Compositional Quantum Natural Language Processing
Theoretical Foundations and Experimental Implementation

Muhammad Hamza Waseem
DPhil Physics, University of Oxford
Research Scientist at Quantinuum

Mathematics and Physics of Quantum Computing and Quantum Learning
Porquerolles, France

May 23-28, 2025

QUANTINUUM

UNIVERSITY OF OXFORD

Advantages of language circuits:

- evolving meanings
- strips off language-dependent overheads
- interpretable, verifiable, safer, fairer...

Advantages of language circuits:

- evolving meanings
- strips off language-dependent overheads
- interpretable, verifiable, safer, fairer...
- cheaper

Advantages of language circuits:

- evolving meanings
- strips off language-dependent overheads
- interpretable, verifiable, safer, fairer...
- cheaper
- the best of both worlds: combinable with reasoning

What is compositionality?

Frege compositionality in formal linguistics:

Frege compositionality in formal linguistics:

Meaning of a whole (cf. sentence) should only depend on meanings of its parts (cf. words) and how they are fitted together (cf. grammar).

Frege compositionality in formal linguistics:

Meaning of a whole (cf. sentence) should only depend on meanings of its parts (cf. words) and how they are fitted together (cf. grammar).

$\Longrightarrow$ bottom-up meaning flow

Frege compositionality in formal linguistics:

Meaning of a whole (cf. sentence) should only depend on meanings of its parts (cf. words) and how they are fitted together (cf. grammar).

$\Longrightarrow$ bottom-up meaning flow

But there is also **Frege's context principle**:

Frege compositionality in formal linguistics:

Meaning of a whole (cf. sentence) should only depend on meanings of its parts (cf. words) and how they are fitted together (cf. grammar).

$\Longrightarrow$ bottom-up meaning flow

But there is also **Frege's context principle**:

Never ask for word meaning in isolation, but only in the context of a sentence.

Frege compositionality in formal linguistics:

Meaning of a whole (cf. sentence) should only depend on meanings of its parts (cf. words) and how they are fitted together (cf. grammar).

$\Longrightarrow$ bottom-up meaning flow

But there is also **Frege's context principle**:

Never ask for word meaning in isolation, but only in the context of a sentence.

$\Longrightarrow$ top-down meaning flow

These Alice's can get disambiguated by context:

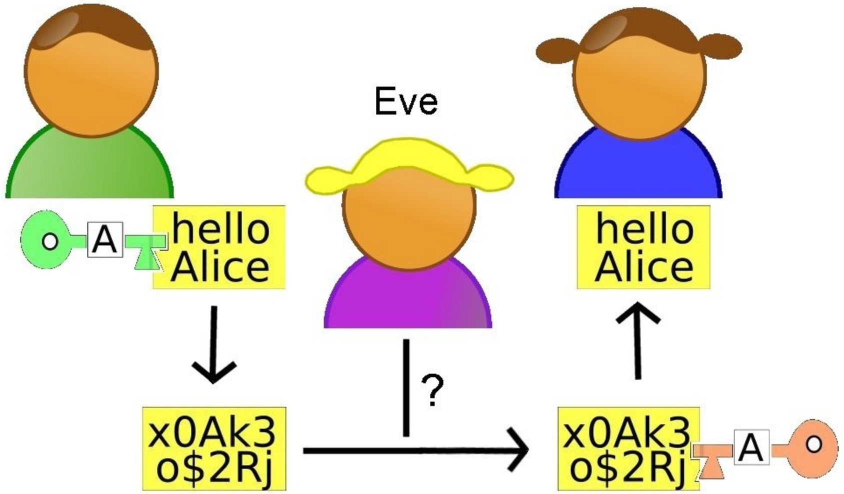

The ambiguity can also intertwine grammar and meaning:

The ambiguity can also intertwine grammar and meaning:

Respectively:

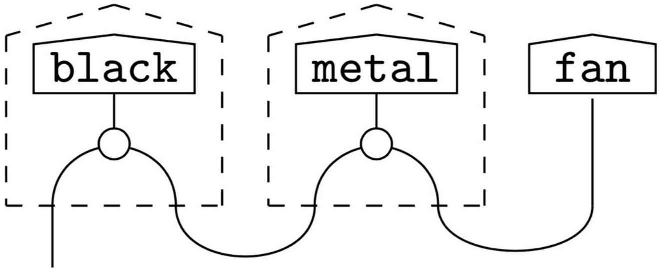

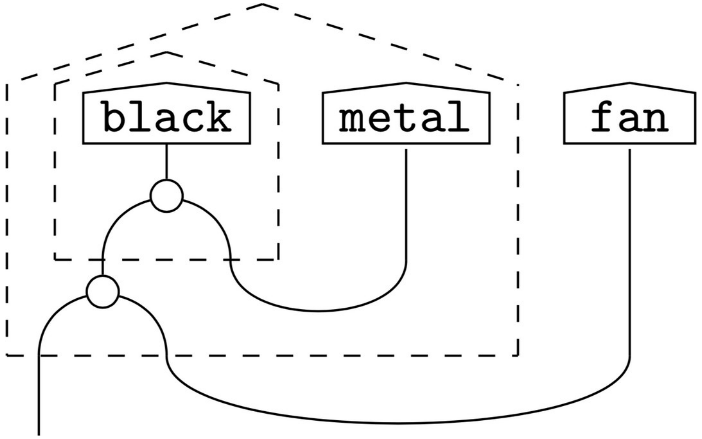

– *process theory (from dodo-book)* –

A *process theory* consists of:

- systems S represented by wires,
- processes P represented by boxes, with systems in S as inputs/outputs,
- composition of processes, represented by wirings, resulting in a process D.

– process theory (from dodo-book) –

A *process theory* consists of:

- systems S represented by wires,
- processes P represented by boxes, with systems in S as inputs/outputs,
- composition of processes, represented by wirings, resulting in a process D.

Could be generalised further e.g.:

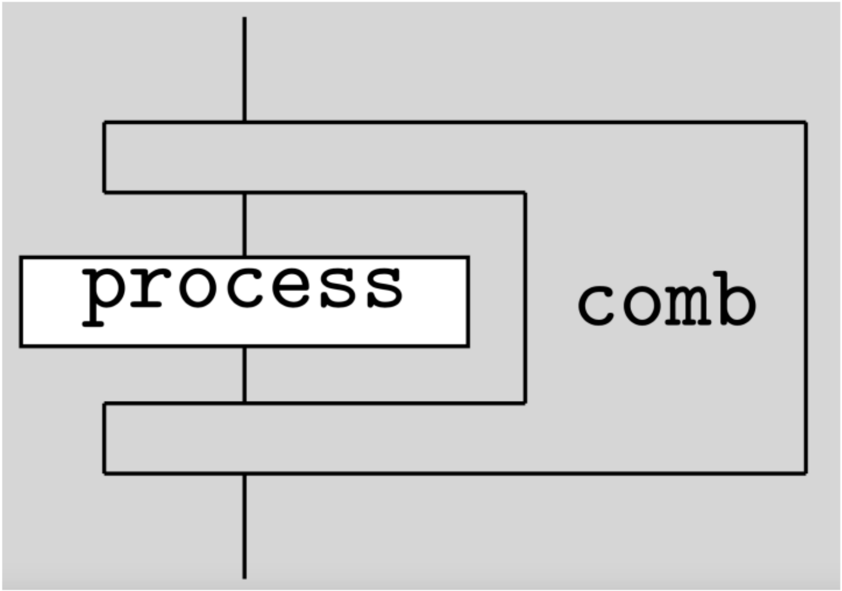

A **Schrödinger compositional theory** is a process theory with diagrams:

