AI

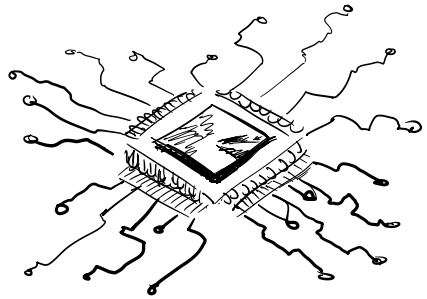

QNLP

NLP

QC

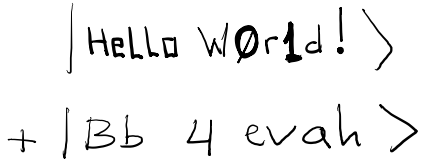

Konstantinos Meichanetzidis

Quantinuum
Cambridge Quantum

NLP Meetup Jan 2021

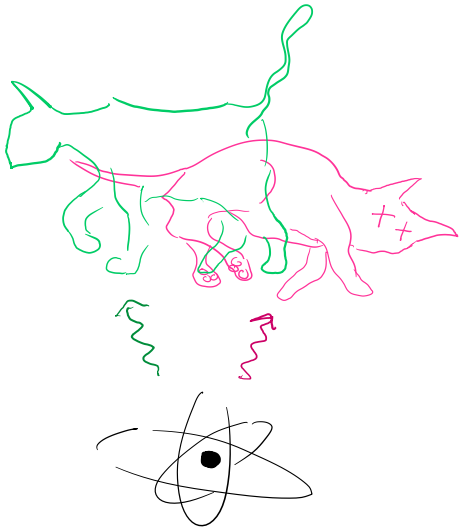

Featuring

Alexis Toumi
Giovanni de Felice
Anna Pearson
Robin Lorenz
Richie Yeung
Ian Fan

and

Dimitri Kartsaklis
Stephen Clark
Bob Coecke

plus

Douglas Brown
Carys Harvey
Richard Wolf

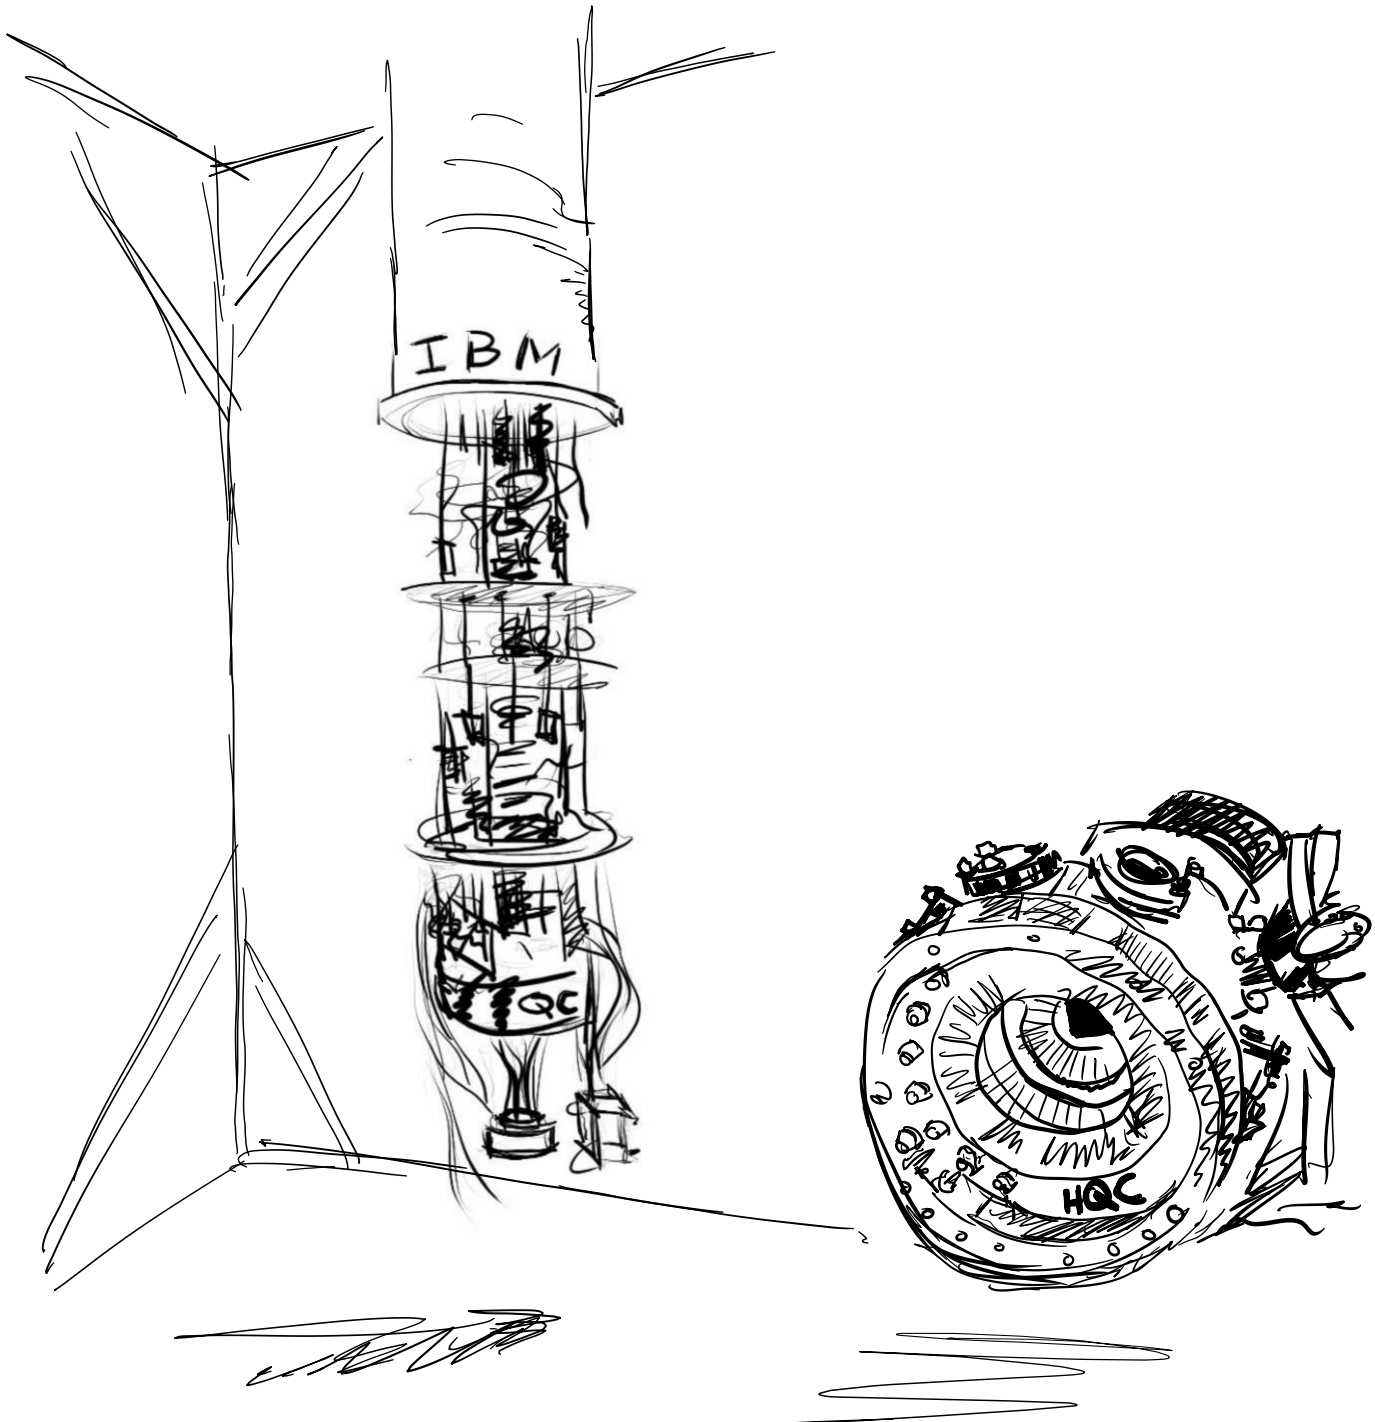

the philosophical slide

NLP ⊆ AI

“The limits of my language mean the limits of my world“
- Ludwig Wittgenstein

“Language faculty is what separates us from other species”
- Noam Chomsky

“NLP is AI-hard”
Turing test : verify intelligent behaviour via language!

