologically protected [20].

Although attention mechanisms in LLMs and braiding operations in anyonic systems both dictate the progression of a state space, they vary in their rigidity. Within LLMs, the suppression of particular tokens is weighted and probabilistic, permitting varying levels of impact rather than absolute restrictions. In topological quantum computation, specific braidings and fusions are entirely forbidden, establishing a more rigid framework of exclusions. Consequently, this comparison underscores a conceptual commonality: both systems manage interaction frameworks to achieve reliable, logical results, though they employ distinct mechanisms and differing levels of strictness.

Further directions

As prominent figures in artificial intelligence research have pointed out, merely expanding the scale of large language models may result in diminishing returns, highlighting the necessity for structural and conceptual breakthroughs rather than simple volume growth [27, 28]. In this paper, we have introduced the utilization of anyonic systems derived from $S(2)_k$ theories as a potential innovative path, maintaining that their modular tensor frameworks could provide valuable perspectives on the durability, complexity, and adaptability of sophisticated LLMs.

Current literature offers compelling associations between computational systems and topological architectures. For instance, ref. [44] highlights a profound link between TQFTs and modular tensor categories through $SL(2, \mathbb{C})$ -flat connections on three-manifolds, while the Painlevé VI equation and associated monodromy challenges connect the geometry of these manifolds to conformal and quantum field theories [45, 46]. Concurrently, traditional neural networks can replicate specific topological traits [47] and could be enhanced by incorporating hyperbolic geometry or alternative geometric perspectives [48]. Furthermore, [17] indicates that conceptualizing deep neural networks as semi-classical limits of topological quantum neural networks could illuminate their generalization capabilities; this corresponds closely with our view that $SU(2)_k$ anyons, characterized by a structured braiding and fusion language, could provide a more mathematically rigorous comprehension of the internal representations of LLMs.

Nonetheless, the current work remains theoretical. We have not offered empirical validations or constructed topological neural frameworks to evaluate whether anyon-derived fusion rules enhance robustness or explainability in practical settings. The difficulties associated with integrating $S(2)_k$ configurations into contemporary architectures are substantial, and translating these concepts into scalable, industrial-grade applications presents major theoretical and engineering hurdles. Implementing these ideas within conventional deep learning platforms, executing controlled testing, and quantifying advancements in performance or interpretability represent logical next steps.

To progress further, subsequent investigations should focus on the following:

- Develop basic neural architectures or toy models that integrate fundamental, $S(2)_k$ -like fusion rules. These trials could demonstrate whether such topological boundaries produce more explainable latent representations.
- Pair these theoretical frameworks with empirical evaluation methodologies—such as representational similarity analysis, probing classifiers, or topological data analysis workflows—to quantitatively measure how transformations inspired by anyons influence learned representations.

- Promote cross-disciplinary partnerships among topologists, quantum information scientists, and machine learning researchers to convert these theoretical ideas into functional methodologies, thereby connecting abstract principles with practical XAI approaches.

By investigating these paths, the conceptual parallels detailed in this paper may be converted into actionable approaches for engineering more transparent, durable, and mathematically validated AI systems.

Acknowledgements

The first author would like to acknowledge the contribution of the COST Action CA21169, supported by COST (European Cooperation in Science and Technology).

Conflict of interest

None

Financial support

None

Ethics statement

None

Received: 15 Sep 2023

Revised: 14 Oct 2023

Accepted: 26 Nov 2023

Published online: 15 March 2024

Rights and permissions

Open Access The author(s) retain copyright. This article is licensed under the Creative Commons Attribution-NonCommercial-ShareAlike 4.0 International License. It may be shared and adapted for non-commercial purposes with appropriate attribution, an indication of changes, and distribution of adaptations under the same license. Third-party material may be subject to separate terms identified in its credit line. View the license at <https://creativecommons.org/licenses/by-nc-sa/4.0/>.

References

Future of Life Institute. Pause giant AI experiments: an open letter. 2023 Mar 22 [cited 2024 Oct 1]. Available from: <https://futureoflife.org/open-letter/pause-giant-ai-experiments/>

Bryant A. AI chatbots: threat or opportunity? *Informatics*. 2023;10(2):49.
<https://doi.org/10.3390/informatics10020049>.

Confalonieri R, Coba L, Wagner B, Besold TR. A historical perspective of explainable artificial intelligence. *Wiley Interdiscip Rev Data Min Knowl Discov*. 2021;11(1):e1391.
<https://doi.org/10.1002/widm.1391>.

Ahmed I, Jeon G, Piccialli F. From artificial intelligence to explainable artificial intelligence in Industry 4.0: a survey on what, how, and where. *IEEE Trans Ind Inform*. 2022;18(8):5031-42.
<https://doi.org/10.1109/TII.2022.3146552>.

Vilone G, Longo L. Classification of explainable artificial intelligence methods through their output formats. *Mach Learn Knowl Extr*. 2021;3(3):615-61.
<https://doi.org/10.3390/make3030032>.