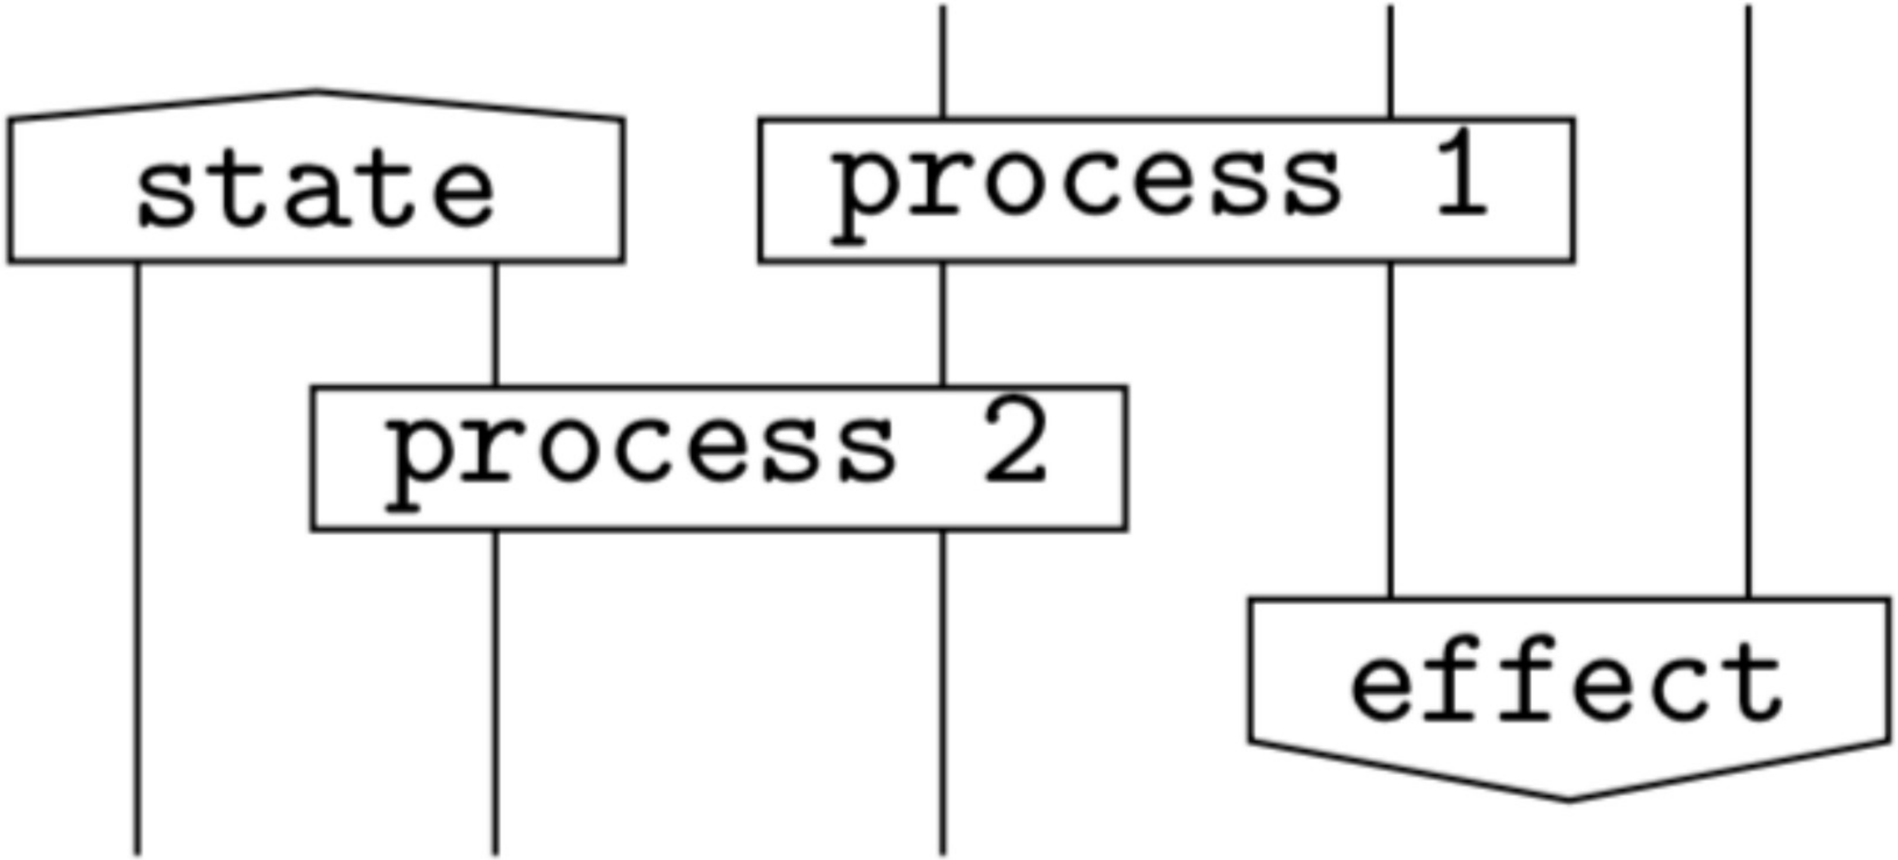

such that:

- Composition is non-trivial, i.e. a whole cannot be decomposed meaningfully.
- All ingredients have clear meaningful ontological counterparts in reality.

Whitehead-compositional theory is a process theory with diagrams:

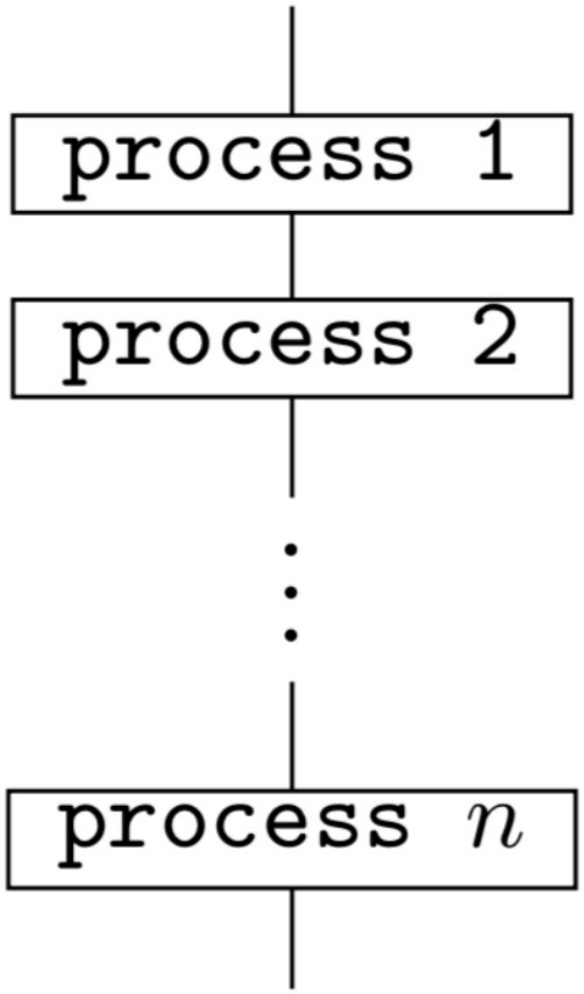

such that:

- All ingredients have clear meaningful ontological counterparts in reality.

arXiv.org > math > arXiv:2110.05327

Search...

Help | Advar

Mathematics > Category Theory

[Submitted on 11 Oct 2021]

Compositionality as we see it, everywhere around us

Bob Coecke

There are different meanings of the term "compositionality" within science: what one researcher would call compositional, is not at all compositional for another researcher. The most established conception is usually attributed to Frege, and is characterised by a bottom-up flow of meanings: the meaning of the whole can be derived from the meanings of the parts, and how these parts are structured together.
Inspired by work on compositionality in quantum theory, and categorical quantum mechanics in particular, we propose the notions of Schrodinger, Whitehead, and complete compositionality. Accounting for recent important developments in quantum technology and artificial intelligence, these do not have the bottom-up meaning flow as part of their definitions.
Schrodinger compositionality accommodates quantum theory, and also meaning-as-context. Complete compositionality further strengthens Schrodinger compositionality in order to single out theories like ZX-calculus, that are complete with regard to the intended model. All together, our new notions aim to capture the fact that compositionality is at its best when it is `real', `non-trivial', and even more when it also is `complete'.
At this point we only put forward the intuitive and/or restricted formal definitions, and leave a fully comprehensive definition to future collaborative work.

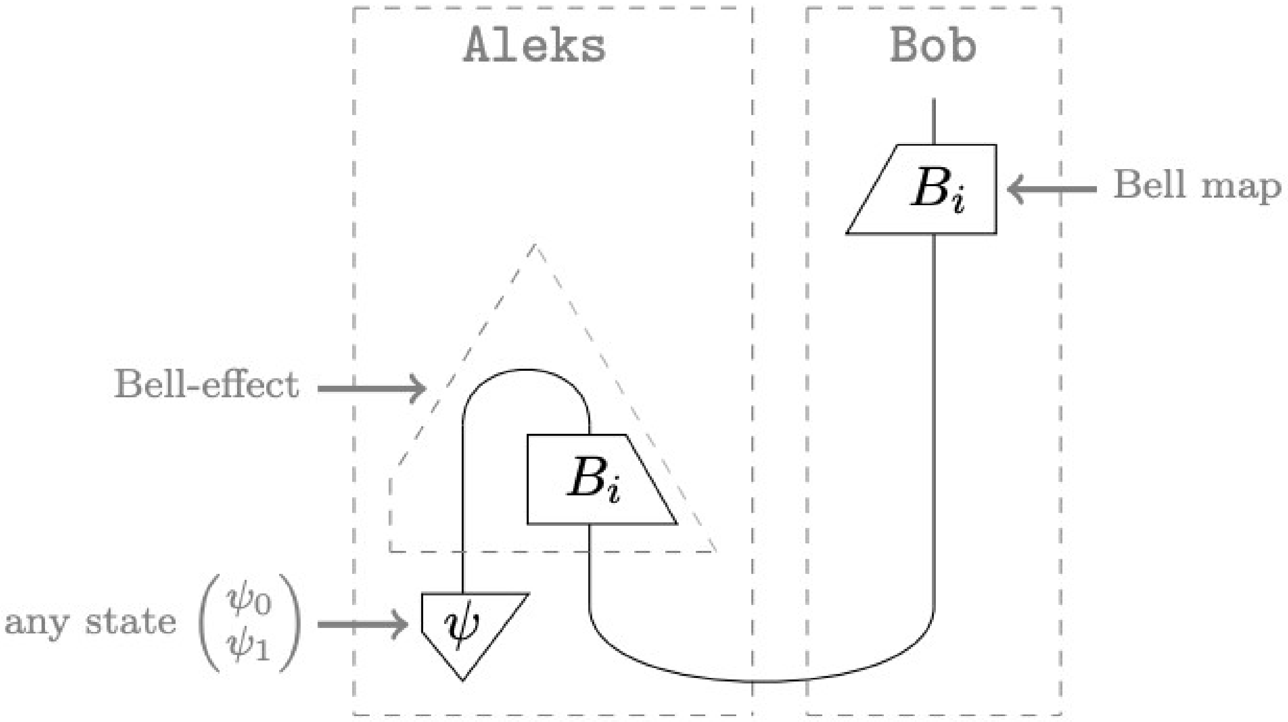
Aleks
Bob
B_i
Bell map
Bell-effect
B_i
any state $\begin{pmatrix}\psi_0\\ \psi_1\end{pmatrix}$
ψ

