Quantum Natural Language Processing 2019

Embracing the NISQ Era

St Aldate's Church, Oxford, UK
December 5th – 6th, 2019

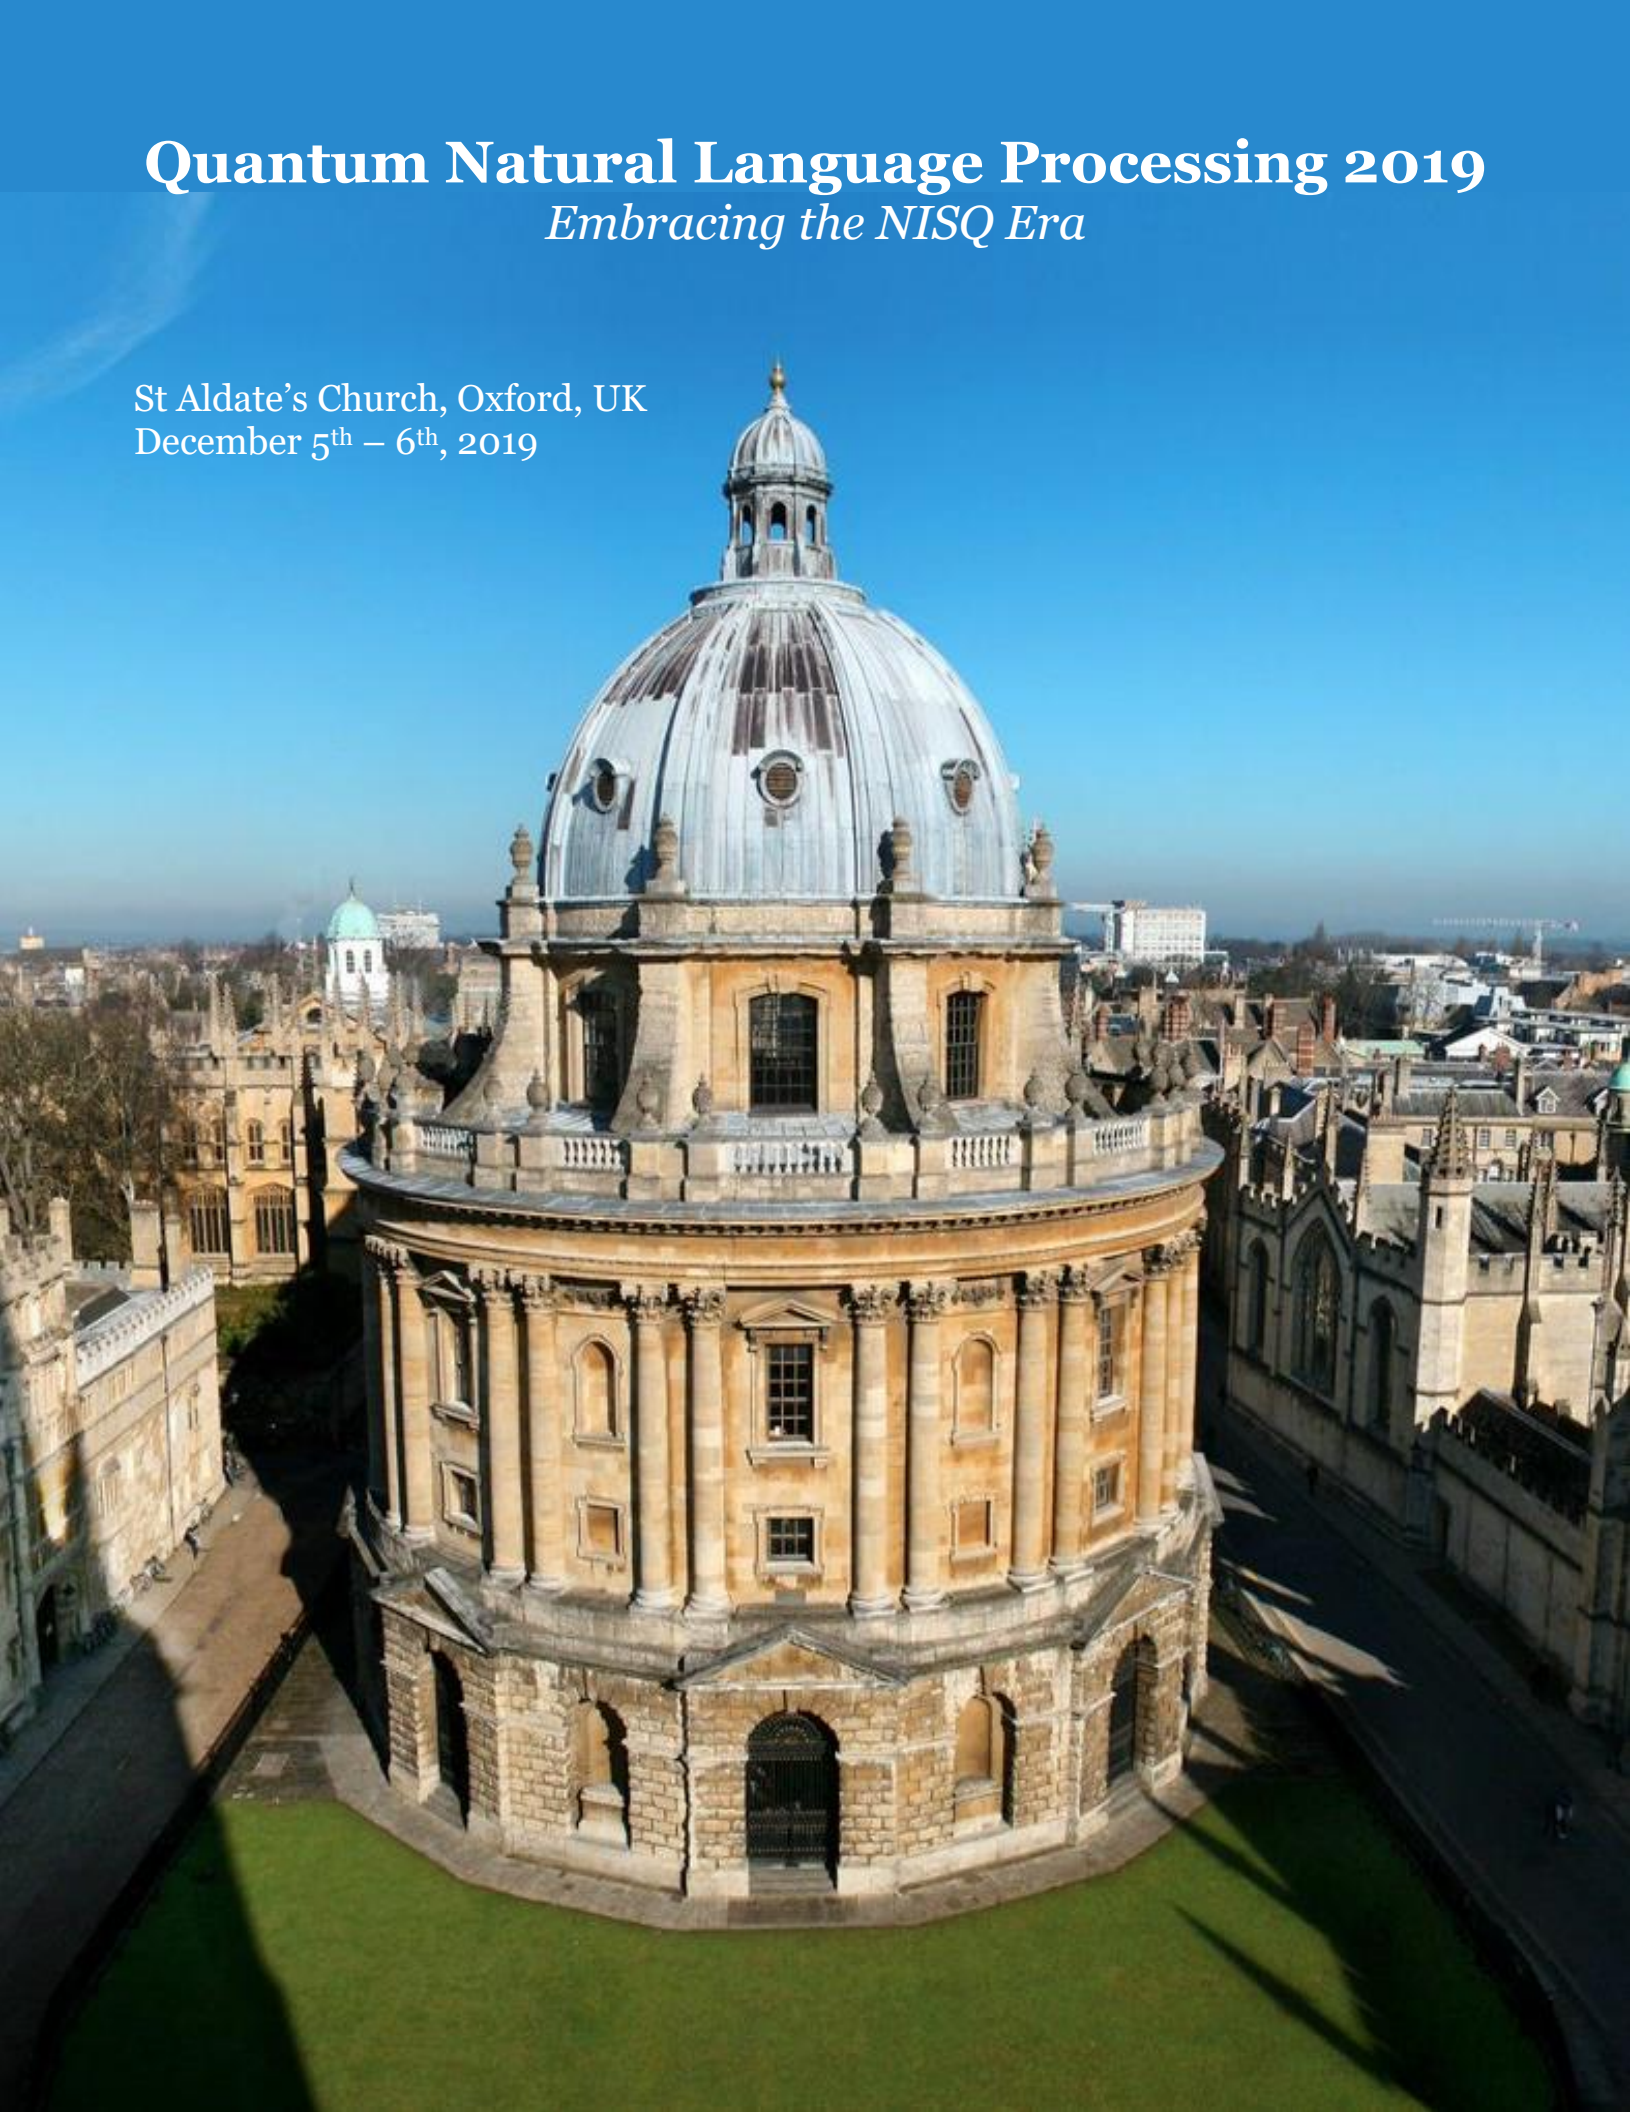

Table of Contents

Sponsors and Organizing Institutions

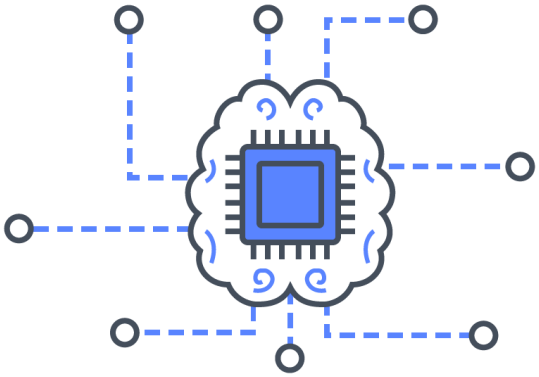

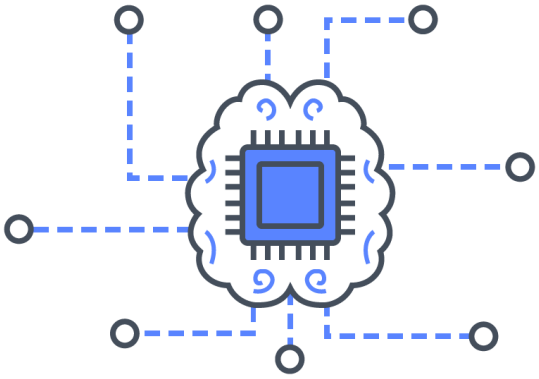

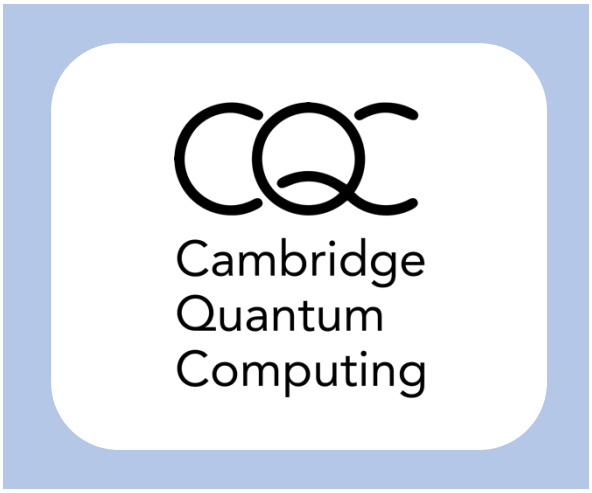

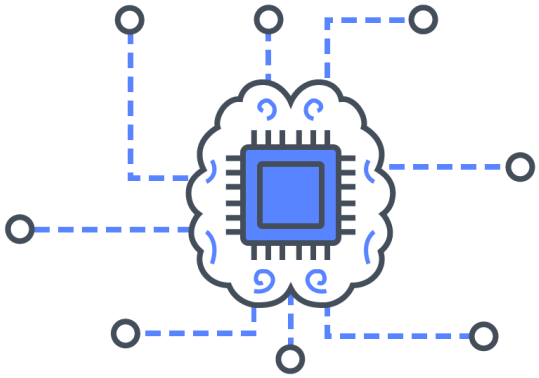

Welcome

The purpose of the conference is to establish a dialogue centered around compositional natural language processing (NLP) and applications on quantum hardware. Applied category theory has provided compositional NLP with mathematical foundations that share similarity to those of quantum theory, suggesting that quantum computers provide a natural setting for compositional NLP tasks [1]. Furthermore, ideas from