Ferraz-Caetano J. The artificial intelligence explanatory trade-off on the logic of discovery in chemistry. *Philosophies*. 2023;8(2):17.
<https://doi.org/10.3390/philosophies8020017>.

Tahir HA, Alayed W, Hassan WU, Haider A. Proposed explainable interference control technique in 6G networks using large language models (LLMs). *Electronics*. 2024;13(22):4375.
<https://doi.org/10.3390/electronics13224375>.

Ferri M. Why topology for machine learning and knowledge extraction? *Mach Learn Knowl Extr*. 2019;1(1):115-20.
<https://doi.org/10.3390/make1010006>.

Holzinger A. Introduction to machine learning and knowledge extraction (MAKE). *Mach Learn Knowl Extr*. 2019;1(1):1-20.
<https://doi.org/10.3390/make1010001>.

Feng JH, Lai EMK, Li W. Geometry of textual data augmentation: insights from large language models. *Electronics*. 2024;13(19):3781.
<https://doi.org/10.3390/electronics13193781>.

Lozić E, Štular B. Fluent but not factual: a comparative analysis of ChatGPT and other AI chatbots' proficiency and originality in scientific writing for humanities. *Future Internet*. 2023;15(10):336.
<https://doi.org/10.3390/fi15100336>.

Gams M, Kramar S. Evaluating ChatGPT's consciousness and its capability to pass the Turing test: a comprehensive analysis. *J Comput Commun*. 2024;12(3):219-37.
<https://doi.org/10.4236/jcc.2024.123014>.

Chen S, Yu S, Zhao S, Lu C. From imitation to introspection: probing self-consciousness in language models. *arXiv [Preprint]*. 2024. arXiv:2410.18819.
<https://doi.org/10.48550/arXiv.2410.18819>.

Pira L, Ferrie C. On the interpretability of quantum neural networks. *arXiv [Preprint]*. 2023. arXiv:2308.11098.
<https://doi.org/10.48550/arXiv.2308.11098>.

Steinmüller P, Schulz T, Graf F, Herr D. eXplainable AI for quantum machine learning. *arXiv [Preprint]*. 2022. arXiv:2211.01441.
<https://doi.org/10.48550/arXiv.2211.01441>.

Marciano A, Chen D, Fabrocini F, Fields C, Greco E, Gresnigt N, et al. Quantum neural networks and topological quantum field theories. *Neural Netw*. 2022;153:164-78.
<https://doi.org/10.1016/j.neunet.2022.05.028>.

Torda T. Interpretability and explainability AI in advanced neural networks for topological quantum field theory and medical application [dissertation]. Rome: Sapienza University of Rome; 2024.

Field B, Simula T. Introduction to topological quantum computation with non-Abelian anyons. *Quantum Sci Technol*. 2018;3(4):045004.

Wang Z. Topological quantum computing. Providence (RI): American Mathematical Society; 2010. p. 115.

Nayak C, Simon SH, Stern A, Freedman M, Das Sarma S. Non-Abelian anyons and topological quantum computation. *Rev Mod Phys*. 2008;80(3):1083-1159.

<https://doi.org/10.1103/RevModPhys.80.1083>.

Biedenharn LC, Lohe MA. Quantum group symmetry and q-tensor algebras. Singapore: World Scientific; 1995.

Kirchner N, Millar D, Ayeni BM, Smith A, Slingerland JK, Pollmann F. Numerical simulation of non-Abelian anyons. *Phys Rev B*. 2023;107:195129.

Xu S, Sun ZZ, Wang K, Xiang L, Bao Z, Zhu Z, et al. Digital simulation of non-Abelian anyons with 68 programmable superconducting qubits. *Chin Phys Lett*. 2023;40(6):060301.

Iqbal M, Tantivasadakarn N, Verresen R, Campbell SL, Dreiling JM, Figgatt C, et al. Non-Abelian topological order and anyons on a trapped-ion processor. *Nature*. 2024;626:505-11.

<https://doi.org/10.1038/s41586-023-06934-4>.

Rowell E, Stong R, Wang Z. On the classification of modular tensor categories. *Commun Math Phys*. 2009;292:343-89.

<https://doi.org/10.1007/s00220-009-0908-z>.

Planat M, Chester D. Topology and dynamics of transcriptome (dys)regulation. *Int J Mol Sci*. 2024;25(9):4971.

<https://doi.org/10.3390/ijms25094971>.

Knight W. OpenAI's CEO says the age of giant AI models is already over. *Wired*. 2023 [cited 2024 Nov 11]. Available from: <https://www.wired.com/story/openai-ceo-sam-altman-the-age-of-giant-ai-models-is-already-over/>

Villalobos P, Ho A, Sevilla J, Besiroglu T, Heim L, Hobbhahn M. Will we run out of data? Limits of LLM scaling based on human-generated data. *arXiv [Preprint]*. 2024. arXiv:2211.04325.