the apps slide

NLP : applications

FORTUNE BUSINESS INSIGHTS

HOME INDUSTRIES SERVICES INSIGHTS ABOUT

Technology & Media / Natural Language Processing (NLP) Market

"Smart Growth

Natural Language Processing (NLP) Market Size, Share & COVID-19 Impact Analysis, By Deployment (On-Premises, Cloud, Hybrid), By Enterprise Size (SMEs, and Large Enterprises), By Technology (Interactive Voice Response (IVR), Optical Character Recognition (OCR), Text Analytics, Speech Analytics, Classification and Categorization), By Industry Vertical (Healthcare, Retail, High Tech, and Telecom, BFSI) and Regional Forecast, 2021-2028

Region : Global | Format: PDF | Report ID: FBI101933 Share

Summary | Table of Content | Market Segmentation | Methodology | Infographics | Request Sample PDF

KEY MARKET INSIGHTS

Listen to Audio Version

The global natural language processing (NLP) market size was USD 16.53 billion in 2020. The global impact of COVID-19 has been unprecedented and staggering, with NLP solutions/services witnessing a positive demand shock across all regions amid the pandemic. Based on our analysis, the global market exhibited a significant growth of 26.5% in 2020 compared to year-on-year growth during 2017-2019. The market is projected to grow from USD 20.98 billion in 2021 to USD 127.26 billion in 2028 at a CAGR of 29.4% during the 2021-2028 period. The steady rise in CAGR is attributable to this market's demand and growth, returning to pre-pandemic levels once the pandemic is over.

Natural language processing is a communication bridge between technology and humans. With its help, the input instructions provided by humans are easy to understand by computers. It is also considered a branch of artificial intelligence (AI) that understands, interprets, and manipulates human languages. The adoption of natural language

- Dialogue
- Text mining
- Translation
- Text ←→ Speech
- Language generation
- Classification
- Bioinformatics
- Literature and art

NLP : state of the art

GPT-3

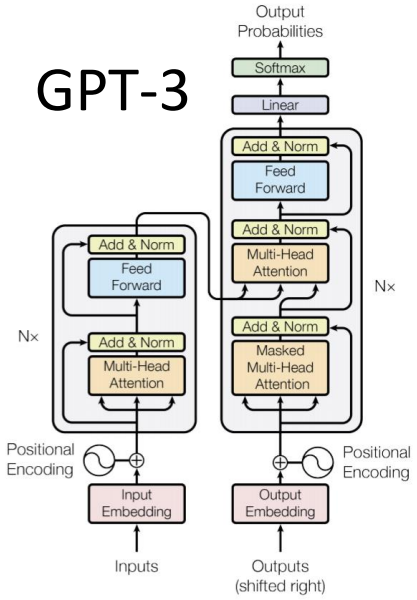

~~large~~ huge language models:

capture prob distros over symbol seqs

No coherent world model!

Human Brain: 200 billion neurons, 125 trillion synapses (just in the cerebral cortex).
New larger language models with 100 trillion parameters.

"but":

nature

Explore content · Journal information · Publish with us · Subscribe

nature › news feature › article

NEWS FEATURE · 03 MARCH 2021

Robo-writers: the rise and risks of language-generating AI

A remarkable AI can write like humans – but with no understanding of what it's saying.

Matthew Hutson

Quanta magazine

Physics Mathematics Biology Computer Science Topics Archive

QUANTIZED COLUMNS

What Does It Mean for AI to Understand?

60

It's simple enough for AI to seem to comprehend data, but devising a true test of a machine's knowledge has proved difficult.

MIT Technology Review Topics

GETTY

Artificial intelligence / Machine learning

OpenAI's new language generator GPT-3 is shockingly good—and completely mindless

The AI is the largest language model ever created and can generate amazing human-like text on demand but won't bring us closer to true intelligence.

by **Will Douglas Heaven** July 20, 2020

A Structural Approach

Untapped potential in formal, rigorous, structural, hybrid approaches:
old-school (symbolic) with new-school methods (distributional)
plus ideas from neuroscience and cognitive science.

“What is the **mechanics of meaning**?”

Focus on alternative models rather than brute compute.

“let’s be physicists about this”

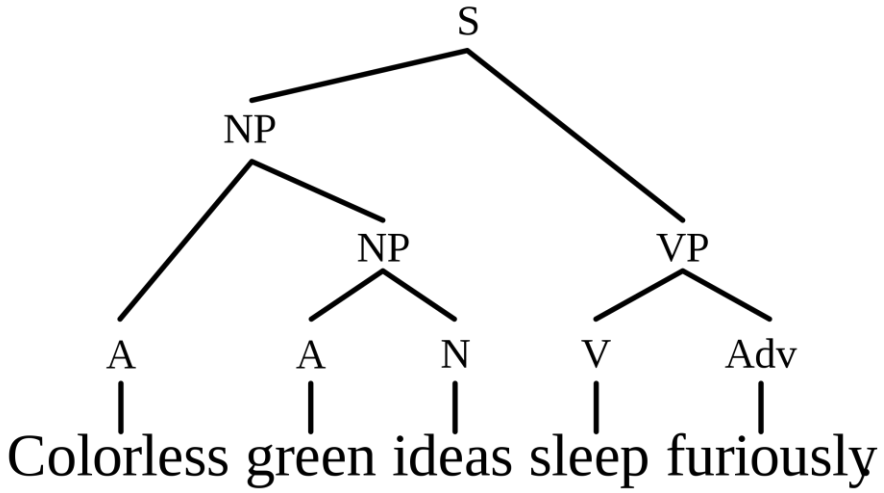

N. Chomsky

Enter the DisCo

Steve

Bob

Categorical Compositional Distributional (DisCoCat) model of meaning (2008)

Bob Coecke
Mehrnoosh Sadrzadeh
Steve Clark
[1003.4394]

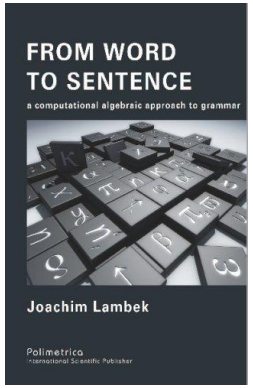

Input:
word meanings

Process:
grammar that *composes* them

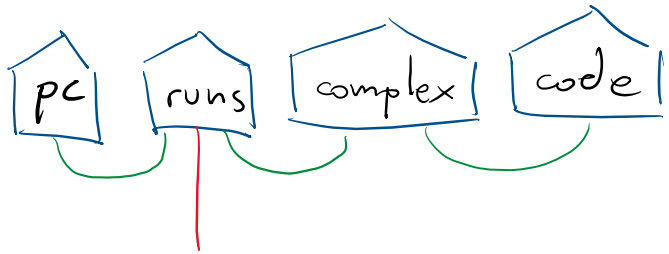

Output:
sentence meaning

String diagrams

states

processes

effects

scalar

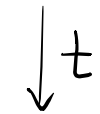

ex. diagram:

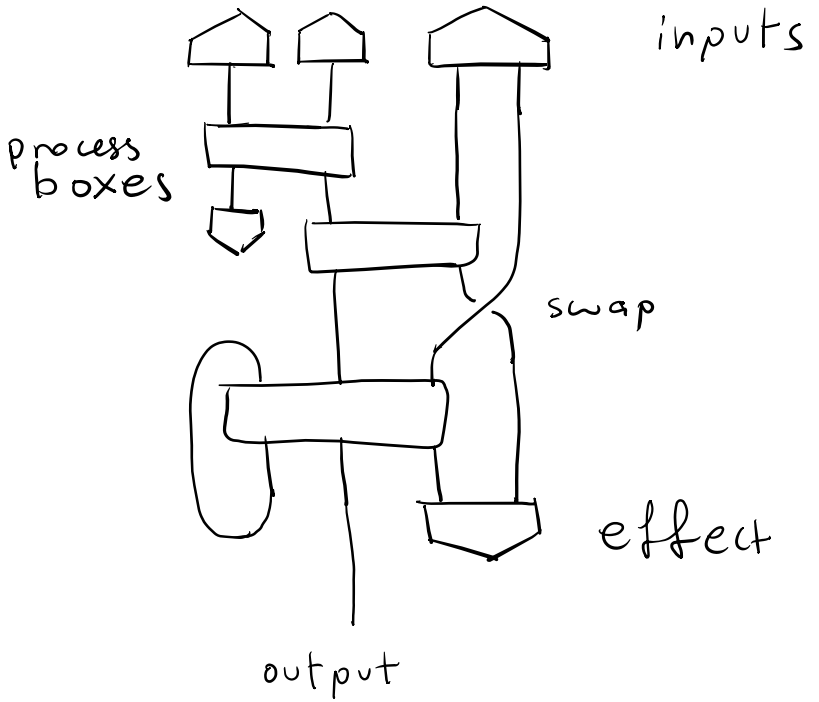

=

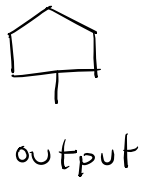

composition:

parallel :