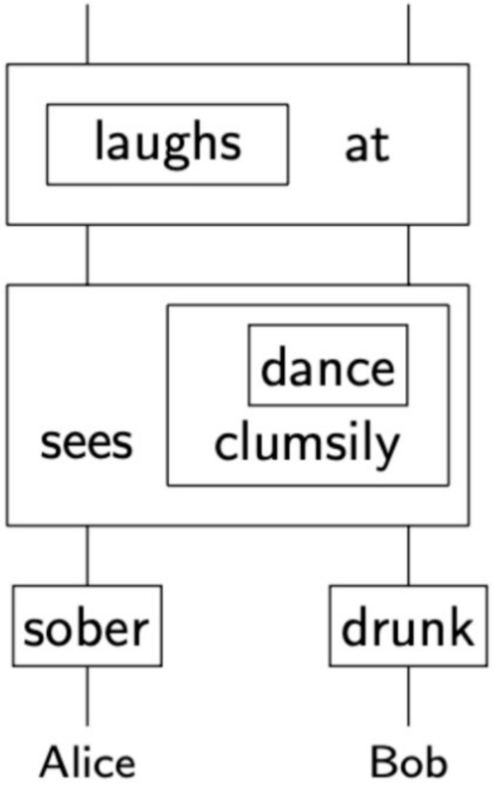
laughs
at
dance
sees
clumsily
sober
drunk
Alice
Bob

Contra utter lack of “interpretability” for AI

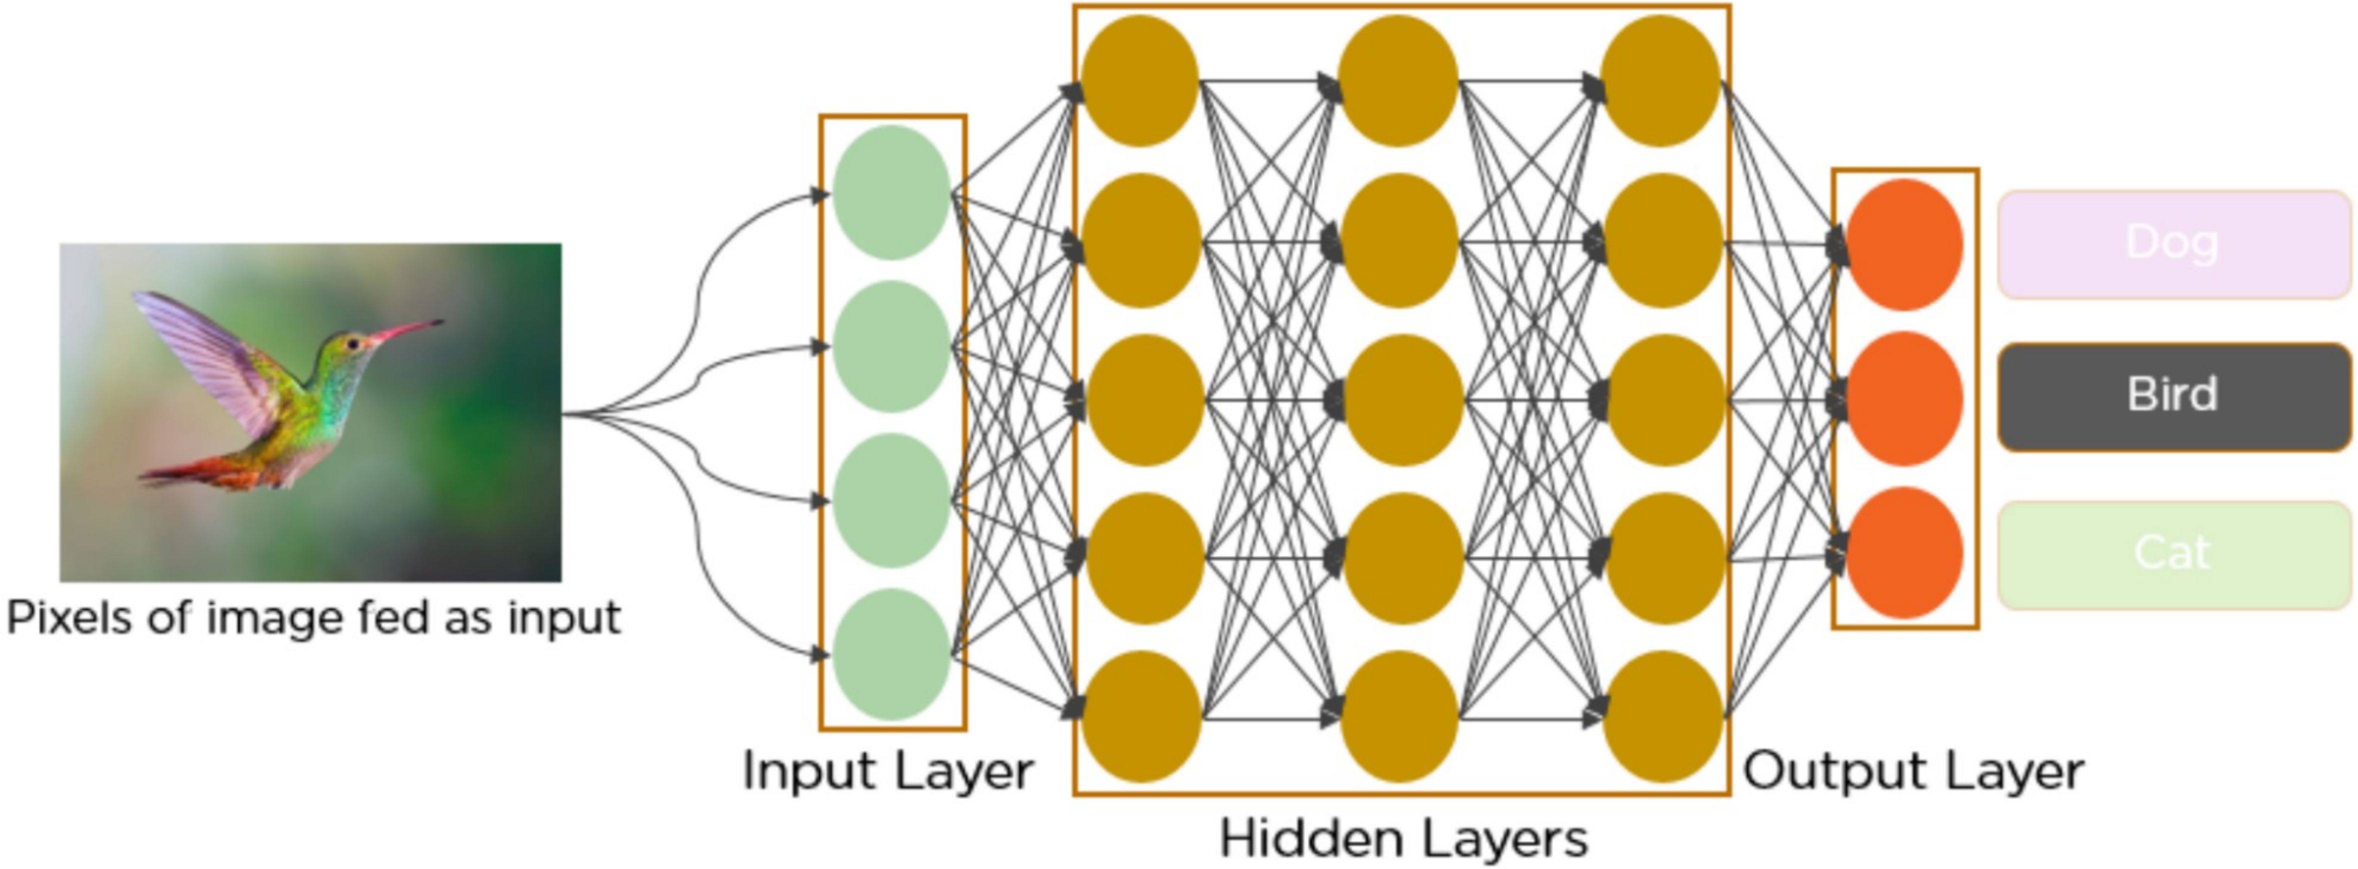

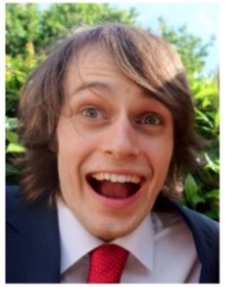



[Submitted on 25 Jun 2024]

Towards Compositional Interpretability for XAI

Sean Tull, Robin Lorenz, Stephen Clark, Ilyas Khan, Bob Coecke

Artificial intelligence (AI) is currently based largely on black-box machine learning models which lack interpretability. The field of eXplainable AI (XAI) strives to address this major concern, being critical in high-stakes areas such as the finance, legal and health sectors.

We present an approach to defining AI models and their interpretability based on category theory. For this we employ the notion of a compositional model, which sees a model in terms of formal string diagrams which capture its abstract structure together with its concrete implementation. This comprehensive view incorporates deterministic, probabilistic and quantum models. We compare a wide range of AI models as compositional models, including linear and rule-based models, (recurrent) neural networks, transformers, VAEs, and causal and DisCoCirc models.

