

# Applying QNLP to sentiment analysis in finance

Jonas Stein<sup>✉</sup>  
LMU Munich

Ivo Christ  
Aqarios GmbH

Nicolas Kraus  
Aqarios GmbH

Maximilian Blathasar Mansky  
LMU Munich

Robert Müller<sup>✉</sup>  
LMU Munich

Claudia Linnhoff-Popien<sup>✉</sup>  
LMU Munich

**Abstract**—As an application domain where the slightest qualitative improvements can yield immense value, finance is a promising candidate for early quantum advantage. Focusing on the rapidly advancing field of Quantum Natural Language Processing (QNLP), we explore the practical applicability of the two central approaches DisCoCat and Quantum-Enhanced Long Short-Term Memory (QLSTM) to the problem of sentiment analysis in finance. Utilizing a novel ChatGPT-based data generation approach, we conduct a case study with more than 1000 realistic sentences and find that QLSTMs can be trained substantially faster than DisCoCat while also achieving close to classical results for their available software implementations.

**Index Terms**—Quantum Computing, Sentiment Analysis, Finance, QNLP, DisCoCat, QLSTM

## I. INTRODUCTION

Predicting the price development of assets is one of the core problems in finance. As the price of any asset is determined by supply and demand, the individual valuation of each buyer and seller constitutes essential information for the task of price prediction. With the advent of social media platforms like twitter, such data became freely available at massive scale. This catalyzed the application of methods from natural language processing (NLP) to allow for automated analysis of investor’s sentiments [1]. While the field of NLP has achieved tremendous success in recent history [2] – even showing substantial evidence to pass the Turing test [3], [4] – the compute used to achieve these advancements is enormous [5].

A particularly promising novel approach to NLP and computational finance in terms of hardware requirements is quantum computing, as it has been shown to need less training data in machine learning [6] and allows for the efficient incorporation of grammatical structure in NLP [7], [8]. Motivated by recent positive experimental results of two quantum NLP (QNLP) methods, namely the Quantum-Enhanced Long Short-Term Memory (QLSTM) neural networks [9], [10] and the quantum native DisCoCat (Distributional Compositional Categorical) based QNLP [11]–[13], we evaluate their applicability to sentiment analysis on financial datasets beyond existing proofs of concept.

This work encompasses the following contributions:

- A comprehensive guide for the application of QLSTMs and DisCoCat.
- Introducing a novel approach to generating datasets for DisCoCat that enables automated data generation with sufficient grammatical structure using ChatGPT.

- A case study comparing the applicability of two state-of-the-art QNLP methods in a finance application.

The remainder of this paper is structured as follows. Sec. II discusses theoretic preliminaries on QLSTMs and DisCoCat. Sec. III shows how both methods can be applied to solve the given task of sentiment analysis. Sec. IV contains the evaluation and subsequent comparison of both approaches on synthetic datasets. Sec. V concludes the findings.

## II. BACKGROUND

A central component in state-of-the-art NLP approaches is choosing how to model text mathematically using *distributional semantics*, i.e., mapping each word onto a vector, which generally leads to similar words being described by vectors with a high similarity [14]. Historically, such *word embeddings* were mainly done in a static manner, meaning that words were mapped onto vectors based on their possible meanings according to, e.g., dictionaries [15]. As NLP evolved, more context aware word embeddings were used, which allows solving potential problems raised from *homonyms* [16]. Homonyms are words that have different meanings based on the context they are used in, e.g., the word “rock” can reference a stone or a genre of music. Today’s state-of-the-art approaches in regards to qualitative performance avoid such problems using *dynamic* embeddings, which is heavily driven by the availability of massive amounts of data and immense computing resources [17].

Expanding on the mathematical model for text, we now examine methods to process it computationally. As text is a form of temporal data with variable length, recurrent neural networks (RNNs) have been extensively used to infer information from the text [18]. However, ordinary RNNs have difficulties accounting for a grammatical semantics that extend beyond locally available information in a sentence [19], e.g., as in “The view from this ugly skyscraper was incredibly beautiful”, where the semantically related words “view” and “beautiful” are separated by almost the whole sentence. Long term dependencies are successfully addressed in a special type of RNNs, i.e., LSTMs, by using the propagation of a cell state in addition to the hidden state over time [20]. Expanding on this, the concept of transformers was introduced, which focus on self-attention while shifting away from the recurrent network structure to allow for parallelization in the otherwise necessarily iterative training procedure [21].