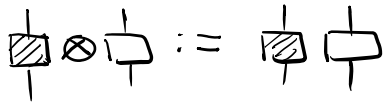

sequential :

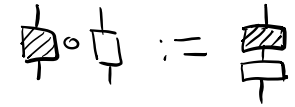

Meaning flow as string diagrams

Parser assigns pregroup types to words :

$$w \mapsto t_{i_1}^{z_{i_1}} \cdot t_{i_2}^{z_{i_2}} \cdot \ldots \cdot t_{i_{|w|}}^{z_{i_{|w|}}} \quad , \quad z_i \in \mathbb{Z} \simeq \{\ldots, lll, ll, l, -, r, rr, rrr, \ldots\} \quad , \quad t_i \in B$$

Types reduce as : $t^z \cdot t^{z+1} \to 1$, $1 \to t^{z+1} \cdot t^z$

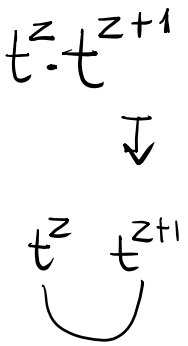

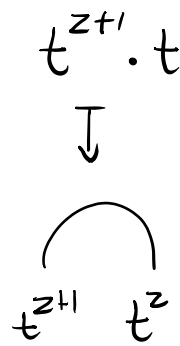

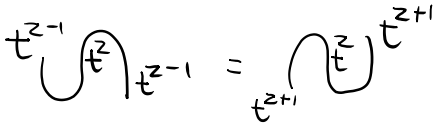

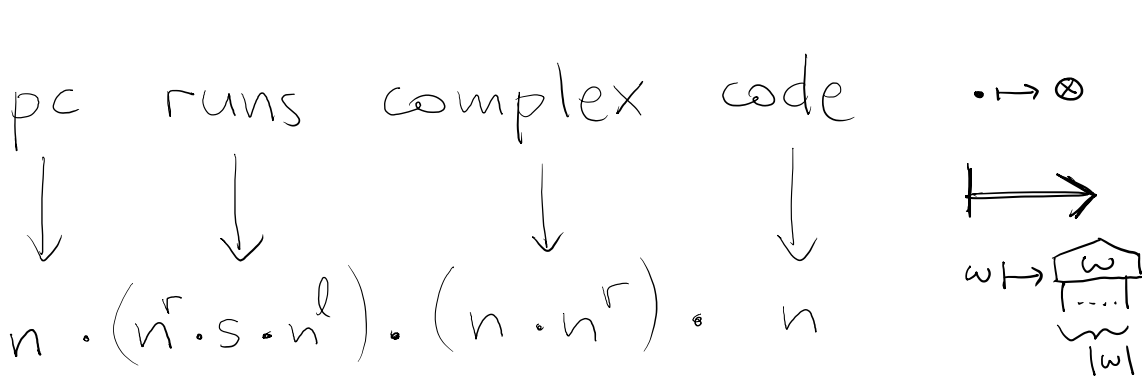
pc runs complex code
n · (n^r · s · n^l) · (n · n^r) · n
• ↦ ⊗
ω ↦ ω
|ω|

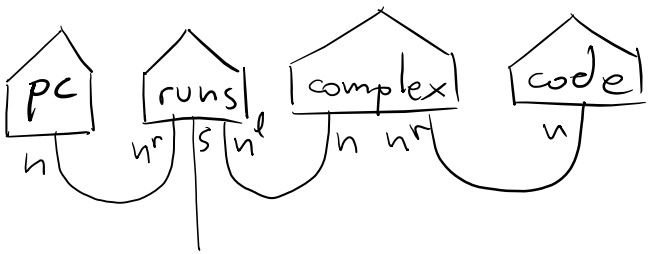
pc
runs
complex
code
n
n^r
s
n^l
n
n^r
n

Reinterpret the String Diagram

"Give semantics to the boxes"

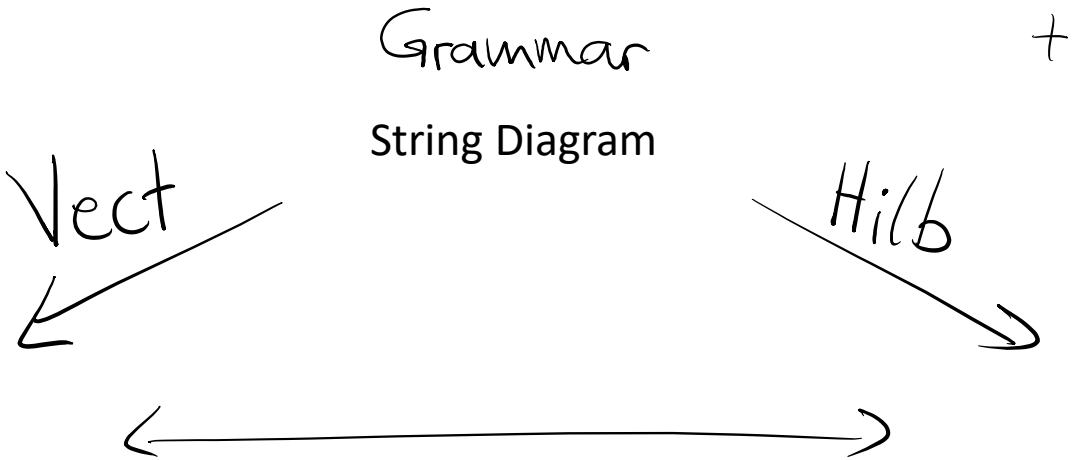

Tensor Network:
runs on CPU

Quantum Circuit:
runs on QPU
with the help of a
quantum compiler (tket)

states : vectors
processes: linear maps
effects : co-vectors

states: quantum states
processes : quantum meps
effects : quantum measurements and post-selections