Next we give a definition of interpretation of a model in terms of its compositional structure, demonstrating how to analyse the interpretability of a model, and using this to clarify common themes in XAI. We find that what makes the standard 'intrinsically interpretable' models so transparent is brought out most clearly diagrammatically. This leads us to the more general notion of compositionally-interpretable (CI) models, which additionally include, for instance, causal, conceptual space, and DisCoCirc models.

We next demonstrate the explainability benefits of CI models. Firstly, their compositional structure may allow the computation of other quantities of interest, and may facilitate inference from the model to the modelled phenomenon by matching its structure. Secondly, they allow for diagrammatic explanations for their behaviour, based on influence constraints, diagram surgery and rewrite explanations. Finally, we discuss many future directions for the approach, raising the question of how to learn such meaningfully structured models in practice.

Compositional Interpretability:

**Compositional structure of the model
reflects
compositional structure of the phenomenon.**

Intrinsic Quantum-Shaped:

If a phenomenon results from quantum systems.

Effective Quantum-Shaped:

If for a phenomenon,
under well-established assumptions,
the quantum formalism is a necessary description.

Effective Quantum-Shaped:

If for a phenomenon,
under well-established assumptions,
the quantum formalism is a necessary description.

These assumptions may either be:

- foundational
- engineering practice

Under the assumptions:

- Compositional interpretability
- Machine learning practice

Under the assumptions:

- Compositional interpretability
- Machine learning practice

Text is effectively quantum-shaped.

Under the assumptions:

- Compositional interpretability
- Machine learning practice

Text is effectively quantum-shaped.

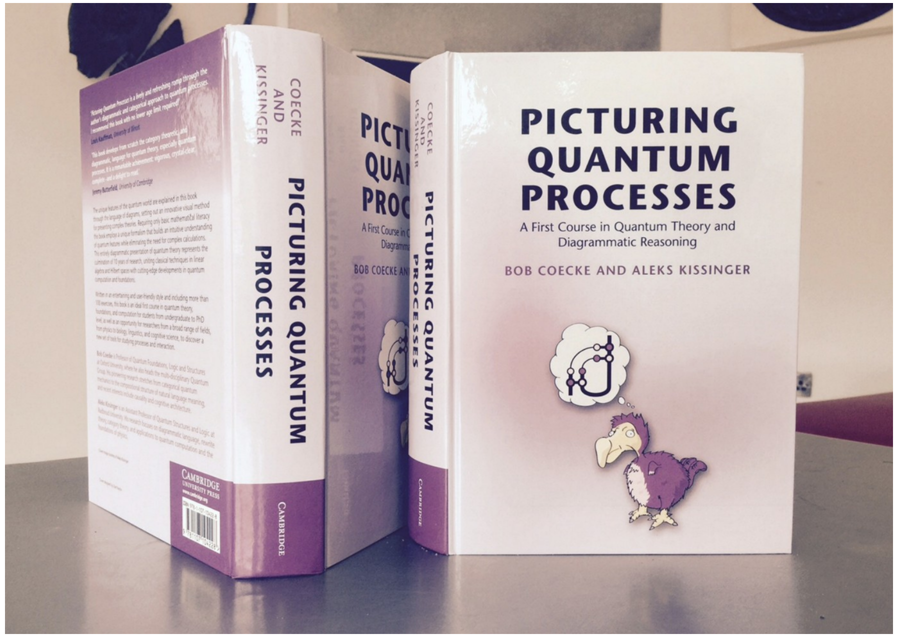

The ZX-calculus

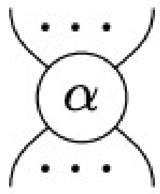

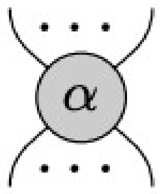

The ZX-calculus

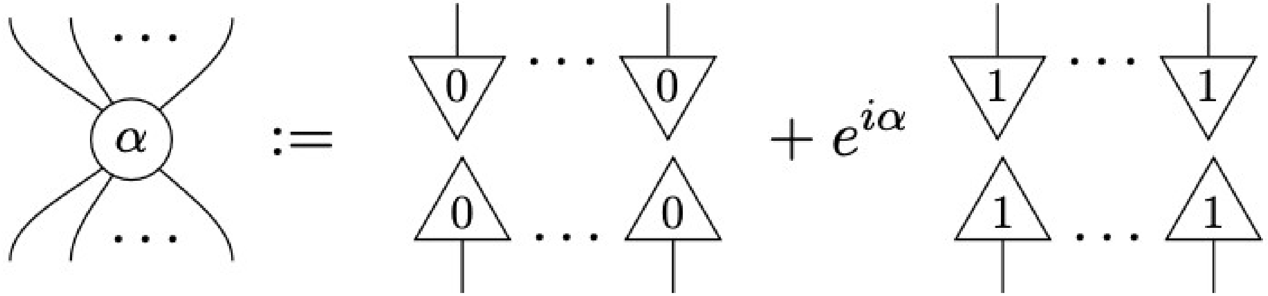

$$= |0\cdots0\rangle\langle0\cdots0| + e^{i\alpha}\,|1\cdots1\rangle\langle1\cdots1|$$

.

The ZX-calculus

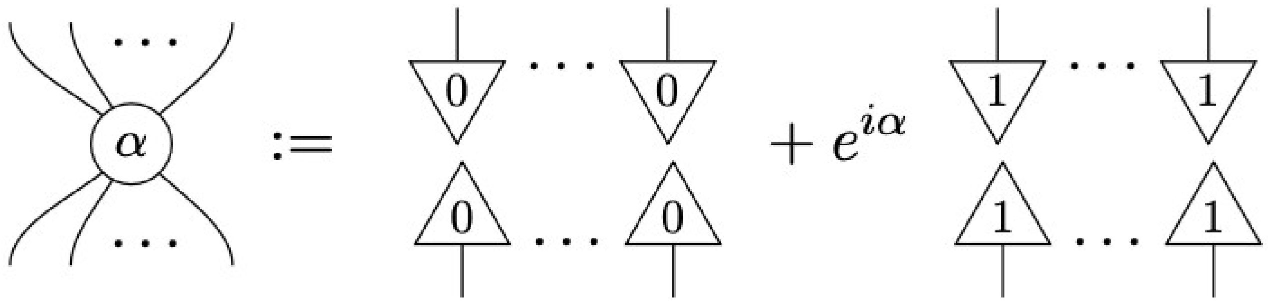

$$= |0\cdots0\rangle\langle0\cdots0| + e^{i\alpha}|1\cdots1\rangle\langle1\cdots1|$$

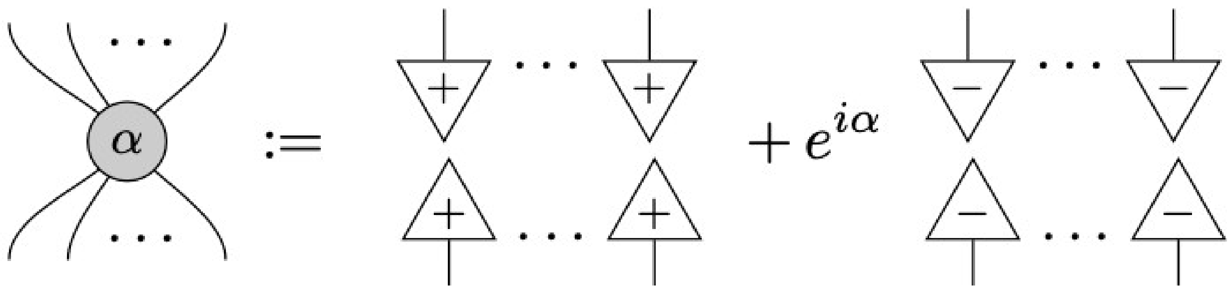

$$= |+\cdots+\rangle\langle+\cdots+| + e^{i\alpha}|-\cdots-\rangle\langle-\cdots-|$$

.