condensed matter physics have made their way into NLP [2,3]. In recent years, Quantum Artificial Intelligence (QAI) has enjoyed fruitful research activity. As noisy intermediate-scale quantum (NISQ) computers have become readily accessible, the time is right to explore the possible advantages of performing NLP on such NISQ devices. This endeavor is to be understood as an emerging subfield of QAI, namely Quantum Natural Language Processing (QNLP). Our ambition is that participants from both the academic and the private sector working on the tangent fields of quantum computing, artificial intelligence and machine learning, as well as linguistics, come together and coordinate to build an active community on the novel and timely field of QNLP.

The Conference Team

Programme Committee

Bob Coecke
Chair

Ilyas Khan
Co-Chair

Konstantinos
Meichanetzidis

Organisation Committee

Destiny Chen

Catie Isham

Conference Location

The Conference will take place in St Aldate's Church, opposite Christ Church and next door to Pembroke College.

St Aldate's Address:

40 Pembroke Street
Oxford OX1 1BP
United Kingdom

Coffee, Tea, and Lunch

Coffee and tea will be available throughout the conference and there will be several breaks to enjoy refreshments. Further, within St Aldate's, there is a quiet area upstairs for participants to rest throughout the day.

Lunch will be provided for all conference participants. Alternatively, there are many local pubs and restaurants within a few minutes walk from St Aldate's Church. See page 6 for recommended restaurant and pubs.

Wi-Fi

Name: StAldates Password: John6.35

Area Map

Local Dining

Restaurants:

- LEON (Mediterranean)
- The Handle Bar Café and Kitchen
- Fernando's Café (Brazilian)
- Bill's Oxford Restaurant
- Turtle Bay Oxford (Caribbean)
- Bbuona (Italian)
- Côte Brasseria (French)
- Sticks'n'Sushi Oxford

Pubs:

- The White Horse
- The Crown
- Eagle and Child (for Tolkien and C.S. Lewis fans)
- The Mitre
- The Chequers
- The St. Aldates Tavern
- Turf Tavern
- The Bear Inn
- The King's Arm

Reception and Conference Dinner

For those joining, the reception and conference dinner on Friday, December 6th will be held at the Old Divinity School, Bodleian Library.

The reception will begin at 6:30 pm followed by dinner at 7:00 pm.

The address for the Old Divinity School, Bodleian Library is as follows:
Bodleian Library
Broad Street
Oxford OX1 3BG
United Kingdom