Reinterpret the String Diagram as a quantum circuit: Define a Functor

ignore adjoints : $t^z \longmapsto d_t = 2^{9t}$

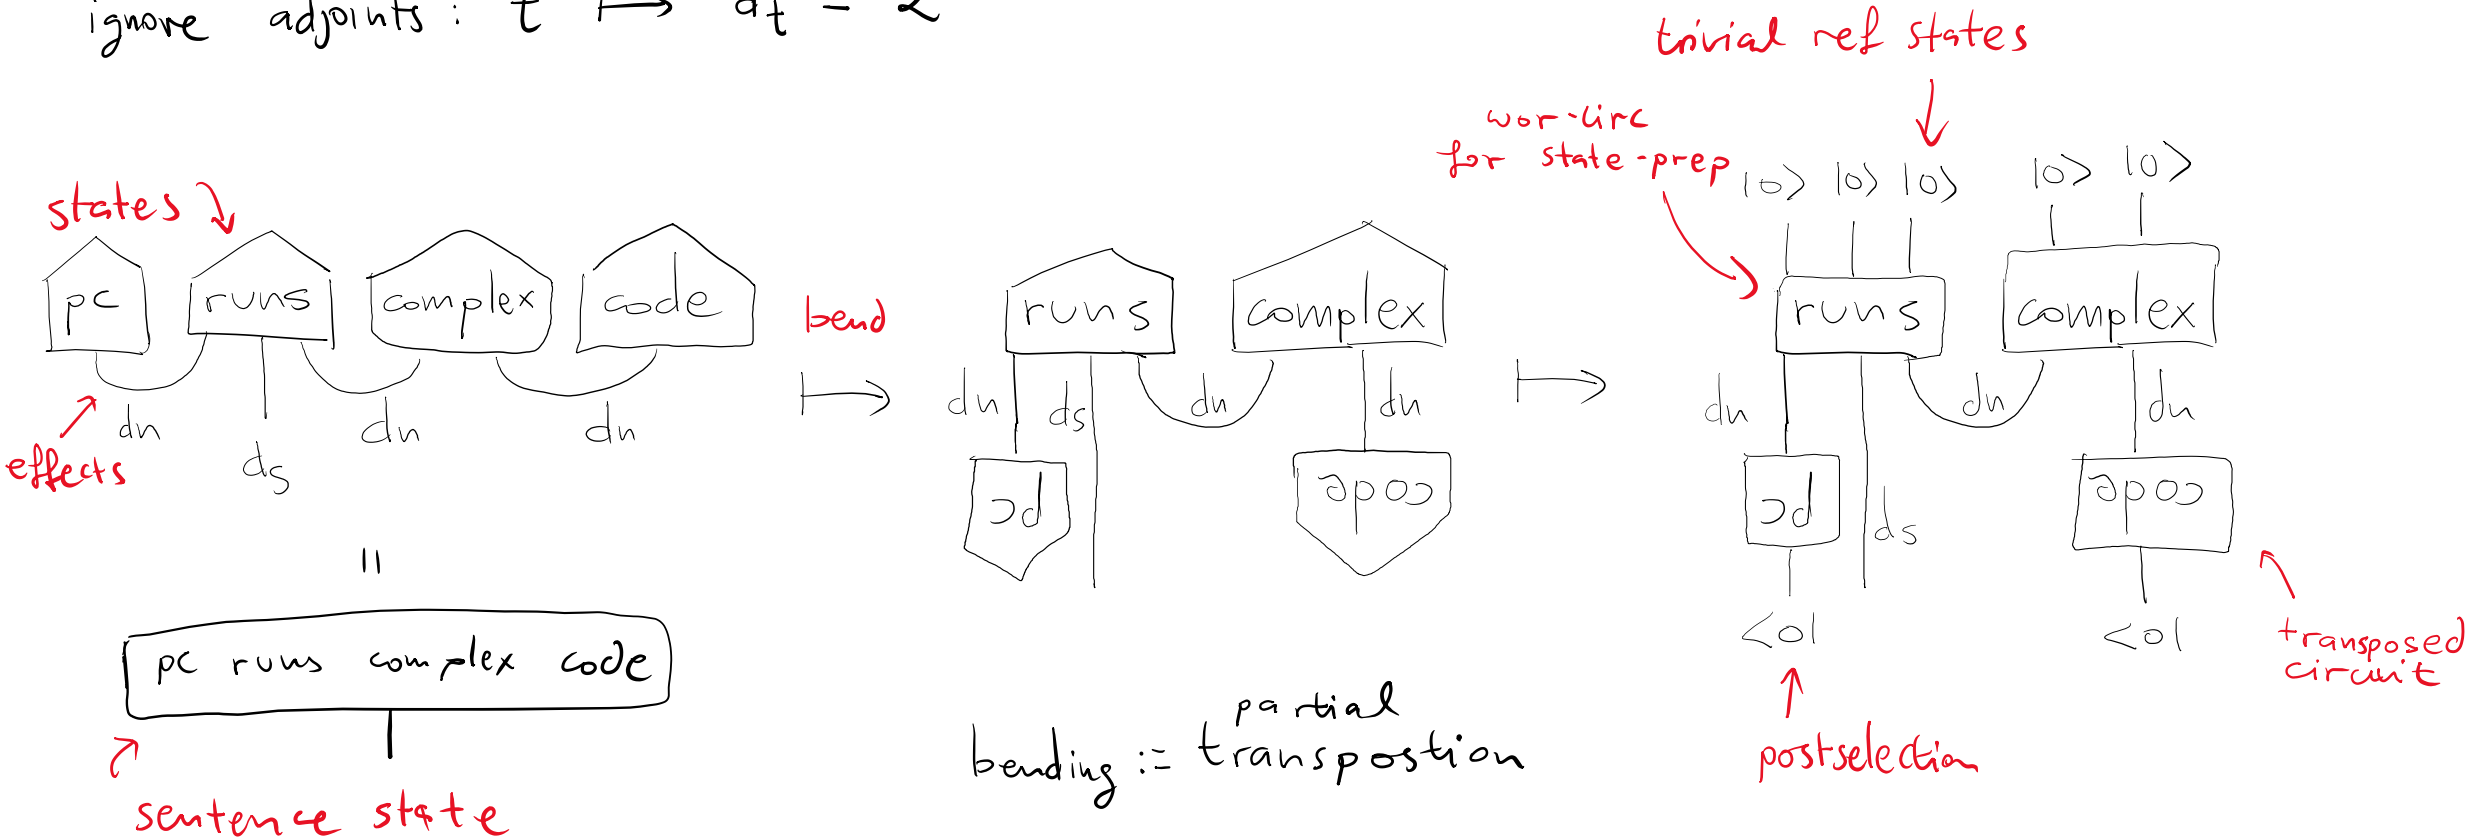

Choices of Quantum Circuit Ansaetze

hyperparameters: $d_n = 2^{q_n}$, $d_s = 2^{q_s}$, d (IQP depth)

ex: $q_n, q_s = 1$, $d = 2$:

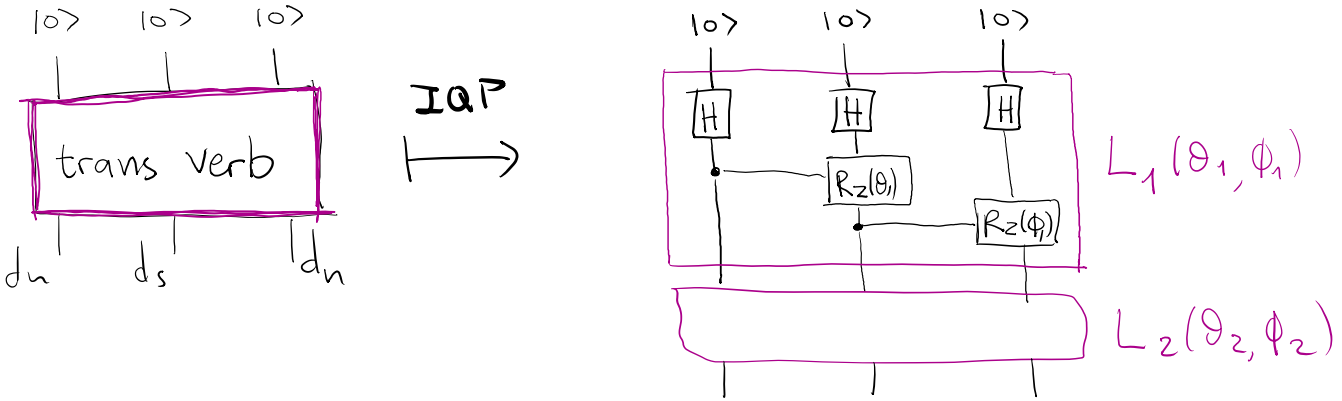

word - circuits

grammar - effects

lambeq: An Efficient High-Level Python Library for Quantum NLP

Dimitri Kartsaklis Ian Fan Richie Yeung Anna Pearson
Robin Lorenz Alexis Toumi Giovanni de Felice
Konstantinos Meichanetzidis Stephen Clark Bob Coecke

Cambridge Quantum Computing
17 Beaumont Street, Oxford, OX1 2NA, UK

{dimitri.kartsaklis;ian.fan;richie.yeung;anna.pearson;
robin.lorenz;alexis.toumi;giovanni.defelice;
kmei;steve.clark;bob.coecke}@cambridgequantum.com

Abstract

We present `lambeq`, the first high-level Python library for Quantum Natural Language Processing (QNLP). The toolkit offers a detailed hierarchy of modules and classes implementing all stages of a pipeline for converting sentences to string diagrams, tensor networks, and quantum circuits ready to be used on a quantum computer. `lambeq` supports syntactical parsing, rewriting and simplification of string diagrams, ansatz creation and manipulation, as well as a number of compositional models for preparing quantum-friendly representations of sentences, employing various degrees of syntax sensitivity. We present the generic architecture and we describe the most important modules in detail, demonstrating the usage with illustrative examples. Further, we test the toolkit in practice by using it in order to perform a number of experiments on simple NLP tasks, implementing both classical and quantum pipelines, with success.