<https://doi.org/10.48550/arXiv.2211.04325>.

Verlinde E. Fusion rules and modular transformations in 2D conformal field theory. *Nucl Phys B*. 1988;300:360-76.

[https://doi.org/10.1016/0550-3213\(88\)90603-7](https://doi.org/10.1016/0550-3213(88)90603-7).

Gils C, Ardonne E, Trebst S, Huse DA, Ludwig AWW, Troyer M, et al. Anyonic quantum spin chains: spin-1 generalizations and topological stability. *Phys Rev B*. 2013;87:235120.

<https://doi.org/10.1103/PhysRevB.87.235120>.

Kirillov AN, Reshetikhin NY. Representations of the algebra $U_q(\mathfrak{sl}(2))$, q-orthogonal polynomials and invariants of links. In: Kac VG, editor. *Infinite dimensional Lie algebras and groups*. Singapore: World Scientific; 1988. p. 285-99.

Johansen EG, Simula T. Fibonacci anyons versus Majorana fermions: a Monte Carlo approach to the compilation of braid circuits in $SU(2)_k$ anyon models. *PRX Quantum*. 2021;2(1):010334.

<https://doi.org/10.1103/PRXQuantum.2.010334>.

Aboumradi W. Quantum computing with anyons: an F-matrix and braid calculator. *arXiv [Preprint]*. 2022. arXiv:2212.0083.

<https://doi.org/10.48550/arXiv.2212.0083>.

Amaral M, Chester D, Fang F, Irwin K. Exploiting anyonic behavior of quasicrystals for topological quantum computing. *Symmetry*. 2022;14(9):1780.

<https://doi.org/10.3390/sym14091780>.

Bauer B, Levailant C. A new set of generators and a physical interpretation for the $SU(3)$ finite subgroup $D(9,1,1;2,1,1)$. *Quantum Inf Process*. 2013;12:2509-21.

<https://doi.org/10.1007/s11128-013-0608-6>.

Levaillant C. On some projective unitary qutrit gates. arXiv [Preprint]. 2013. arXiv:1401.0506.
<https://doi.org/10.48550/arXiv.1401.0506>.

Planat M, Aschheim R, Amaral MM, Irwin K. Informationally complete characters from quark and lepton mixings. *Symmetry*. 2020;12(6):1000.
<https://doi.org/10.3390/sym12061000>.

Lillicrap TP, Santoro A, Marris L, Akerman CJ, Hinton G. Backpropagation and the brain. *Nat Rev Neurosci*. 2020;21(6):335-46.
<https://doi.org/10.1038/s41583-020-0277-3>.

Hofer C, Kwitt R, Niethammer M, Uhl A. Deep learning with topological signatures. In: *Advances in Neural Information Processing Systems 30 (NeurIPS 2017)*. 2017. p. 1633-43.
<https://doi.org/10.48550/arXiv.1707.04041>.

Zheng Z, Wang Y, Huang Y, Song S, Yang M, Tang B, et al. Attention heads of large language models: a survey. arXiv [Preprint]. 2024. arXiv:2409.03752.
<https://doi.org/10.48550/arXiv.2409.03752>.

Khurana D, Koli A, Khatter K, Singh S. Natural language processing: state of the art, current trends and challenges. *Multimed Tools Appl*. 2023;82(3):3713-44.
<https://doi.org/10.1007/s11042-022-13428-4>.

Otter DW, Medina JR, Kalita JK. A survey of the usages of deep learning for natural language processing. *IEEE Trans Neural Netw Learn Syst*. 2021;32(2):604-24.
<https://doi.org/10.1109/TNNLS.2020.2979670>.

Raynal M. *Concurrent programming: algorithms, principles, and foundations*. Berlin: Springer; 2013.
<https://doi.org/10.1007/978-3-642-32027-9>.

Cui SX, Qiu Y, Wang Z. From three dimensional manifolds to modular tensor categories. *Commun Math Phys*. 2023;397:1191-235.

Boalch P. From Klein to Painlevé via Fourier, Laplace and Jimbo. *Proc Lond Math Soc*. 2005;90(1):167-208.

Planat M, Chester D, Irwin K. Dynamics of Fricke-Painlevé VI surfaces. *Dynamics*. 2024;4(1):1-13.
<https://doi.org/10.3390/dynamics4010001>.

Deng DL, Li X, Das Sarma S. Machine learning topological states. *Phys Rev B*. 2017;96:195145.
<https://doi.org/10.1103/PhysRevB.96.195145>.

Sanborn S, Mathe J, Papillon M, Buracas D, Lillemark HJ, Shewmake C, et al. Beyond Euclid: an illustrated guide to modern machine learning with geometric, topological, and algebraic structures. arXiv [Preprint]. 2024. arXiv:2407.09468.
<https://doi.org/10.48550/arXiv.2407.09468>.