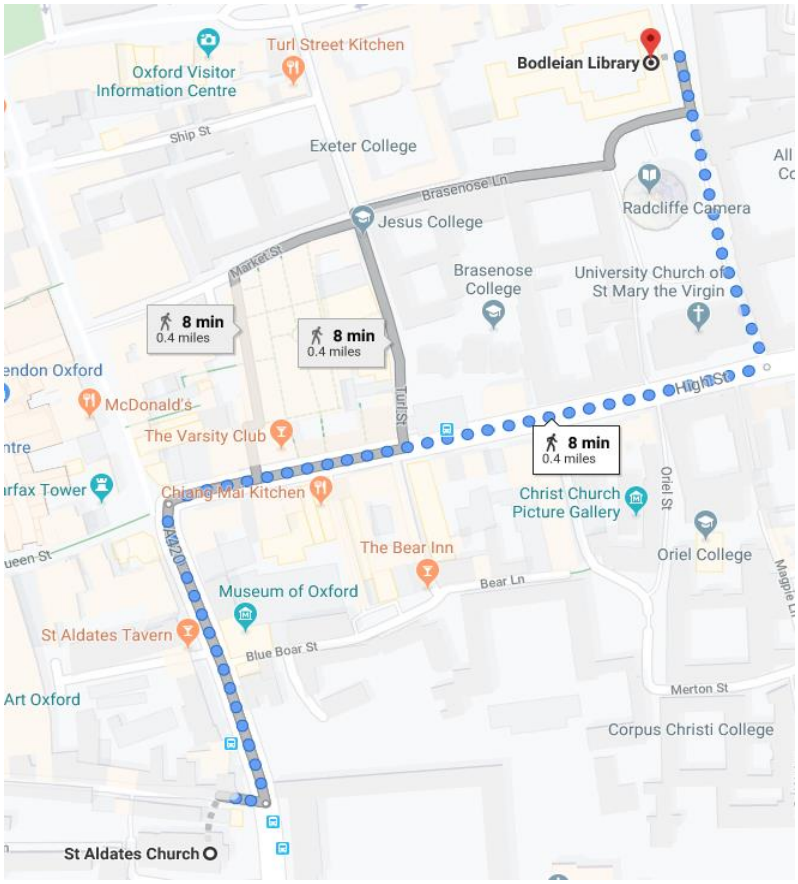

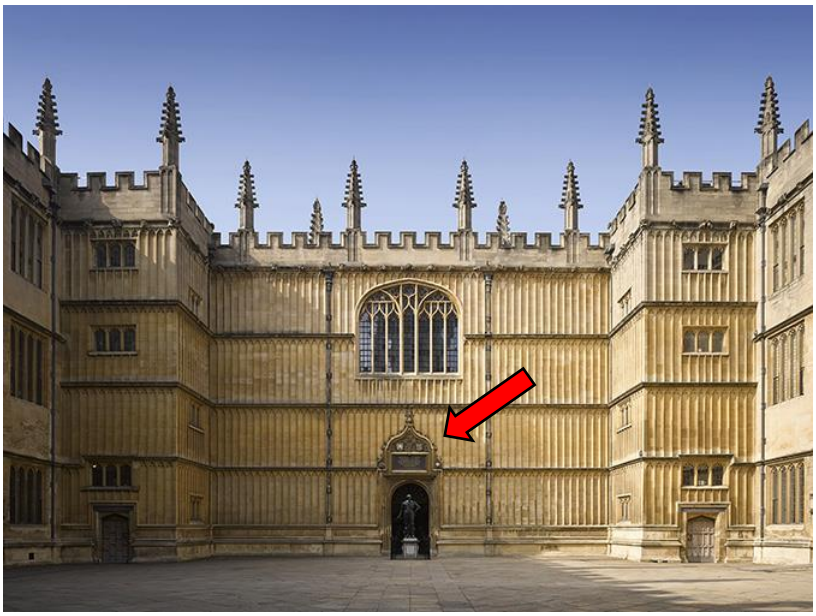

Walking Directions:

Please walk down Catte Street. When you have arrived in the courtyard of the Bodleian Library, please go around the Christmas Tree, and enter the dinner venue using the glass double doors behind the statue.

Programme

Quantum Natural Language Processing 2019
Embracing the NISQ Era

Day 1
Thursday, December 5th, 2019

Time	Talk	Speaker
09:30 – 09:50 am	Registration	
09:50 – 10:00 am	Welcoming Remarks	Bob Coecke – *CQC, University of Oxford* Niels Nielsen – *CQC*
10:00 – 10:40 am	Calculating Sentence Similarity Using a Hybrid Classical-Quantum Workflow	Lee O'Riordan – *Irish Centre for High-End Computing*
10:40 – 11:20 am	Rigorous Statements on Near-Term Quantum Computing	Jens Eisert – *University of Cambridge*
11:20 – 11:50 am	Coffee & Tea Break	
11:50 – 12:30 pm	A Quantum Search Decoder for Natural Language Processing	Johannes Bausch – *University of Cambridge*
12:30 – 13:10 pm	Universality: From Spin Models to Automata	Gemma De las Cuevas – *University of Innsbruck*
13:10 – 14:40 pm	Lunch Break	
14:40 – 15:20 pm	Why NLP is Quantum-Native, and Beyond!	Bob Coecke – *CQC, University of Oxford*
15:20 – 16:00 pm	Compositional Hyponymy with Positive Operators	Martha Lewis – *University of Amsterdam*
16:00 – 16:30 pm	Coffee & Tea Break	
16:30 – 17:10 pm	Compositional NLP and Quantum Mechanics	Dimitri Kartsaklis – *Apple Inc.*
17:10 – 17:50 pm	Towards Natural Language Processing on Quantum Hardware	Giovanni De Felice –CQC, *University of Oxford* Konstantinos Meichanetzidis – *CQC, University of Oxford* Alexis Toumi – CQC, *University of Oxford*