Open-source
Apache 2.0 license

[2110.04236]
https://github.com/CQCL/lambeq

Forbes

Cambridge Quantum Makes Quantum Natural Language Processing A Reality

Paul Smith-Goodson Contributor
Moor Insights and Strategy Contributor Group
Cloud
Analyst-in-residence, Quantum Computing

Listen to article 6 minutes

Online demo:
https://qnlp.cambridgequantum.com/generate.html

λambeq enables the automatic deployment of large-scale compositional language models on quantum computers

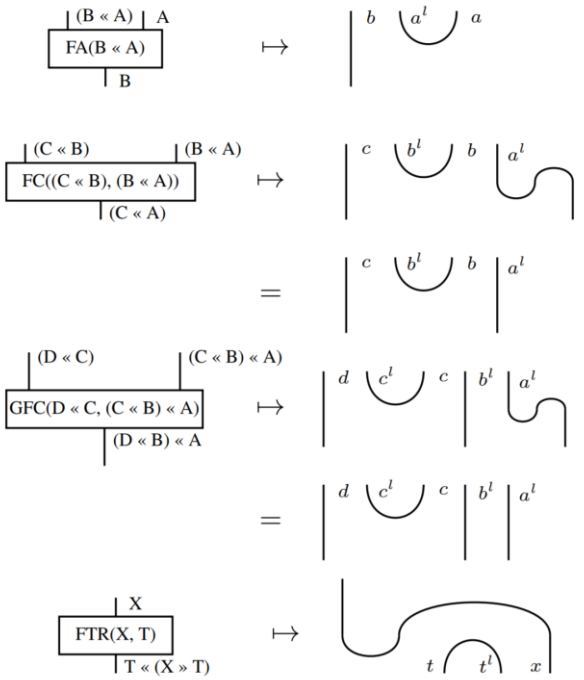

syntax & grammar | middleware | string diagrams

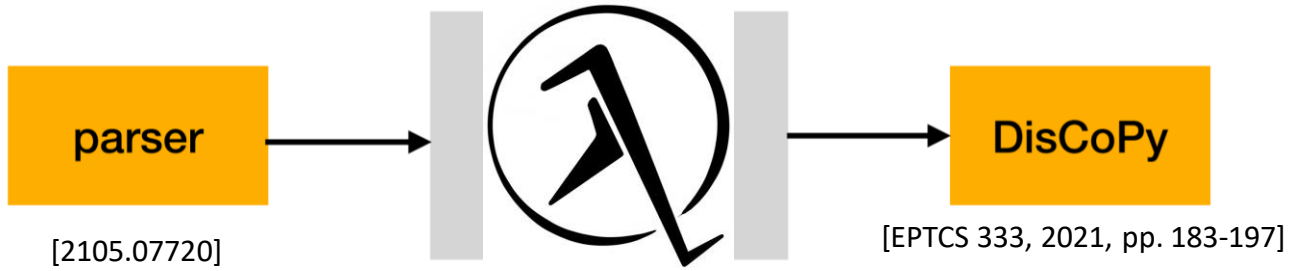

[2105.07720]

[EPTCS 333, 2021, pp. 183-197]