The ZX-calculus

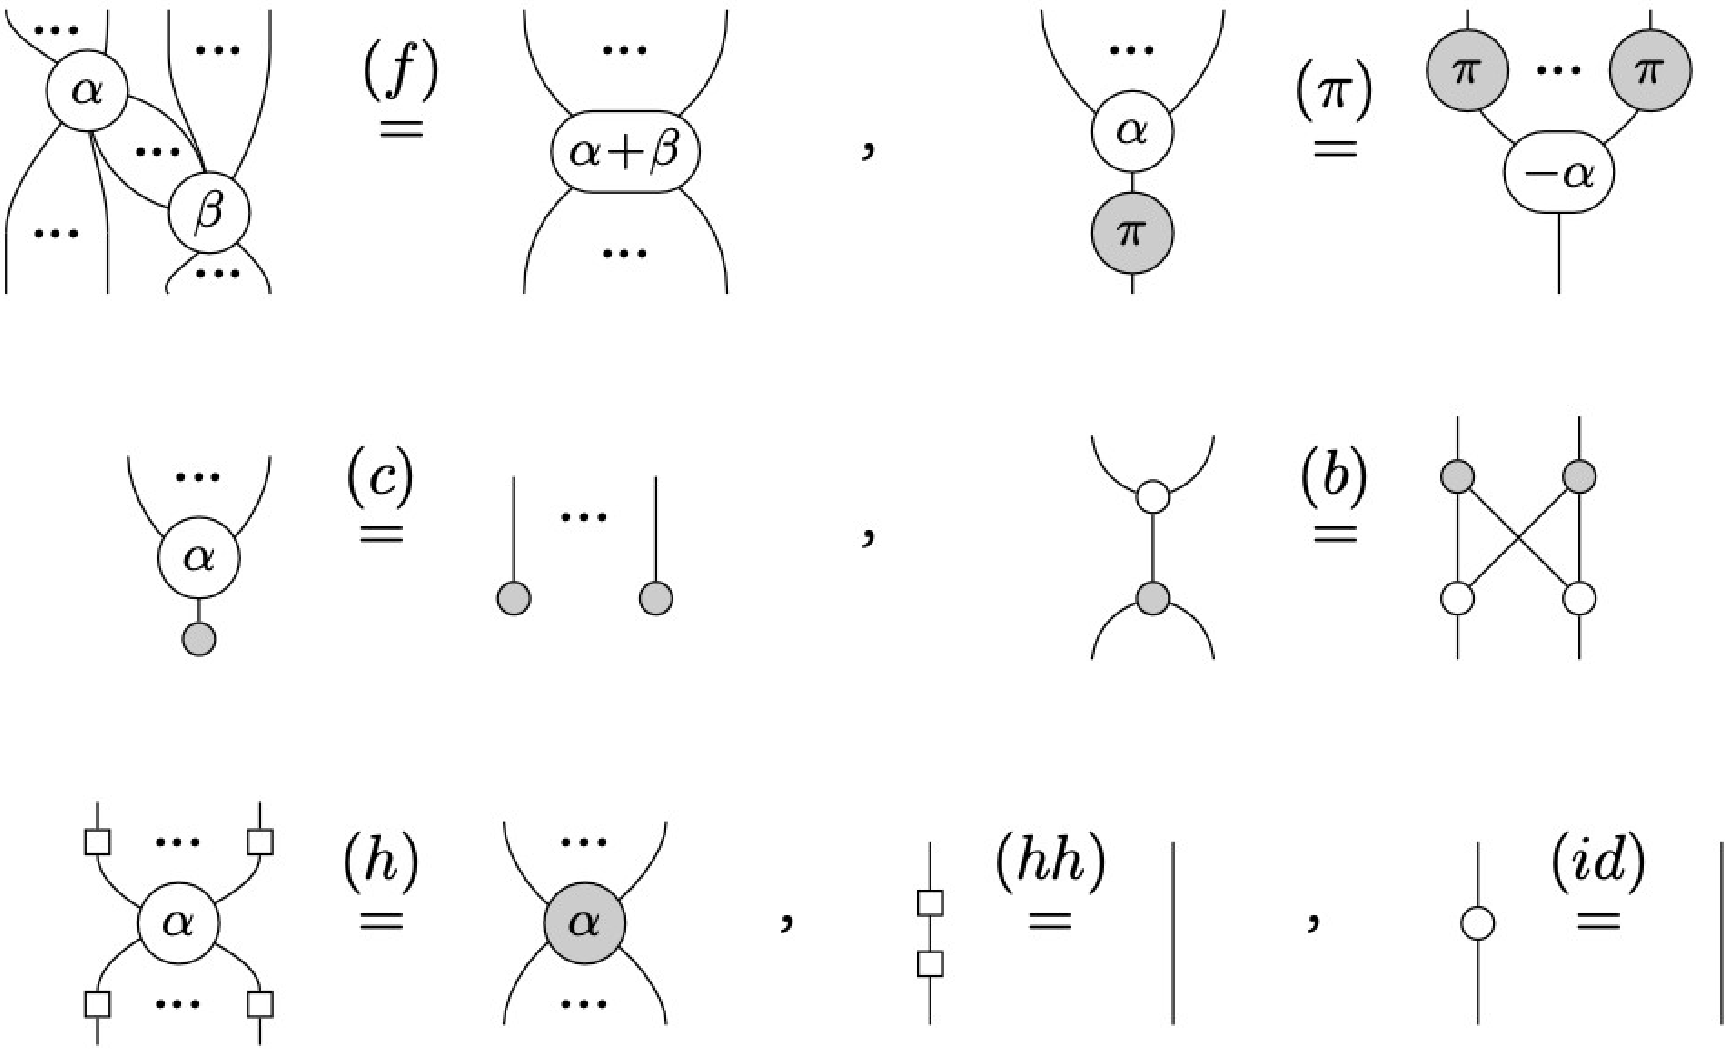

DisCoCat on Quantum Computers

The grammar/meaning-blend is exponentially expensive!

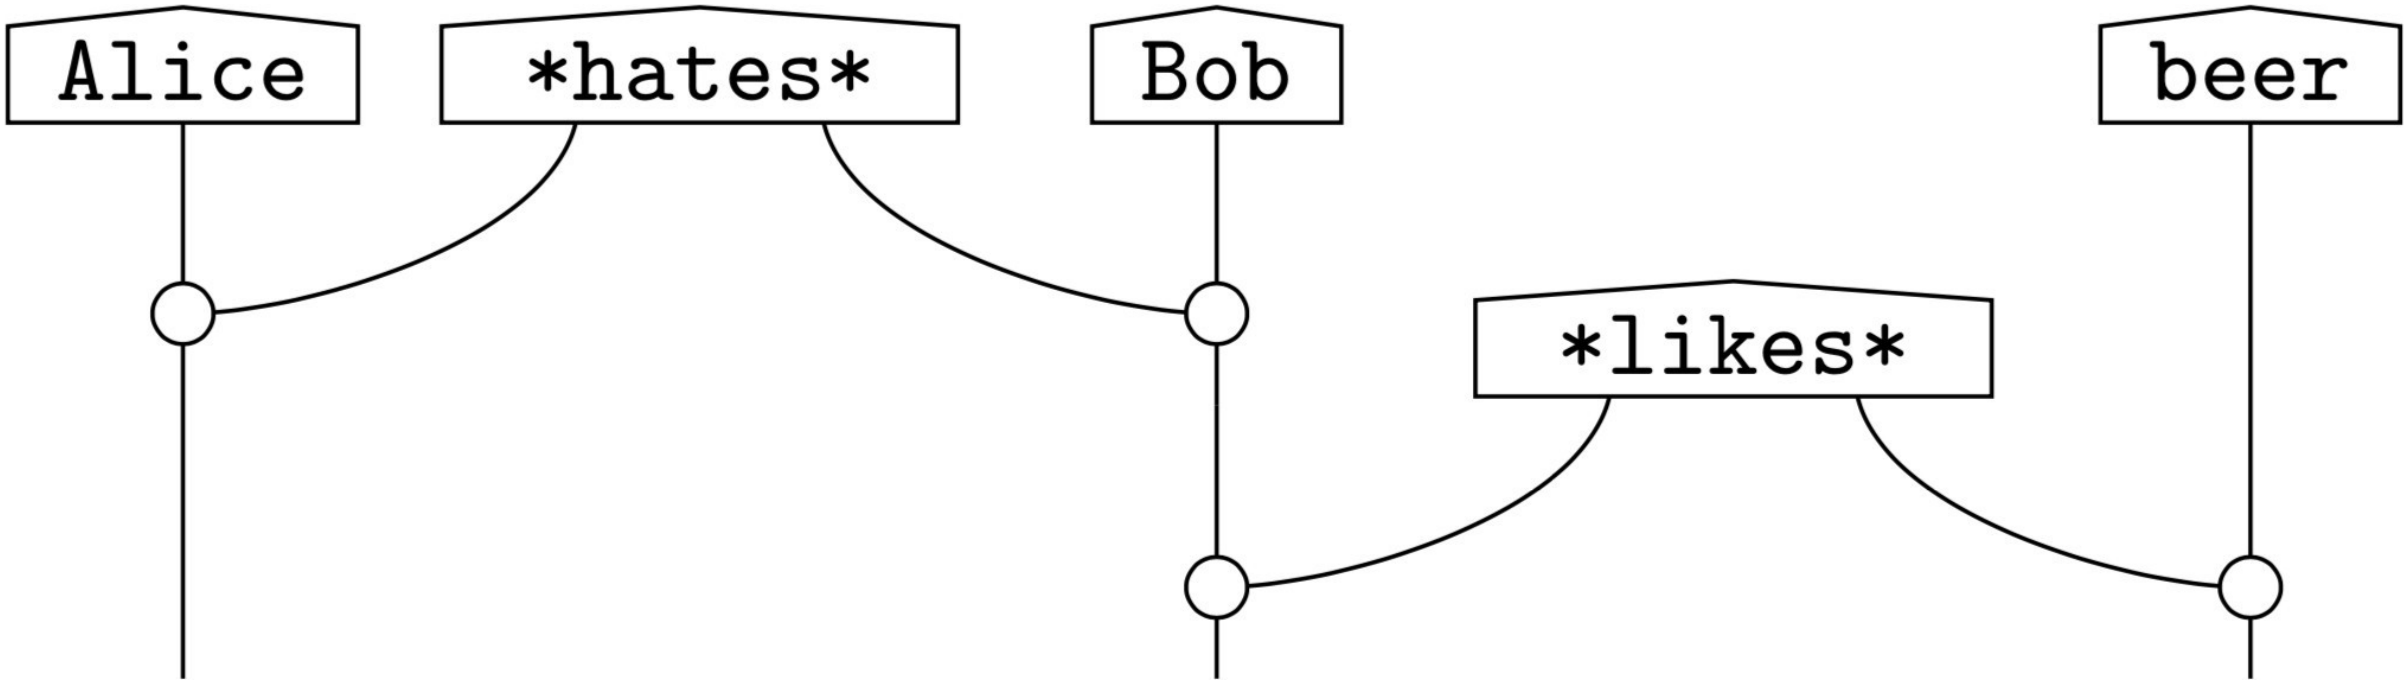

Do it on a ‘hypothetical’ quantum computer!

arXiv.org > cs > arXiv:1608.01406

Search...