Quantum Natural Language Processing 2019
Embracing the NISQ Era

Day 2
Friday, December 6th, 2019

Time	Talk	Speaker
09:30 – 10:10 am	Which Language Operations to Implement First with Quantum Computers?	Dominic Widdows – *Grab*
10:10 – 10:50 am	Qiskit OpenPulse: Building Quantum Applications from the Ground Up	Nate Earnest-Noble – *IBM*
10:50 – 11:10 am	Coffee & Tea Break	
11:10 – 11:50 am	Working with PyZX	Aleks Kissinger – *University of Oxford*
11:50 – 12:30 pm	A Divide-and-Conquer Hybrid Method for Smaller Quantum Computers	Vedran Dunjko – *Leiden Institute of Advanced Computer Science*
12:30 – 14:00 pm	Lunch Break	
14:00 – 14:40 pm	Hybrid Quantum Algorithms for Machine Learning	Mattia Fiorentini – *CQC*
14:40 – 15:00 pm	Quantum Enhanced Bayesian Inference for NLP Tasks	Linda Anticoli – *CQC*
15:00 – 15:40 pm	A toy model for QNLP takes baby steps into the NISQ era	Stefano Gogioso – *University of Oxford*
15:40 – 16:10 pm	Coffee & Tea Break	
16:10 – 16:55 pm	An Overview of the Key Players and Developments in Quantum Computing	Ilyas Khan – *CQC* Denise Ruffner – *CQC*
16:55 – 17:25 pm	Application Oriented Benchmarks for NISQ Devices	Ross Duncan – *CQC, University of Strathclyde*
18:30 – 19:00 pm	Reception (By Invitation Only)	
19:00 – 21:00 pm	Dinner (By Invitation Only)	

Talk Titles and Abstracts

Calculating Sentence Similarity Using a Hybrid Classical-Quantum Workflow

Speaker: Lee O'Riordan – *ICHEC*

Abstract: Natural language processing (NLP) is often used to perform tasks like sentiment analysis, relationship extraction and word sense disambiguation. The "distributional compositional semantics" (DisCo) formalism incorporates the grammatical structure of sentences of a language into the analysis algorithms, giving grammatically informed algorithms that compute the meaning of sentences. This algorithm has been noted to offer significant improvements to the quality of results, as compared to traditional "bag-of-words" methods. However, the main challenge in its implementation is the need for large classical computational resources. Here, we take inspiration from the DisCo-formalism and implement a hybrid classical-quantum sentence similarity calculation workflow using the Intel Quantum Simulator (Intel-QS). By appropriately analysing and preparing our corpus data, we demonstrate a workflow for calculating the similarity between sentences using quantum states.

Rigorous Statements on Near-Term Quantum Computing

Speaker: Jens Eisert – *Free University of Berlin*

Abstract: Recent years have witnessed an enormous interest in assessing the potential of noisy intermediate scale quantum (NISQ) devices. Significant progress has indeed been made in the development and the understanding of variational hybrid methods that have both a classical and a quantum component, specifically in variational quantum eigensolvers and quantum approximate optimization algorithms. In this talk, we will ask questions of what statements one can expect to be shown with rigorous methods in this emergent field of near-term quantum computing, and what "proof pockets" are within reach. The main part of the talk will be concerned with a stochastic gradient descent method for hybrid quantum-classical optimization that features a recovery guarantee: It overcomes the prejudice that in order to estimate gradients, one should measure expectation values

[1]. Instead, single-shot measurements are significantly more efficient, while maintaining convergence. We will also briefly discuss questions of expressivities, exploring the expressive power of circuits, with surprising exponential separations featuring here [2]. If time allows, we will discuss in an outlook how one could close the conceptual and mathematical loopholes when showing a quantum advantage ("supremacy") [3].

[1] Stochastic gradient descent for hybrid quantum-classical optimization, R. Sweke, F. Wilde, J. Meyer, M. Schuld, P. K. Fährmann, B. Meynard-Piganeau, J. Eisert, arXiv:1910.01155.
[2] Expressive power of tensor-network factorizations for probabilistic modeling, with applications from hidden Markov models to quantum machine learning, I. Glasser, R. Sweke, N. Pancotti, J. Eisert, J. I. Cirac, arXiv:1907.03741.
[3] Closing gaps of a quantum advantage with short-time Hamiltonian dynamics, J. Haferkamp, D. Hangleiter, A. Bouland, B. Fefferman, J. Eisert, J. Bermejo-Vega, arXiv:1908.08069.