Pipeline

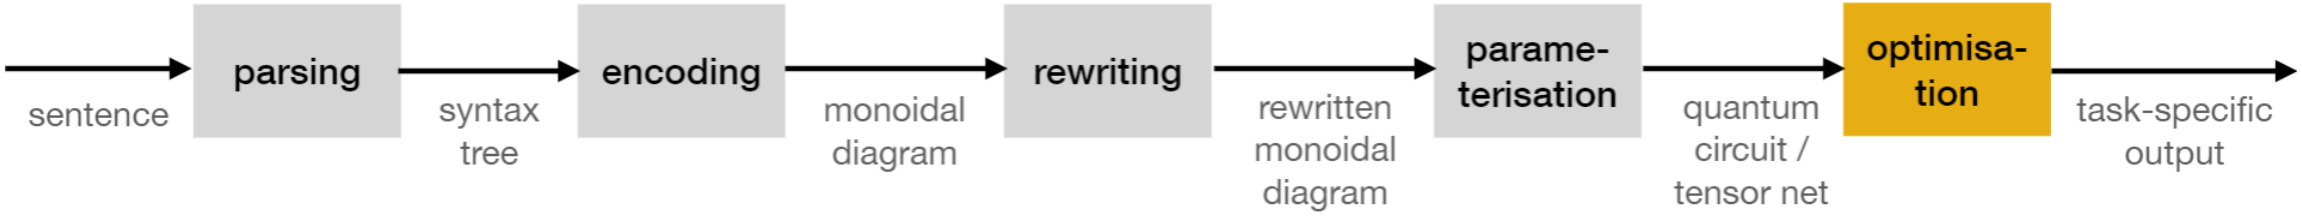

Example:
sentence classification

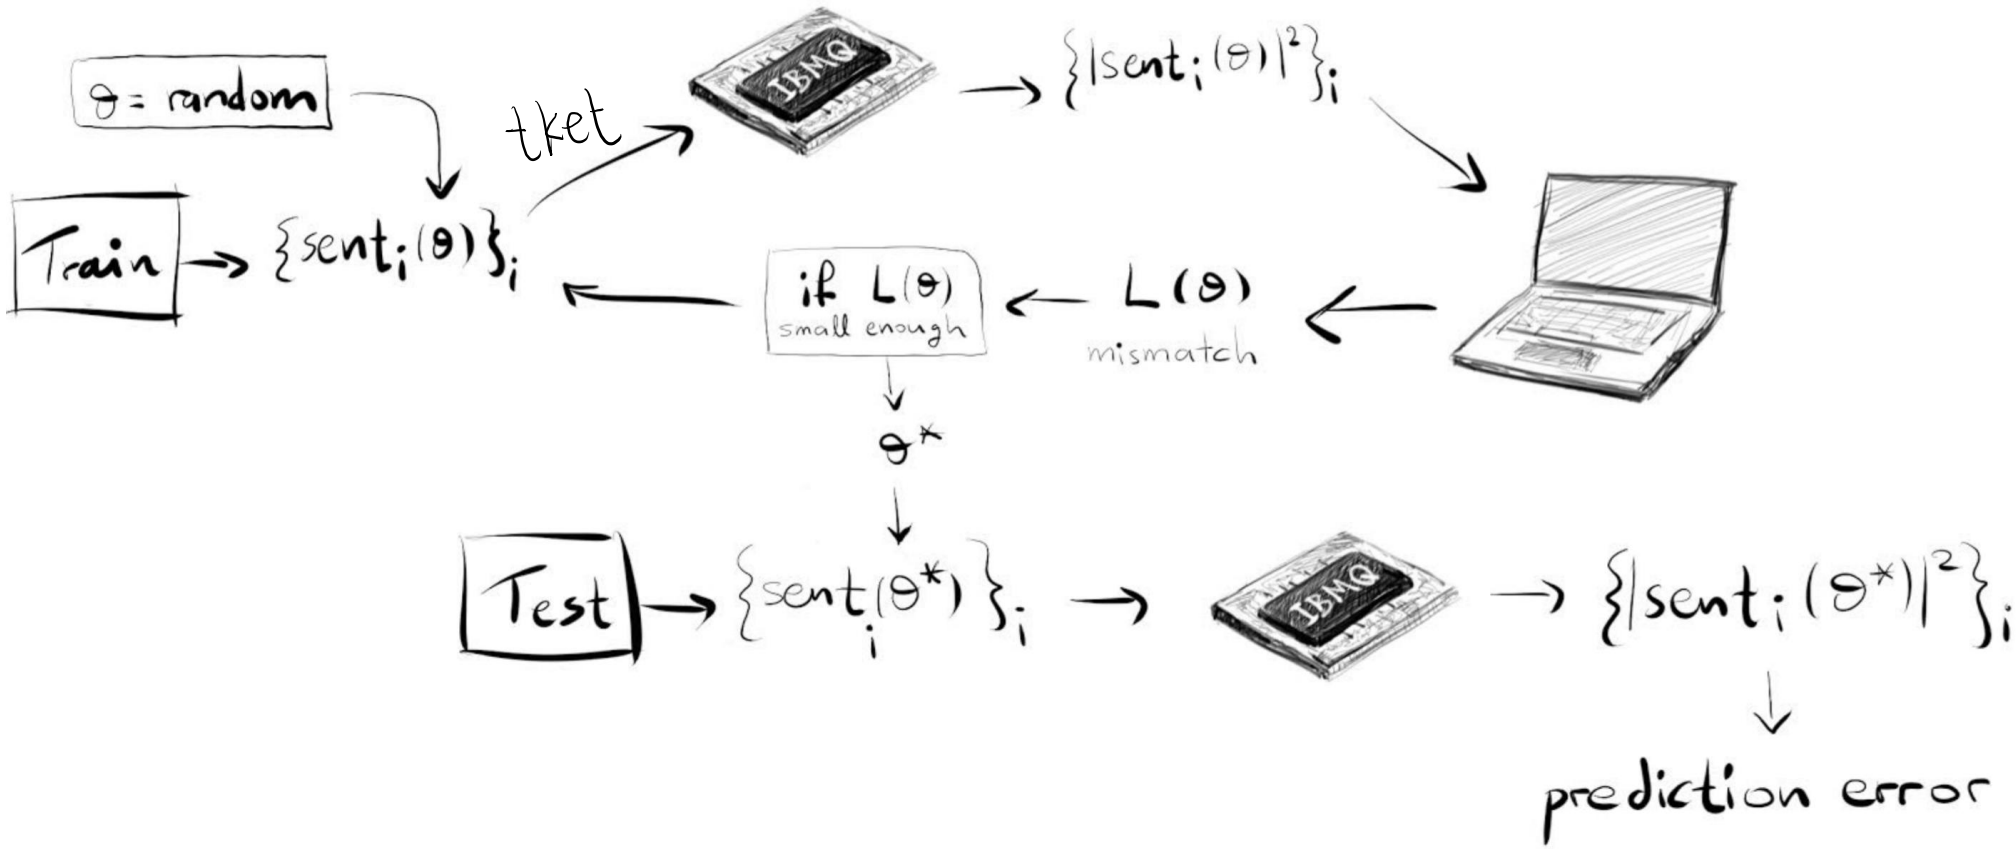

Tasks

Truth value: $q_n = 1, q_s = 0$
Diagram is a scalar

False:
(Juliet kills Romeo who dies, 0)
(Romeo kills Juliet, 0)
(Romeo who kills Juliet dies, 0)
…

True:
(Juliet dies, 1)
(Romeo who dies loves Juliet, 1)
(Romeo who kills Romeo dies, 1)
…

Topic: $q_n = 1, q_s = 1$
Diagram is a state

Cooking:
(Skilful man prepares sauce, 0)
(Woman cooks tasty meal , 0)
(Skilful person prepares meal, 0)
…

Technology:
(Skilful woman debugs program , 1)
(Man prepares useful application, 1)
(Person debugs useful software, 1)
…

Notes:

Intra-sentence correlations are ‘quantum’: due to grammar
Inter-sentence correlations are ‘classical’: due to shared words

Classical Simulation

False:
(Juliet kills Romeo who dies, 0)
(Romeo kills Juliet, 0)
(Romeo who kills Juliet dies, 0)
…
True:
(Juliet dies, 1)
(Romeo who dies loves Juliet, 1)
(Romeo who kills Romeo dies, 1)
…

basinhopping over Nelder-Mead

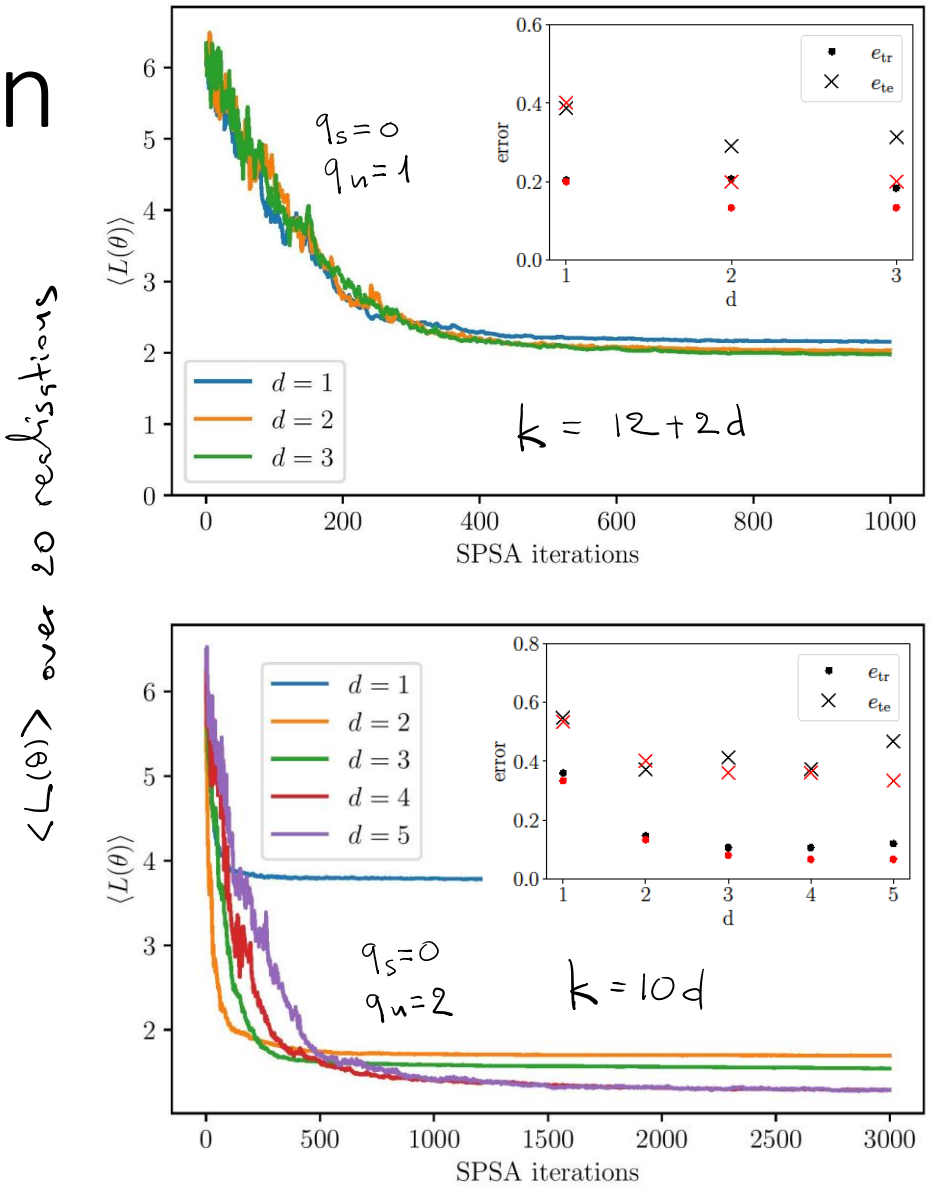

Classical Simulation

Cooking:
(Skilful man prepares sauce, 0)
(Woman cooks tasty meal , 0)
(Skilful person prepares meal, 0)
…
Technology:
(Skilful woman debugs program , 1)
(Man prepares useful application, 1)
(Person debugs useful software, 1)
…