Help | Advanc

Computer Science > Computation and Language

[Submitted on 4 Aug 2016]

Quantum Algorithms for Compositional Natural Language Processing

William Zeng (Rigetti Computing), Bob Coecke (Univesity of Oxford)

We propose a new application of quantum computing to the field of natural language processing. Ongoing work in this field attempts to incorporate grammatical structure into algorithms that compute meaning. In (Coecke, Sadrzadeh and Clark, 2010), the authors introduce such a model (the CSC model) based on tensor product composition. While this algorithm has many advantages, its implementation is hampered by the large classical computational resources that it requires. In this work we show how computational shortcomings of the CSC approach could be resolved using quantum computation (possibly in addition to existing techniques for dimension reduction). We address the value of quantum RAM (Giovannetti,2008) for this model and extend an algorithm from Wiebe, Braun and Lloyd (2012) into a quantum algorithm to categorize sentences in CSC. Our new algorithm demonstrates a quadratic speedup over classical methods under certain conditions.

Comments: In Proceedings SLPCS 2016, arXiv:1608.01018

Questions

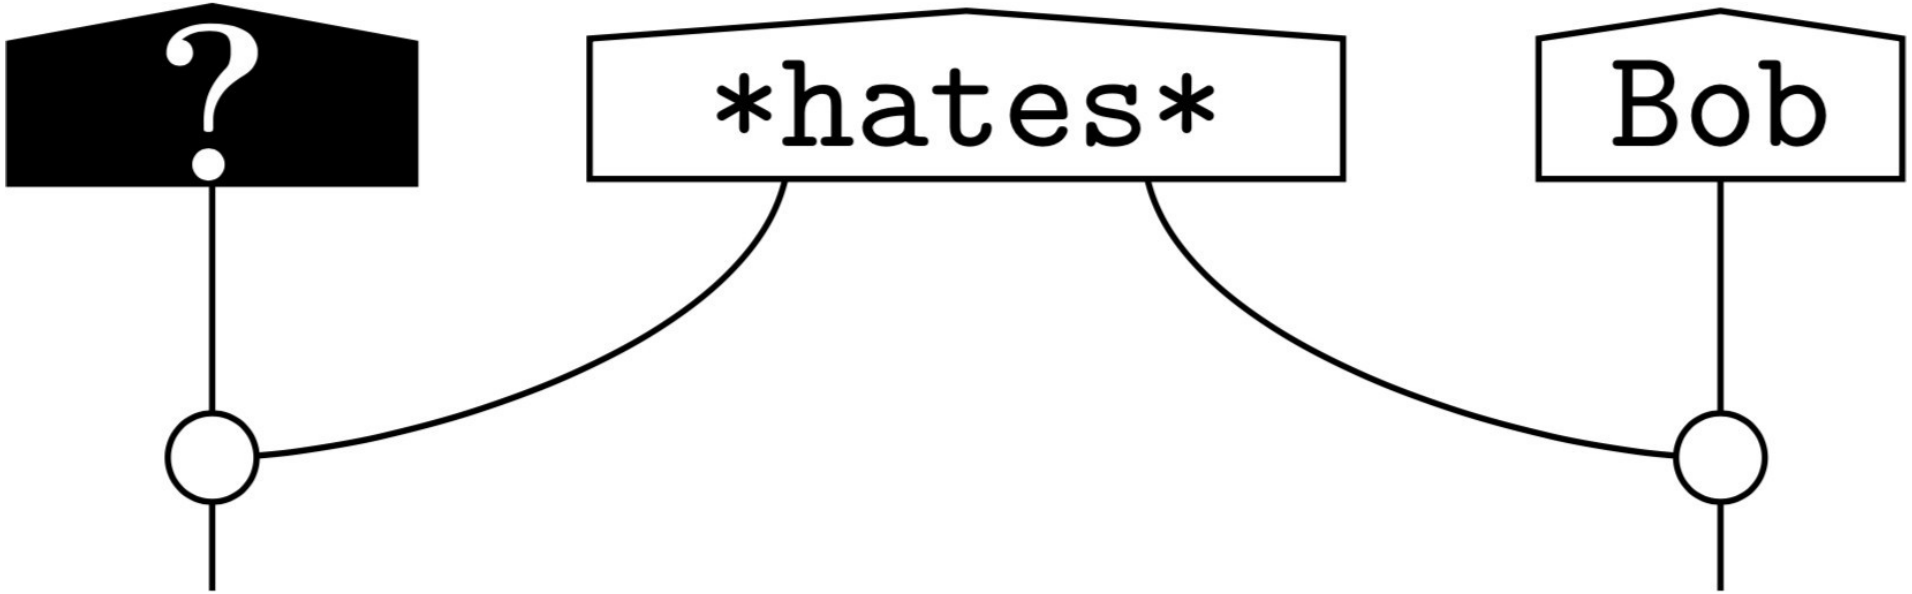

Answers

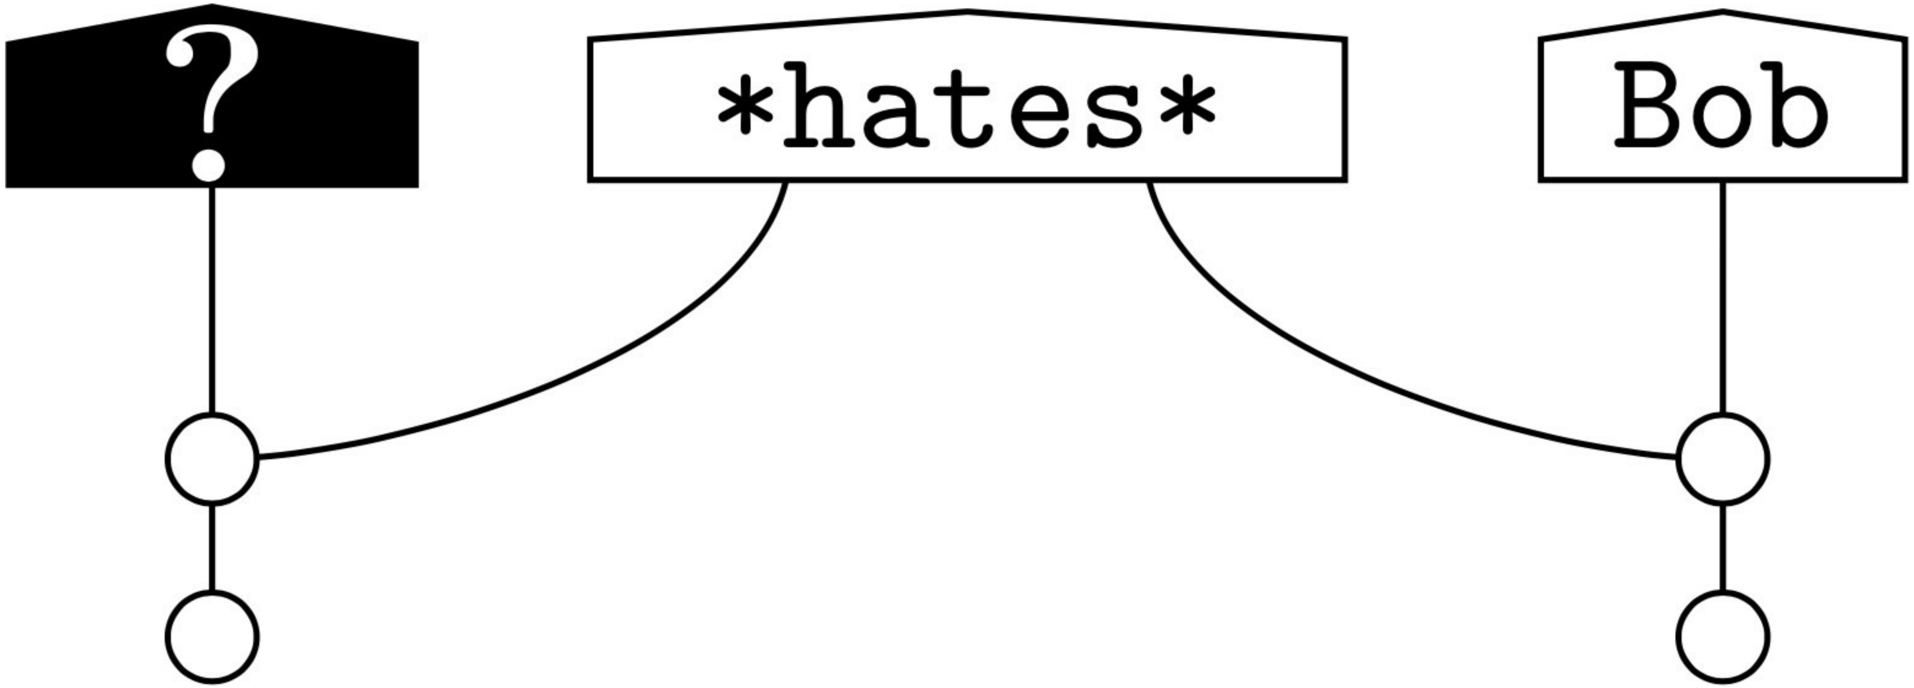

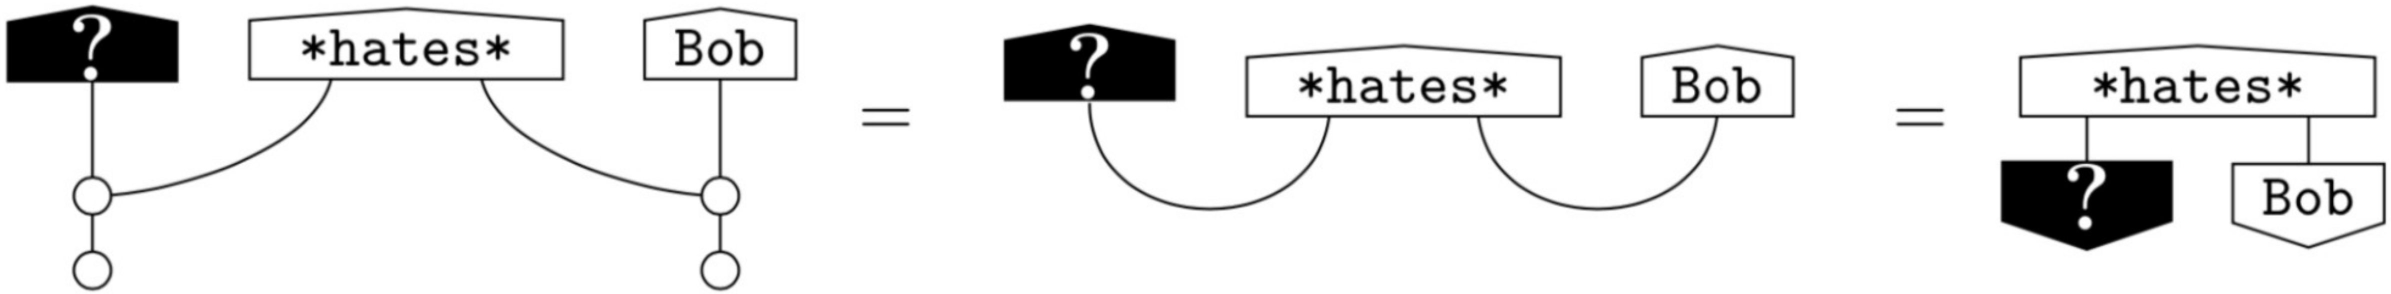
?
hates
Bob
=
?
hates
Bob
=
hates
?
Bob

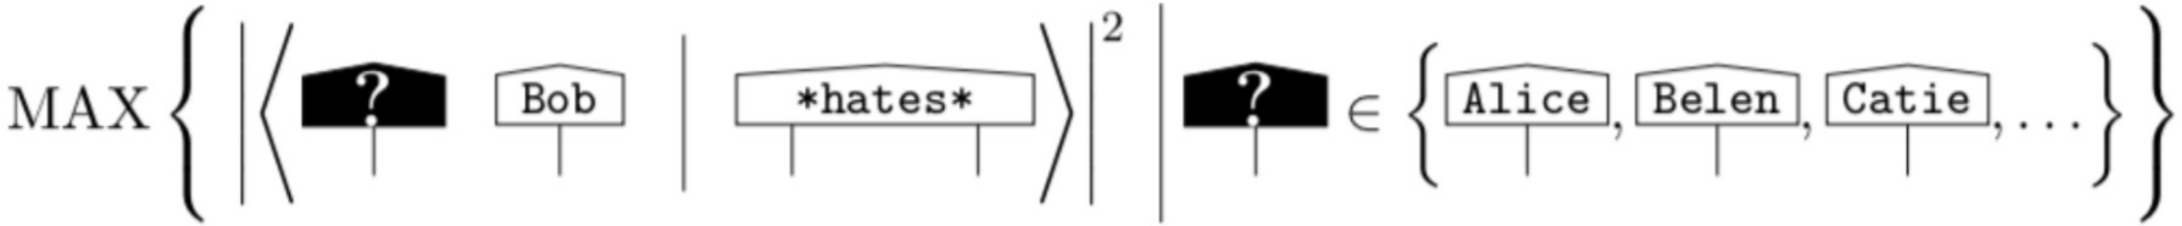
MAX
?
Bob
hates
2
?
∈
Alice
Belen
Catie
, . . .

Do it on existing quantum computers!




arXiv.org > quant-ph > arXiv:2012.03755

Search...

Help | Adv

Quantum Physics

[Submitted on 7 Dec 2020]

Foundations for Near-Term Quantum Natural Language Processing

Bob Coecke, Giovanni de Felice, Konstantinos Meichanetzidis, Alexis Toumi