A Quantum Search Decoder for Natural Language Processing

Speaker: Johannes Bausch – *University of Cambridge*

Abstract: Probabilistic language models, e.g. those based on an LSTM, often face the problem of finding a high probability prediction from a sequence of random variables over a set of words. This is commonly addressed using a form of greedy decoding such as beam search, where a limited number of highest-likelihood paths (the beam width) of the decoder are kept, and at the end the maximum-likelihood path is chosen. The resulting algorithm has linear runtime in the beam width. However, the input is not necessarily distributed such that a high-likelihood input symbol at any given time step also leads to the global optimum. Limiting the beam width can thus result in a failure to recognise long-range dependencies.

In practice, only an exponentially large beam width can guarantee that the global optimum is found: for an input of length n and average parser branching ratio R, the baseline classical algorithm needs to query the input on average Rn times.

In this work, we construct a quantum algorithm to find the globally optimal parse with high constant success probability. Given the input to the decoder is distributed like a power-law with exponent k>0, our algorithm yields a runtime R^f(R,k), where f≤1/2, and f→0 exponentially quickly for growing k. This implies that our algorithm always yields a super-Grover type speedup, i.e. it is more than quadratically faster than its classical counterpart. We further modify our procedure to recover a quantum beam search variant, which enables an even stronger empirical speedup, while sacrificing accuracy. Finally, we apply this quantum beam search decoder to Mozilla's implementation of Baidu's DeepSpeech neural net, which we show to exhibit such a power law word rank frequency, underpinning the applicability of our model.

Based on: https://arxiv.org/abs/1909.05023

Universality: From Spin Models to Automata

Speaker: Gemma de Las Cuevas – *University of Innsbruck*

Abstract: Why is it so easy to generate complexity? Because essentially every non-trivial system is universal, that is, capable of exploring all complexity in its domain. I will discuss this concept of universality in two domains: for spin models and for automata (or, equivalently, formal languages). I will explain the first step toward linking them rigorously, by which we describe spin hamiltonians as automata. The latter leads to a new complexity measure of hamiltonians, with a different threshold between “easy” and “hard” than the computational complexity of the ground state energy problem.

Why NLP is Quantum-Native, and Beyond!

Speaker: Bob Coecke – *CQC, University of Oxford*

Compositional Hyponymy with Positive Operators

Speaker: Martha Lewis – *University of Amsterdam*

Abstract: In natural language processing words are commonly represented as vectors. However, vector representations do not intrinsically incorporate the hierarchical relationships that obtain between many words. We model words as positive operators. These have an ordering which we interpret as modelling hierarchical information. We describe a simple way of building positive operators for words, and give methods for composing these words representations to form phrases and sentences. We test the methods on simple sentence level entailment datasets.

Compositional NLP and Quantum Mechanics

Speaker: Dimitri Kartsaklis – *Apple Inc.*

Abstract: In this talk I will review some of the developments on compositional models of natural language inspired by quantum mechanics, coming from personal and other people's research in the Quantum Group of the University of Oxford. We'll examine how the two areas connect in an abstract level, and we will see how insights from quantum mechanics can actually find real-world applications in NLP, from a modelling perspective. We will attempt to evaluate these models form a point of view that stands right between theory and (current) practice.

Application Oriented Benchmarks for NISQ Devices

Speaker: Ross Duncan – *CQC, University of Strathclyde*

Abstract: Noisy Intermediate-Scale Quantum (NISQ) computers are rather poor approximations to the quantum computers we all learned about from Nielsen and Chuang. But how bad are they? The answer depends on what you want to do with them. In this talk I'll present the results of volumetric benchmark experiments on recent superconducting devices.

Towards Natural Language Processing on Quantum Hardware

Speakers: Giovanni De Felice – *University of Oxford;*
Konstantinos Meichanetzidis – *CQC, University of Oxford*;
Alexis Toumi – *University of Oxford*

Abstract: We demonstrate proof-of-concept applications of compositional NLP using quantum circuits. The small-scale of our examples constitutes them well-suited for NISQ devices. Further, we discuss extendability of our methods to various NLP tasks. We conclude with a live demo of our in-development Python library.

Which Language Operations to Implement First with Quantum Computers?

Speaker: Dominic Widdows – *Grab*

Abstract: Vector representations have been used in natural language processing for half a century and have flourished particularly strongly in the past decade. The list of linguistic phenomena and operations modelled using vectors by now includes similarity and synonymy, negation and disjunction, conditionals and inference, hyponymy and classification, semantic composition, analogy, ambiguity resolution, sentiment analysis, and of course orthography and spelling correction.