In combination with substantial compute and large amounts of data, current high performance NLP approaches have shown the ability to learn grammatical semantics without any explicit knowledge about grammar. A completely different approach from RNNs, LSTMs and Transformers is DisCoCat, which natively integrates grammatical information given in *syntax trees*<sup>1</sup> using *category theory* to unify distributional semantics with the *principle of compositionality*<sup>2</sup>. In essence, DisCoCat exploits a mathematical similarity between grammar, expressed through syntax trees, and quantum computing [11].

Expanding on this brief overview of NLP, we now show how quantum computing techniques can be productively applied to enhance LSTMs and DisCoCat based NLP.

### A. Using QLSTMs for QNLP

(Q)LSTMs can be viewed as an extension of conventional (Q)RNNs with two core differences: (1) the introduction of an additional cell state  $c_t$  for long term memory storage besides the hidden state  $h_t$  and (2), the replacement of the single neural network with a specific composition of neural networks (i.e., a (Q)LSTM cell) as shown in Fig. 1. [9], [20]

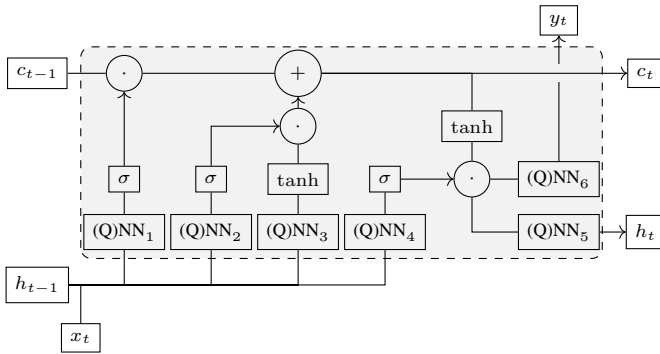


Fig. 1: Structure of a (Q)LSTM in which the internal processing pipeline is indicated by arrows. Inputs are the cell state  $c_{t-1}$  of the previous time step, the corresponding hidden state  $h_{t-1}$  and the current state of the time series data  $x_t$  to be processed.  $c_t$  and  $h_t$  denote the computed cell- and hidden states passed to the next cell. The output of the cell is  $y_t$ .

In this figure,  $\sigma$  denotes a sigmoid activation function, analogue notation is used for  $\tanh$ . The operations  $\odot$  and  $\oplus$  denote element-wise multiplication and addition of vectors.  $x_t$  denotes the data input at point  $t$  and  $y_t$  the output at the same point in time,  $h_t$  denotes the hidden state similar to the hidden state in conventional RNNs and  $c_t$  denotes the cell state of the LSTM. [24]

The LSTM cell is divided into three neural networks acting as gates:

- 1) The *forget gate*: The first (Q)NN determines the amount of information to be forgotten about the state  $c_{t-1}$ . The

<sup>1</sup>In the context of grammar, a syntax tree describes the parts of speech such as nouns, verbs, adjectives, etc., and puts them in structural context [22].

<sup>2</sup>This principle states that the meaning of an expression is fully described by the meaning and composition of its parts [23].

lower the value of  $\sigma((\text{QNN}_1(h_{t-1}, x_t)))$ , the more the resulting entries of  $c_{t-1}$  approach zero.

- 2) The *input gate*: Involving the second and third QNN, the update gate is split into deciding how much information is to be updated ( $(\text{QNN}_2)$ ) and which information is to be updated ( $(\text{QNN}_3)$ ).
- 3) The *output gate*: Determines the information to be output based on all states  $x_t$ ,  $h_{t-1}$  and the processed  $c_{t-1}$ .

In the presented structure, additional postprocessing of the generated state from the output gate is conducted individually to allow for yielding different outputs  $h_t$  and  $y_t$  (in some implementations,  $y_t = h_t$  is used). Choosing suitable dimensionalities for  $h$  and  $c$ , as well as architectures for all (Q)NNs, (Q)LSMTs can be used to process time series completely analog to standard (Q)RNNs [25].