We provide conceptual and mathematical foundations for near-term quantum natural language processing (QNLP), and do so in quantum computer scientist friendly terms. We opted for an expository presentation style, and provide references for supporting empirical evidence and formal statements concerning mathematical generality.
We recall how the quantum model for natural language that we employ canonically combines linguistic meanings with rich linguistic structure, most notably grammar. In particular, the fact that it takes a quantum-like model to combine meaning and structure, establishes QNLP as quantum-native, on par with simulation of quantum systems. Moreover, the now leading Noisy Intermediate-Scale Quantum (NISQ) paradigm for encoding classical data on quantum hardware, variational quantum circuits, makes NISQ exceptionally QNLP-friendly: linguistic structure can be encoded as a free lunch, in contrast to the apparently exponentially expensive classical encoding of grammar.
Quantum speed-up for QNLP tasks has already been established in previous work with Will Zeng. Here we provide a broader range of tasks which all enjoy the same advantage.
Diagrammatic reasoning is at the heart of QNLP. Firstly, the quantum model interprets language as quantum processes via the diagrammatic formalism of categorical quantum mechanics. Secondly, these diagrams are via ZX-calculus translated into quantum circuits. Parameterisations of meanings then become the circuit variables to be learned.
Our encoding of linguistic structure within quantum circuits also embodies a novel approach for establishing word-meanings that goes beyond the current standards in mainstream AI, by placing linguistic structure at the heart of Wittgenstein's meaning-is-context.