The corresponding mathematical toolbox includes superposition, subspaces, projections, tensor products, density matrices, self-adjoint and positive operators, and spectral decompositions, all of which are also involved in quantum physics.
This talk will attempt to give a guided tour of this language zoo, in the hope that audience members with particular expertise in quantum physics and computing will be able to comment on which parts of this bestiary might be most promising candidates for early implementation on quantum hardware.

Examples will draw especially form short informal language, including fragments from smartphone communications such as text messages and mobile search queries, where the need for effective language processing is especially pressing in some languages in emerging markets.

Qiskit OpenPulse: Building Quantum Applications from the Ground Up

Speaker: Nate Earnest-Noble – *IBM*

Abstract: NISQ devices show promise to make impact on real world problems in the near-term. However, these devices suffer from short coherent times and imperfects gates, which ultimately limits the extent to which these devices can perform. Due to this issue, error mitigation – a general term describing how to reduce the effects of imperfect qubits and/or gates – has become an increasingly important area of research. In this talk I will present the newly developed Qiskit OpenPulse, which allows users to program quantum computers at the pulse level, unlocking the users' ability to program a quantum computer without the need to make use of predefined gates. I will specifically cover how different gates are realized with superconducting qubits and how this allows for the implementation of Richardson extrapolation, a general error mitigation technique that is realized through extending the duration of a quantum gate.

Working with PyZX

Speaker: Aleks Kissinger – *University of Oxford*

Abstract: ZX-diagrams is a useful tool for representing, transforming, simplifying, decomposing, and generally reasoning about quantum circuits. However, even for small, NISQ-era circuits (e.g. tens or hundreds of gates) it is not very practical to work with diagrams by hand. In this talk, I will give a brief overview of some of the techniques developed by myself and others over the past few years for transforming, simplifying, and checking circuits using the ZX-calculus and the PyZX tool.

A Divide-and-Conquer Hybrid Method for Smaller Quantum Computers

Speaker: Vedran Dunjko – *LIACS*

Abstract: Theory shows that arbitrary-sized quantum computers may offer computational advantages for many problems. However, quantum computers on a reasonable horizon will be restricted in many ways, including size.

Can quantum computers of smaller size (limited to M qubits) genuinely speed up interesting algorithms, even when the problem size (n) is much larger than the computer itself (n>>M)?

We describe a positive result: a hybrid divide and conquer strategy which allows us to make better use of smaller quantum computers. Our approach works best for a class of algorithm often employed in artificial intelligence applications. In this talk we will discuss result, and its implications on the booming fields of quantum machine learning and quantum AI.

Hybrid Quantum Algorithms for Machine Learning

Speaker: Mattia Fiorentini – *CQC*

Abstract: Hybrid quantum-classical systems make it possible to utilize existing quantum computers to their fullest extent. Within this framework, parameterized quantum circuits can be regarded as machine learning models with remarkable expressive power. With an increasing number of experimental demonstrations carried out on actual quantum hardware and with software being actively developed, this rapidly growing field is poised to have a broad spectrum of real-world applications. In this talk, I will be presenting presents the components of these models and discusses their application to a variety of data-driven tasks, such as supervised learning and generative modelling.

Quantum Enhanced Bayesian Inference for NLP Tasks

Speaker: Linda Anticoli – *CQC, University of Udine*

Abstract: Bayesian inference is largely used in unsupervised methods for natural language processing (NLP) such as, e.g., inference for probabilistic context free grammars (PCFG), unsupervised part-of-speech (POS) tagging, and structured prediction of sequences and trees. We want to investigate whether Bayesian methods for such NLP tasks may benefit from a quantum speedup or not, e.g. using quantum Gibbs sampling or quantum enhanced Monte Carlo methods.

A toy model for QNLP takes baby steps into the NISQ era

Speaker: Stefano Gogioso – *University of Oxford*

Abstract: We present preliminary results for an implementation of a toy model of compositional distributional semantics on both real and simulated NISQ devices. Current limitations and projected capabilities will be discussed. Diagrams and other pretty pictures will be shown.

An Overview of the Key Players and Developments in Quantum Computing

Speakers: Ilyas Khan – *CQC*; Denise Ruffner – *CQC*

Application Oriented Benchmarks for NISQ Devices

Speaker: Ross Duncan – *CQC*

Abstract: Noisy Intermediate-Scale Quantum (NISQ) computers are rather poor approximations to the quantum computers we all learned about from Nielsen and Chuang. But how bad are they? The answer depends on what you want to do with them. In this talk I'll present the results of volumetric benchmark experiments on recent superconducting devices.

Notes

Notes