### B. Using DisCoCat for QNLP

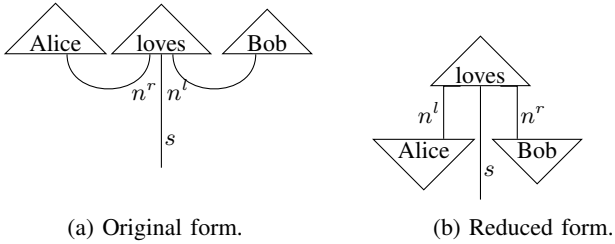
Following the original DisCoCat approach proposed by Coecke et al. [23], [26], there are three steps to generate a mathematical model of a sentence that can be processed using quantum computers:

- 1) Parse the sentence into a pregroup expression.
- 2) Construct the corresponding DisCoCat diagram (incl. possible reductions).
- 3) Translate the DisCoCat diagram into a quantum circuit according to possibly parameterized word embeddings.

A pregroup  $(A, 1, \cdot, -^l, -^r, \leq)$  can be understood as a special non-commutative group, i.e., essentially a group with differentiating between left and right inverses and an order relation. By assigning a type (i.e., a pregroup element) to each word in a sentence and concatenating these with suitable left and right inverses depending on their grammatical function in the sentence, we can express the sentence in form of a pregroup expression. Considering the sentence “Alice loves Bob” and following the established pregroup grammar for the English language [26], all nouns are identified with the pregroup element  $n$  and the verb “loves” is identified as  $n^r \cdot s \cdot n^l$ , meaning that it expects a noun from the left and right and yields a complete sentence  $s$  when concatenated with these:

$$\begin{aligned} \text{“Alice loves Bob”} &\mapsto n \cdot (n^r \cdot s \cdot n^l) \cdot n & (1) \\ &= (n \cdot n^r) \cdot s \cdot (n^l \cdot n) = 1 \cdot s \cdot 1 = s \end{aligned}$$

Expanding on this model of grammar, we now introduce the distributional semantics aspect of DisCoCat to construct the DisCoCat diagram (step 2). This is done by mapping each atomic type  $a \in A$  onto a vector space, i.e.,  $n$  to  $N$  and  $s$  to  $S$ , and their concatenation to tensor product spaces ( $n^r \cdot s \cdot n^l$  to  $N \otimes S \otimes N$ ). Therefore, a pregroup expression is then represented as a concatenation of functions, e.g., “loves” becomes a bilinear map  $N \times N \rightarrow S$  while adjectives (of type  $n \cdot n^l$ ) are represented by a linear map  $N \rightarrow N$ . Using the diagrammatic calculus of compact closed categories, such computations can be represented in so-called DisCoCat diagrams, as shown in Fig. 2a. Following the established choices of maps and



(a) Original form. (b) Reduced form.  
Fig. 2: DisCoCat diagrams for “Alice loves Bob”.

using complex vector spaces, these diagrams can be reshaped according to mathematical properties as, e.g., the isomorphism between row and column vectors through taking the conjugate transpose. A possible reduction of the example in Fig. 2a is depicted in Fig. 2b.

To finally conduct step 3, i.e., mapping the resulting DisCoCat diagram onto a quantum circuit,  $\nabla$  boxes are substituted by quantum states  $|\psi\rangle$  based on the chosen word embedding, while  $\triangle$  boxes correspond to their adjoints  $\langle\psi|$ . Reading the DisCoCat diagrams from top to bottom and defining straight wires  $|$  as identity operations, this leads to the quantum circuit shown in Fig. 3. The different structures shown in equation (1), the diagrammatic representation in Fig. 2 and the quantum circuit in Fig. 3 are equivalent in a categorical sense.

Note that this QNLP approach requires classical postselection, i.e., to extract the meaning of the sentence  $|s\rangle$ , wires 1 and 3 have to be measured in the  $|0\rangle^{\otimes N}$  state. In practice, this can contribute to a potentially immense increase in overall runtime, especially when scaling up the number of qubits per DisCoCat wire, i.e., in this case, the dimensionality of  $N$ .

### III. METHODOLOGY

In the following, we describe our process of data generation, the choice for the classical baseline facilitating meaningful evaluation of the results, as well as the implementation of the QLSTM and the DisCoCat approach.