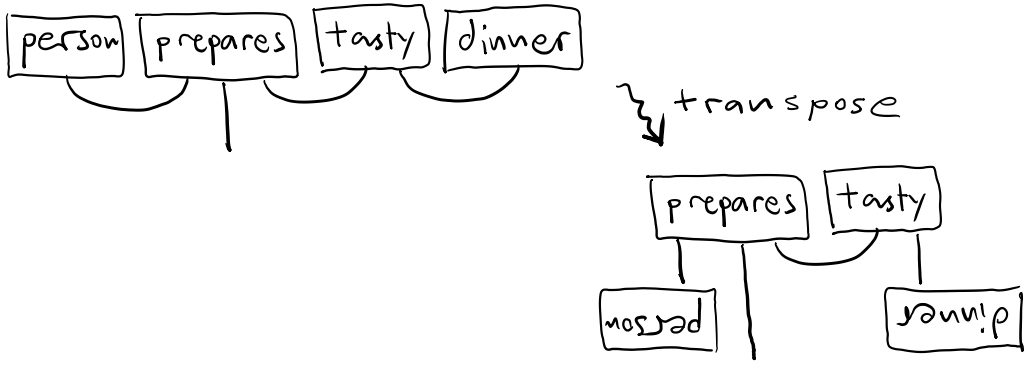

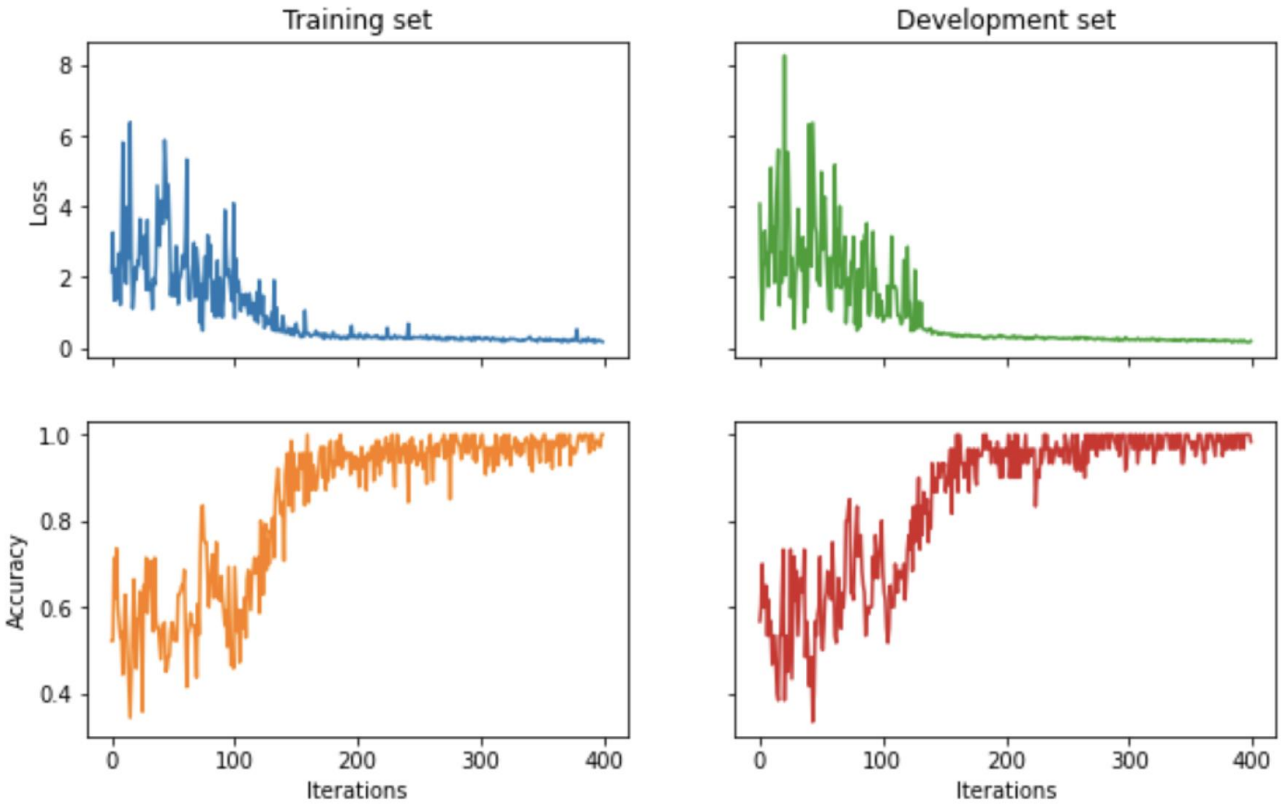

Quantum Runs on IBM

Truth Value

[2012.03756]

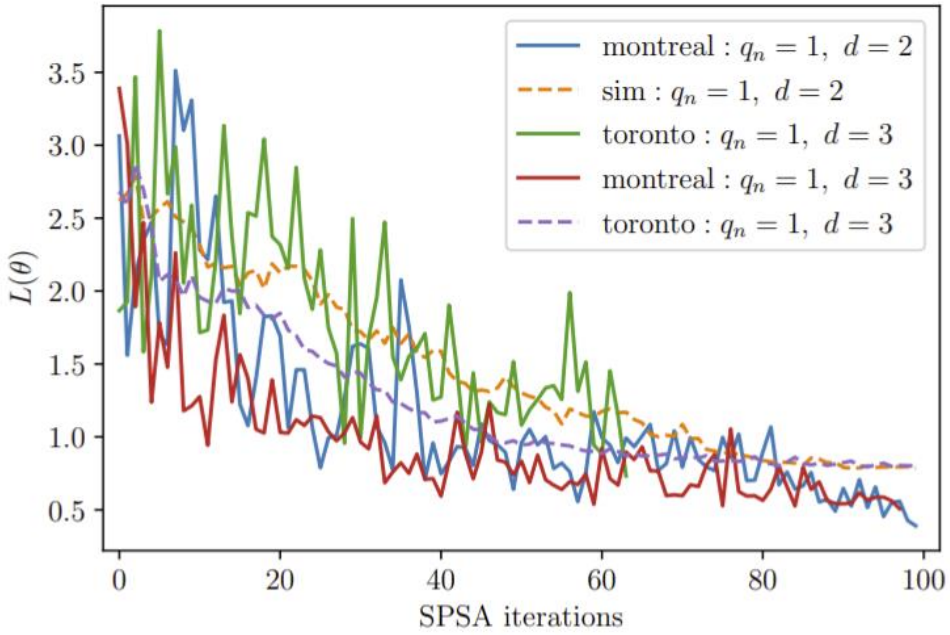

Ibmq_montreal

Topic

[2102.12846]

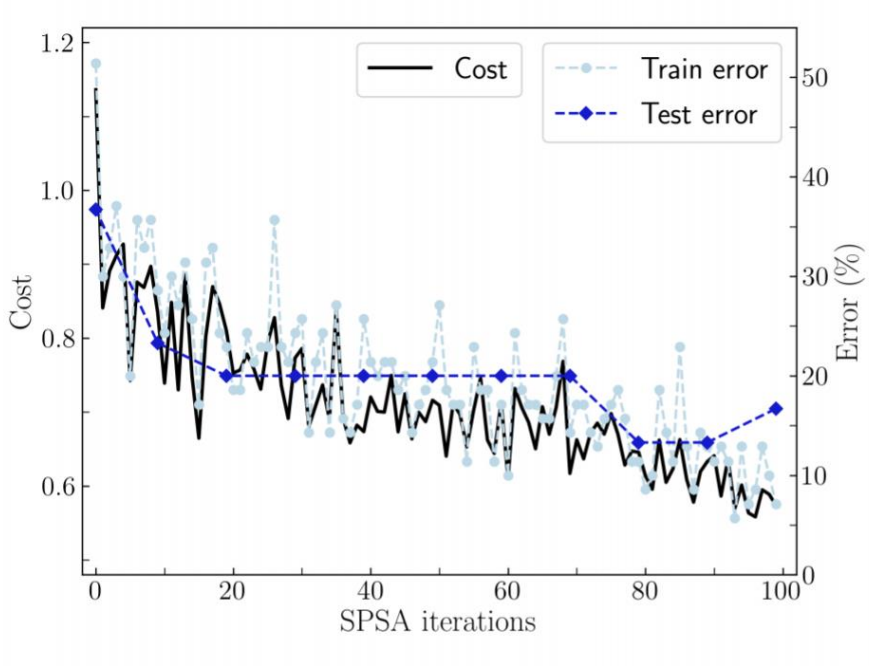

Ibmq_bogota
And system H1

Quantum Runs on H1

Same topic classification task.

Train classically,
not even considering noise.

Then Test on H1.

Acc 97%.

Just this past week, H1-2 measured a Quantum Volume of 2,048 (211), setting a new bar on the highest quantum volume ever measured on a quantum computer. The performance of the H1 generation of quantum computers continues to achieve the 10X per year increase that was announced in March 2020.

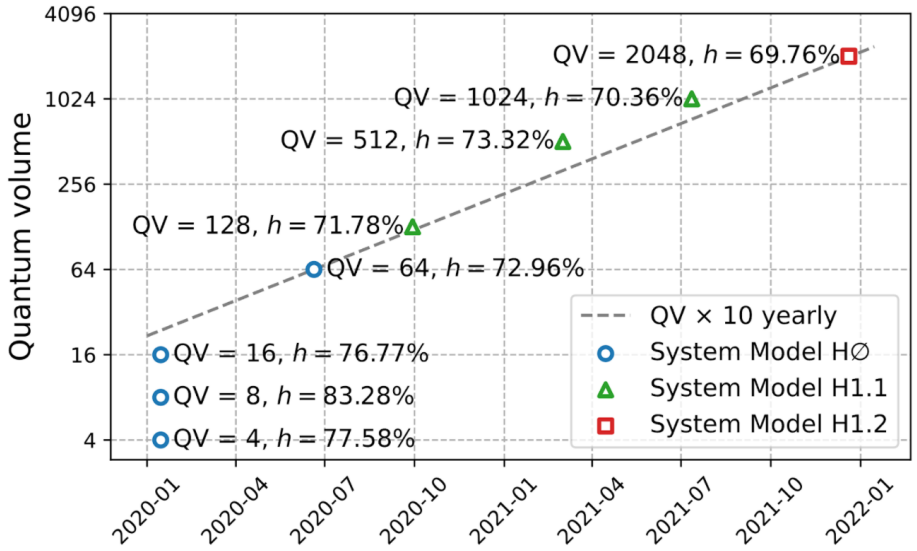

The Data

The average single-qubit gate fidelity for this milestone was 99.996(2)%, the average two-qubit gate fidelity was 99.77(9)%, and state preparation and measurement (SPAM) fidelity was 99.61(2)%. We ran 2,000 randomly generated quantum volume circuits with 5 shots each, using standard optimization techniques to yield an average of 122 two-qubit gates per circuit.