#### A. Dataset generation

Aiming to use the same datasets for the QLSTM and DisCoCat approaches, the sentences must inherit a grammatical structure that is comprehensible to the syntax parser employed in DisCoCat. To the authors’ awareness, no real world datasets inheriting this property while also being small enough for current quantum circuit simulators are available to date. Motivated by the recent success of large language models like ChatGPT [27], we propose to employ such means to create suitable synthetic datasets. To generate simple sentences containing sentiments on finance topics, we interacted with ChatGPT in the following way:

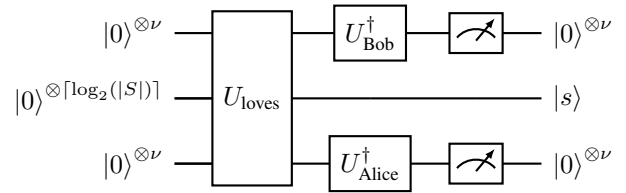


Fig. 3: A quantum circuit for “Alice loves Bob” when using an arbitrary word embedding  $U_{\text{word}}$ , where  $\nu := \lceil \log_2(|N|) \rceil$ .

Generate sentences with a maximum length of five words discussing financial topics or stocks in a positive, neutral or negative way. At the end of each sentence, mention its respective label with negative being 0, neutral being 1 and positive being 2. [Query]

Inflation fears rattle markets (Negative - 0)  
Interest rates stay steady (Neutral - 1)  
[ChatGPT] Apple reports record profits (Positive - 2)

To generate more complex statements, we altered the input to: “Generate detailed sentences discussing financial topics or stocks in a positive, neutral or negative way. [...]”. An example of a positive reply to this was “The rise of online banking has made it easier and more convenient for customers to manage their finances”.

An overview of the generated datasets is displayed in Tab. I.

Complexity	Class			∅ word count	vocabulary size
	-	o	+		
Low	34%	18%	48%	4.9	913
Moderate	37%	17%	46%	18.4	1608

TABLE I: Data class distribution of both generated datasets encompassing roughly 1000 sentences each.

#### B. Classical Baseline LSTM

To allow for a meaningful comparison of results, we employ a LSTM as the classical baseline. Following Sec. II-A, a hyperparameter search the low (moderate) complexity datasets lead to the use of the ReLU activation function, a trainable word embedding layer with size 5 (10), one unidirectional LSTM layer, one fully connected hidden layer of size 8 (16), followed by a dropout layer with rate 0 (0.1).

#### C. QLSTM Implementation

Analogously, QLSTM implementation also follows the general architecture described in Sec. II-A but differs from the classic baseline by exchanging the LSTM with a QLSTM cell. As a result of a basic hyperparameter tuning for the low (moderate) complexity datasets, we choose: A three qubit one hot encoding for  $y_t$ , four qubits and one ansatz layer in each QNN, whereas the ansatz corresponds to that proposed in the

original QLSTM paper<sup>3</sup> for the low (moderate) complexity datasets [9]. The shape of  $h$  and  $c$  inside the (Q)LSTM cell are determined by concatenation  $v_t$  of the previous hidden state  $h_{t-1}$  and current input vector  $x_t$ .

#### D. DisCoCat Implementation

For implementing the DisCoCat approach, we follow the procedure stated in Sec. II-B. In doing so, the first key component is a suitable pregroup parser. As no stable pregroup parsers exists at the time of conducting this research, we use the common, closest suited substitution: A CCG parser [28]. In our implementation we used the BobCatParser from the Lambeq [29] python library for this task. For translating the DisCoCat diagrams into quantum circuits, we follow [12] by choosing these ansätze: Single qubit words are embedded with the standard Euler parameterization (i.e.,  $R_x(\theta_3)R_z(\theta_2)R_x(\theta_1)|0\rangle$ ) and words spanning over  $d > 1$  qubits are embedded with  $d$  many IQP layers [30], which initially apply a Hadamard gate to each qubit and then apply parameterized, controlled  $R_z$  gates in a chain-like pattern iteratively for all neighboring wires.

As the only available software implementation for DisCoCat<sup>4</sup> merely supported CPU-based circuit simulation (opposed to the available GPU-support for the (Q)LSTM implementation) lead to major wallclock time issues, we restricted the DisCoCat approach to binary classification by dropping the neutral data class data. While this generally makes the results less comparable, its impact is arguably negligible in the encountered case where QLSTM outperforms DisCoCat. The fundamental benefit of this is, that  $|s\rangle$  can be represented with a single qubit, as it only needs to carry binary information hence requiring the least possible computational effort.

#### IV. EVALUATION

To evaluate the applicability of QLSTMs and DisCoCat to sentiment analysis in finance, we employ both approaches as specified in Sec. III on the generated datasets using an 80/10/10 train/validation/test split and employing the binary cross entropy loss. As displayed in Fig. 4, the classical and the QLSTM are capable of learning the data in merely 20 epochs with validation accuracies over 80%, hence showing very promising performance. For the moderate complexity dataset, much longer training times appear to be necessary, especially for the QLSTM, where it takes about 100 epochs before training progress starts to manifest, indicating possible barren plateaus. A more extensive hyperparameter search might circumvent this issue, but as the simulation already took 10 hours with a small sized cloud compute approach, classical circuit simulation appears to become a substantial bottleneck.

<sup>3</sup>The ansatz consists of a data input layer, a variational layer and a measurement layer. The data input layer starts with Hadamard gates on every qubit  $q_i$  (where  $i \in \{1, \dots, n\}$ ), then proceeds with  $R_y$  gates rotating for  $\arctan(x_i)$  and finishes analogously with  $R_z$  rotating for  $\arctan(x_i^2)$ , where  $x_i$  denotes the scalar data inputs. The variational layer starts with cyclic CNOTs where every qubit  $q_i$  iteratively acts as the control for the  $q_{i+k \bmod n}$ -th qubit, where  $k \in \{1, 2\}$ .

<sup>4</sup>For details, see [https://github.com/ichrist97/qnlp\\_finance](https://github.com/ichrist97/qnlp_finance).

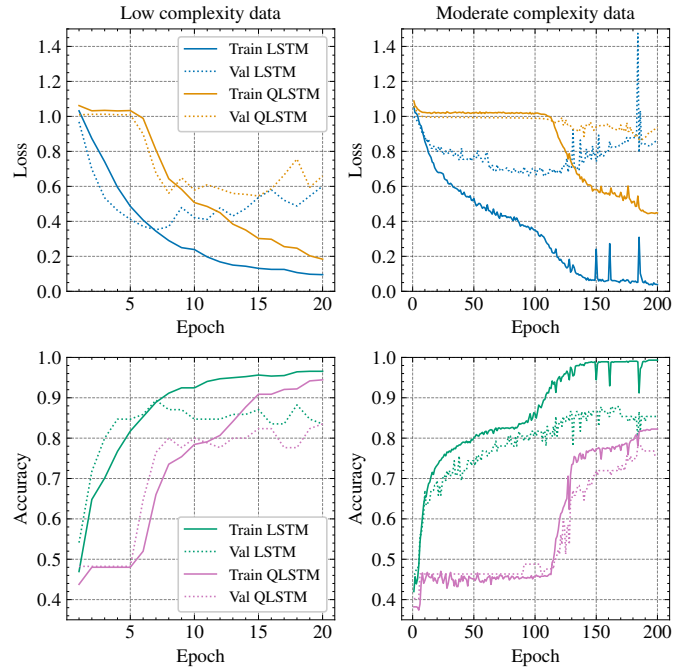


Fig. 4: Training curve for LSTM and QLSTM on both datasets.

In course of the conducted hyperparameter optimization, this effect appears to be strongly correlated with the number of ansatz layers employed, which is sensible, as an increased number of layers leads to a bigger number of parameters – which is a natural cause of barren plateaus [31]. This correspondence indicates a time/quality trade-off which should be investigated more closely in future work.

For the DisCoCat approach, only the low complexity data was used, as the hardware requirements were substantially higher with 100 epochs taking roughly 82 hours of wallclock time with the same setup used for the QLSTM evaluation. Examining the results of DisCoCat displayed in Fig. 5, we can observe that while a slight trend of improvement in the loss and accuracy is visible over the course of all 100 epochs, it is barely better than random guessing, which would achieve a 50% accuracy. This indicates that QLSTMs are a lot faster to train than DisCoCat based approaches for realistic data sizes and complexities when using the employed CPU based circuit simulator. However, we argue against generalizing these results towards “QLSTMs outperform DisCoCat“, as a clear learning progress of DisCoCat is observable. Extrapolating from similar studies [12], we expect reasonably accurate results for DisCoCat given a more efficient software implementation.