[https://www.quantinuum.com/pressrelease/demonstrating-benefits-of-quantum-upgradable-design-strategy-system-model-h1-2-first-to-prove-2-048-quantum-volume]

Future

Other grammar reductions instead of pregroups
Ex. Trees:

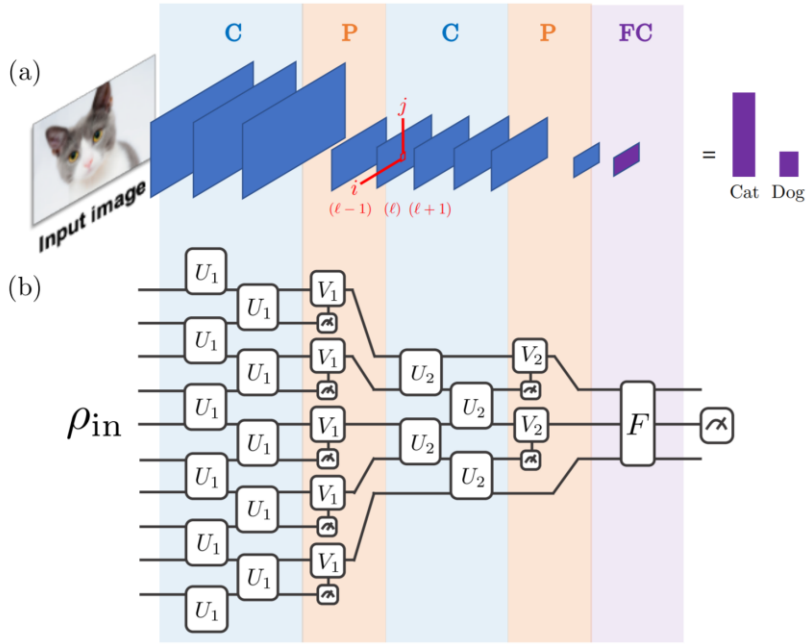

[1810.03787]

By staying in CCG land,
we can define Syntax-informed QCNNs
by giving PQC ansaetze
to the 2-to-1 boxes in the tree

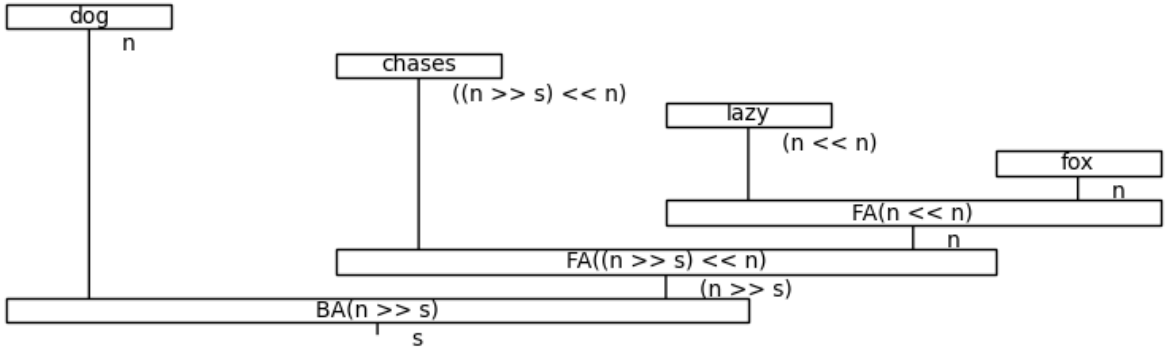

Future

benchmark more **baselines**:
key question is
“what does grammar really buy you?”

what’s the quantum version of attention?

language generation

example baselines:

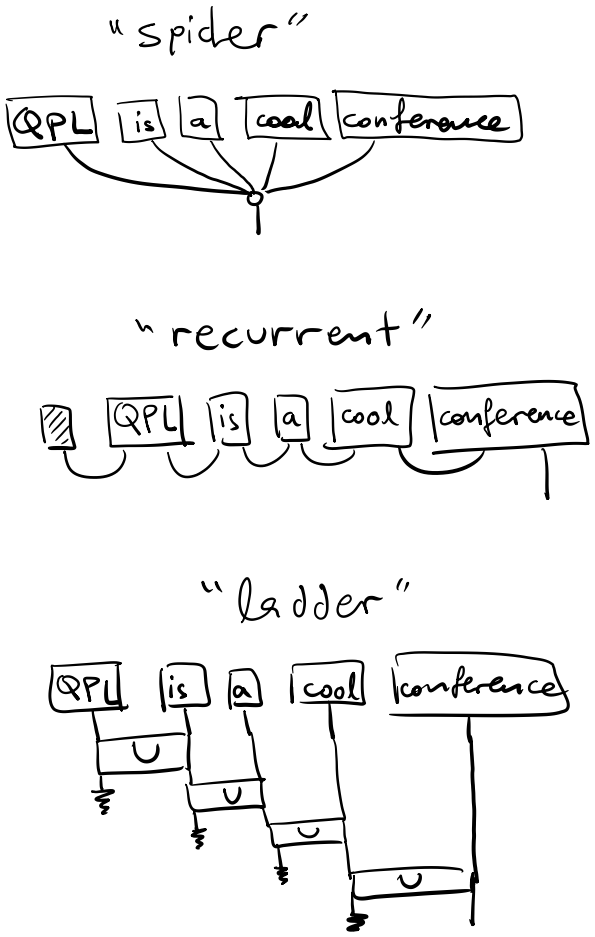

Future

DisCoCirc: sentence composition;
Bob Coecke: "*get quadvantage a la quantum simulation*".
So IF this model works well for some tasks, one would need a QC to run it.

"Nature is quantum so we should use quantum systems to simulate her"

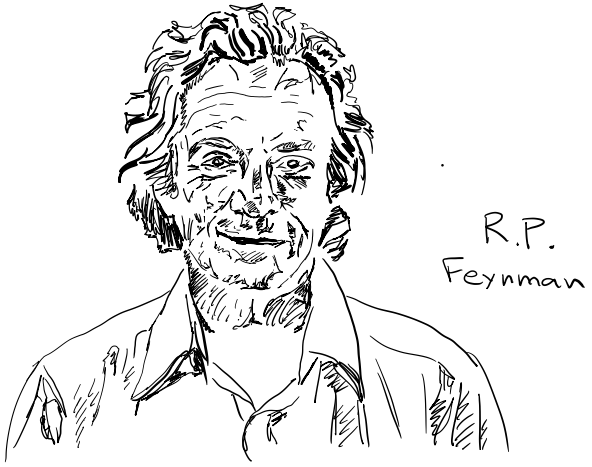

R.P. Feynman

define functor into qcircs

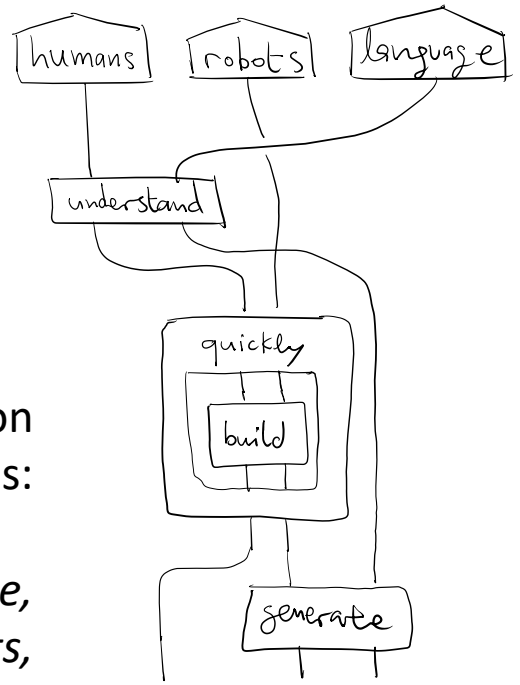

giving boxes qcirc semantics gives a classically hard model

Circuit from composition of three sentences:

`Humans understand language, humans quickly build robots, robots generate language'.

Thx :)