#### V. CONCLUSION

In sight of current QC hard- and software limitations, QLSTMs seem to be better suited when solving QNLP tasks like sentiment analysis on the employed datasets, which appear to be a lot more realistic than those used in existing proofs of concept for both approaches. Notably, the employed ChatGPT-based data generation approach enabled producing realistic

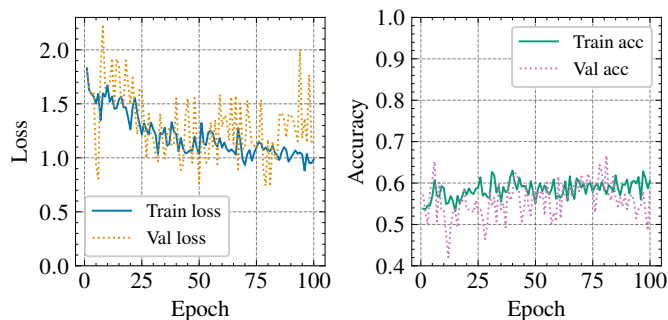


Fig. 5: Training curve for DisCoCat on low complexity data.

data, while preserving necessary grammatical correctness, which facilitates more practical testing. A future possibility to overcome the experienced circuit simulation bottleneck might be the support of GPU based circuit simulation for the employed implementation using pennylane and Lambeq.

#### ACKNOWLEDGMENT

This paper was partially funded by the German Federal Ministry for Economic Affairs and Climate Action through the funding program "Quantum Computing – Applications for the industry" based on the allowance "Development of digital technologies" (contract number: 01MQ22008A). Furthermore, this paper was partially funded by the German Federal Ministry of Education and Research through the funding program "Network of interdisciplinary education and training concepts in quantum technologies" based on the allowance "Quantum technologies – from fundamentals to market" (contract number: 13N16000).

#### REFERENCES

- [1] Shen Ao. Sentiment analysis based on financial tweets and market information. In *2018 International Conference on Audio, Language and Image Processing (ICALIP)*, pages 321–326, 2018.
- [2] Narendra Patwardhan, Stefano Marrone, and Carlo Sansone. Transformers in the real world: A survey on nlp applications. *Information*, 14(4), 2023.
- [3] A. M. TURING. I.—COMPUTING MACHINERY AND INTELLIGENCE. *Mind*, LIX(236):433–460, 10 1950.
- [4] Siddharth Nath, Abdullah Marie, Simon Ellershaw, Edward Korot, and Pearse A Keane. New meaning for NLP: the trials and tribulations of natural language processing with GPT-3 in ophthalmology. *Br. J. Ophthalmol.*, 106(7):889–892, jul 2022.
- [5] Deepak Narayanan, Mohammad Shoeybi, Jared Casper, Patrick LeGresley, Mostofa Patwary, Vijay Korthikanti, Dmitri Vainbrand, Prethvi Kashinkunti, Julie Bernauer, Bryan Catanzaro, Amar Phanishayee, and Matei Zaharia. Efficient large-scale language model training on gpu clusters using megatron-lm. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC '21*, New York, NY, USA, 2021. Association for Computing Machinery.
- [6] Matthias C Caro, Hsin-Yuan Huang, Marco Cerezo, Kunal Sharma, Andrew Sornborger, Lukasz Cincio, and Patrick J Coles. Generalization in quantum machine learning from few training data. *Nature communications*, 13(1):4919, 2022.
- [7] Raffaele Guarasci, Giuseppe De Pietro, and Massimo Esposito. Quantum natural language processing: Challenges and opportunities. *Applied Sciences*, 12(11), 2022.
- [8] Nikitas Stamatopoulos, Guglielmo Mazzola, Stefan Woerner, and William J. Zeng. Towards Quantum Advantage in Financial Market Risk using Quantum Gradient Algorithms. *Quantum*, 6:770, July 2022.

- [9] Samuel Yen-Chi Chen, Shinjae Yoo, and Yao-Lung L. Fang. Quantum long short-term memory. In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8622–8626, 2022.
- [10] Riccardo Di Sipio, Jia-Hong Huang, Samuel Yen-Chi Chen, Stefano Mangini, and Marcel Worring. The Dawn of Quantum Natural Language Processing. In *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022.
- [11] Konstantinos Meichanetzidis, Alexis Toumi, Giovanni de Felice, and Bob Coecke. Grammar-Aware Question-Answering on Quantum Computers. 2020.
- [12] Robin Lorenz, Anna Pearson, Konstantinos Meichanetzidis, Dimitri Kartsaklis, and Bob Coecke. Qnlp in practice: Running compositional models of meaning on a quantum computer. *J. Artif. Int. Res.*, 76, apr 2023.
- [13] Victor Martinez and Guillaume Leroy-Meline. A multiclass Q-NLP sentiment analysis experiment using DisCoCat, 2022.
- [14] Faiza Khan Khattak, Serena Jeblee, Chloé Pou-Prom, Mohamed Abdalla, Christopher Meaney, and Frank Rudzicz. A survey of word embeddings for clinical text. *Journal of Biomedical Informatics*, 100:100057, 2019. Articles initially published in *Journal of Biomedical Informatics*: X 1-4, 2019.
- [15] Daniel Jurafsky and James H. Martin. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. Prentice Hall PTR, USA, 1st edition, 2000.
- [16] Kazuhiro Kimura, Takashi Uzuoka, and Shin-ya Amano. Association-based natural language processing with neural networks. In *30th Annual Meeting of the Association for Computational Linguistics*, pages 224–231, 1992.
- [17] OpenAI. Gpt-4 technical report, 2023.
- [18] Kanchan M Tarwani and Swathi Edem. Survey on recurrent neural network in natural language processing. *Int. J. Eng. Trends Technol*, 48(6):301–304, 2017.
- [19] Gongbo Tang, Mathias Müller, Annette Rios, and Rico Sennrich. Why self-attention? a targeted evaluation of neural machine translation architectures. In *2018 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2018.
- [20] Sepp Hochreiter and Jürgen Schmidhuber. Long Short-Term Memory. *Neural Computation*, 9(8):1735–1780, 11 1997.
- [21] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention Is All You Need. In *Advances in Neural Information Processing Systems 30 (NIPS 2017)*. Curran Associates, Inc., 2017.
- [22] J. Lambek. *From Word to Sentence: A Computational Algebraic Approach to Grammar*. Open access publications. Polimetrica, 2008.
- [23] Bob Coecke, Mehrnoosh Sadrzadeh, and Stephen Clark. Mathematical foundations for distributed compositional model of meaning. *lambek festschrift. Linguistic Analysis*, 36:345–384, 2010.
- [24] Ian Goodfellow, Joshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.
- [25] Harshal Dhake, Yashwant Kashyap, and Panagiotis Kosmopoulos. Algorithms for hyperparameter tuning of lstms for time series forecasting. *Remote Sensing*, 15(8):2076, 2023.
- [26] William Zeng and Bob Coecke. Quantum algorithms for compositional natural language processing. In Dimitrios Kartsaklis, Martha Lewis, and Laura Rimell, editors, *Proceedings of the 2016 Workshop on Semantic Spaces at the Intersection of NLP, Physics and Cognitive Science*, Glasgow, Scotland, 11th June 2016, volume 221 of *Electronic Proceedings in Theoretical Computer Science*, pages 67–75. Open Publishing Association, 2016.
- [27] OpenAI. Introducing ChatGPT, 2022.
- [28] Richie Yeung and Dimitri Kartsaklis. A ccg-based version of the disoccat framework. In *Proceedings of the 2021 Workshop on Semantic Spaces at the Intersection of NLP, Physics, and Cognitive Science (SemSpace)*, pages 20–31, 2021.
- [29] Dimitri Kartsaklis, Ian Fan, Richie Yeung, Anna Pearson, Robin Lorenz, Alexis Toumi, Giovanni de Felice, Konstantinos Meichanetzidis, Stephen Clark, and Bob Coecke. lambeq: An Efficient High-Level Python Library for Quantum NLP. 2021.
- [30] Vojtěch Havlíček, Antonio D Córcoles, Kristan Temme, Aram W Harrow, Abhinav Kandala, Jerry M Chow, and Jay M Gambetta. Supervised learning with quantum-enhanced feature spaces. *Nature*, 567(7747):209–212, 2019.

- [31] Andrea Skolik, Jarrod R McClean, Masoud Mohseni, Patrick van der Smagt, and Martin Leib. Layerwise learning for quantum neural networks. *Quantum Machine Intelligence*, 3:1–11, 2021.