Comments: 43 pages, lots of pictures

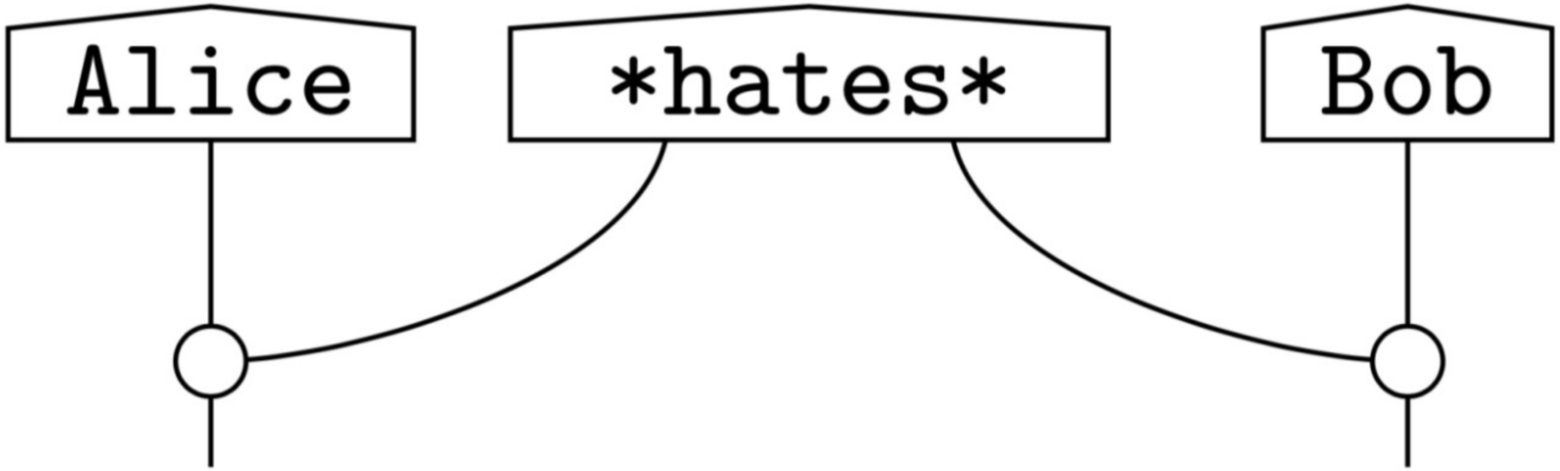
Alice
hates
Bob

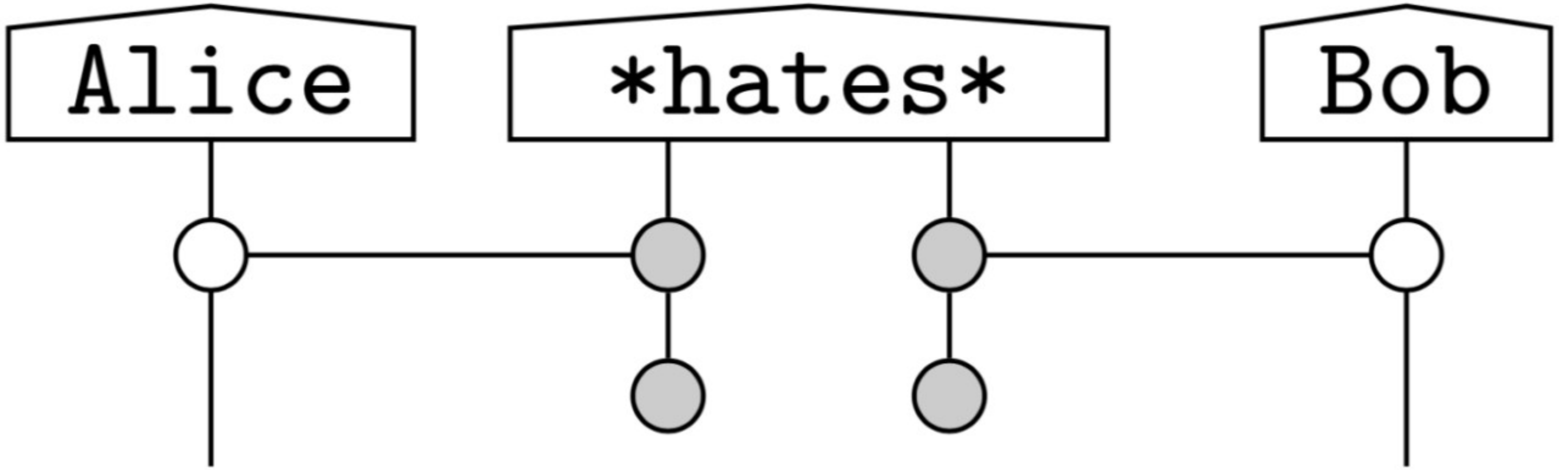
Alice
hates
Bob

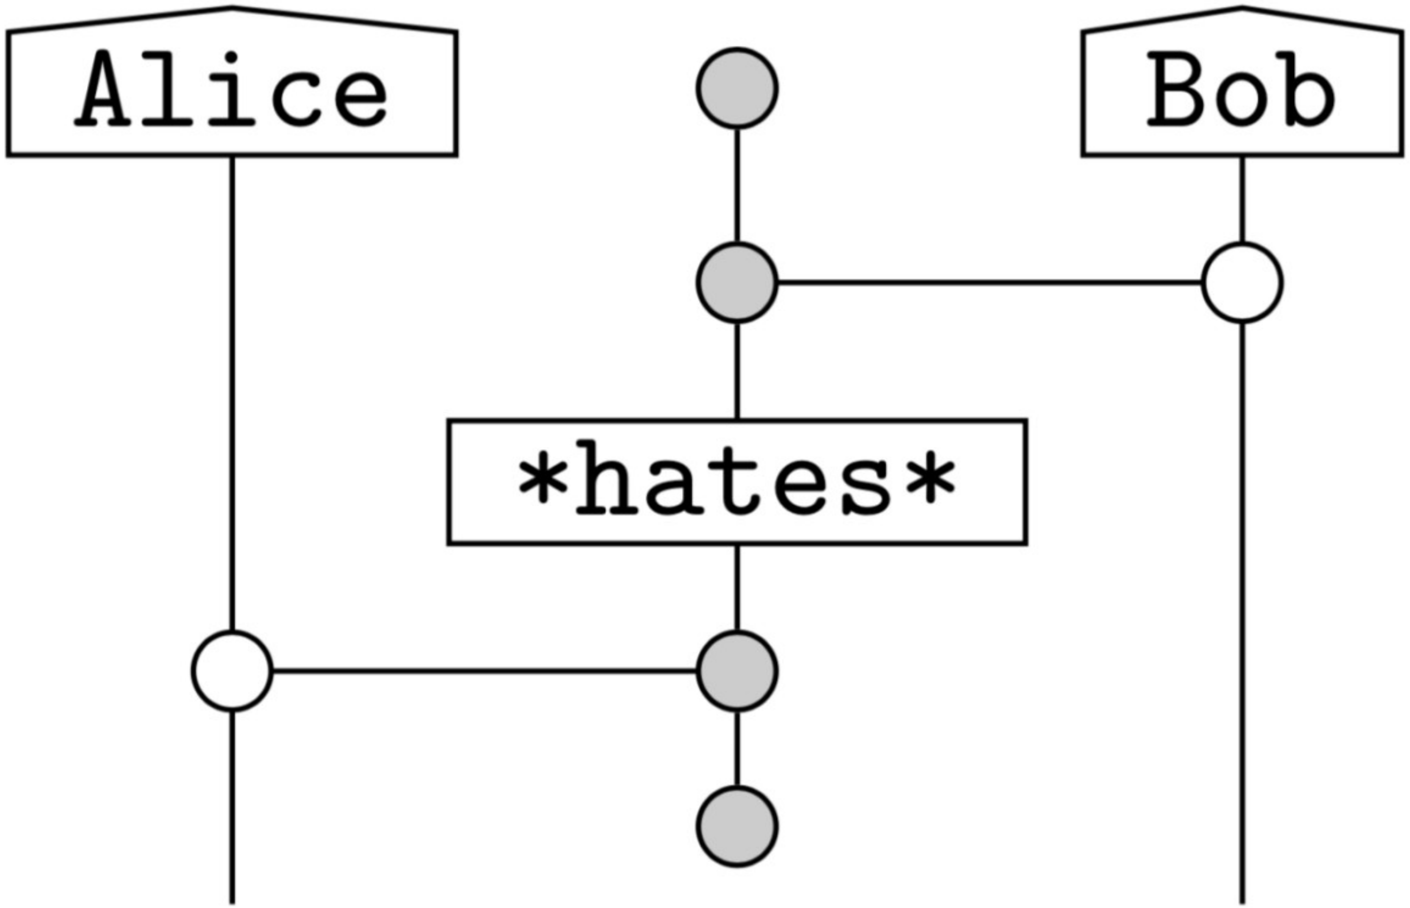
Alice
Bob
hates

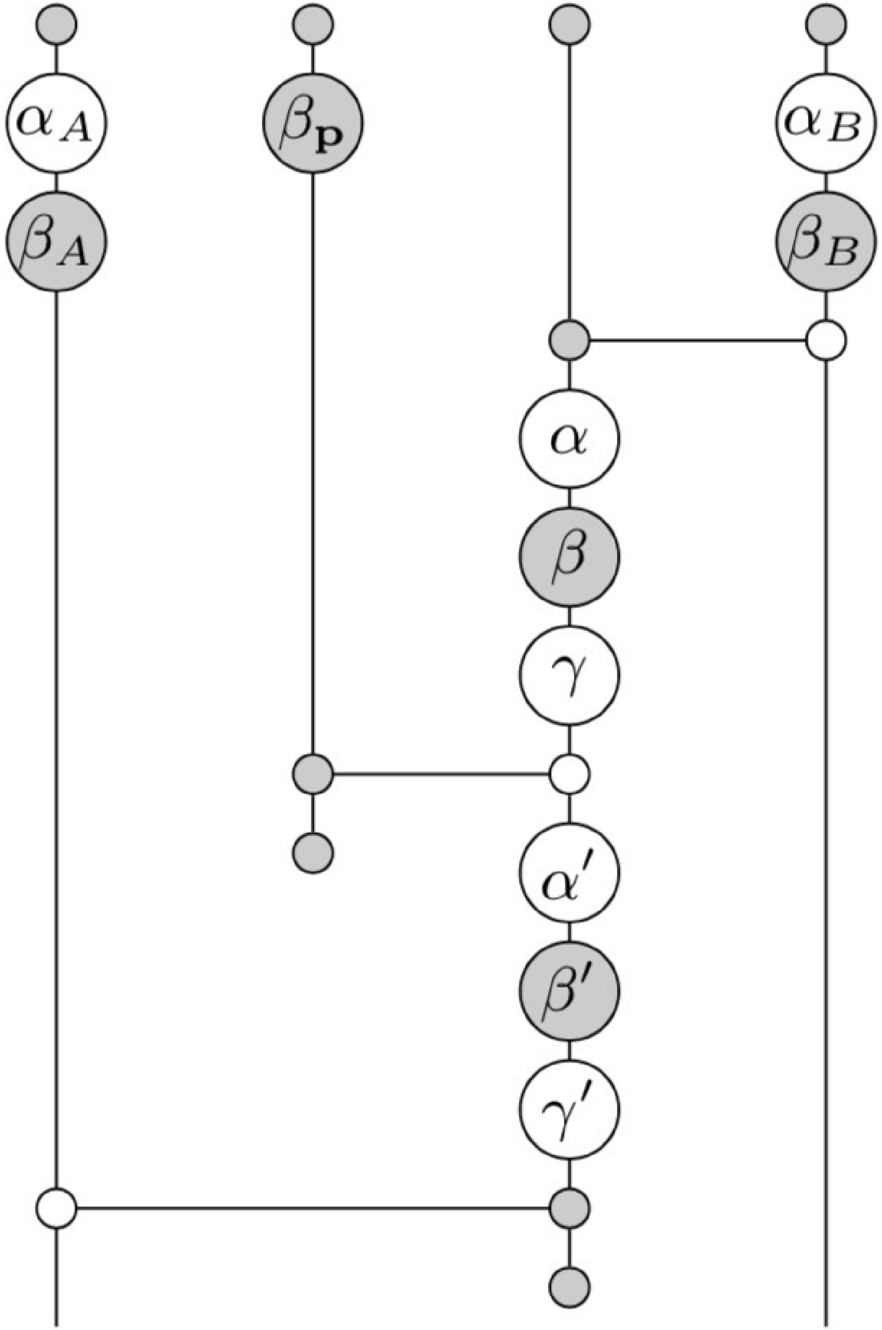
α_A
β_A
$\beta_{\mathbf{p}}$
α_B
β_B
α
β
γ
α'
β'
γ'

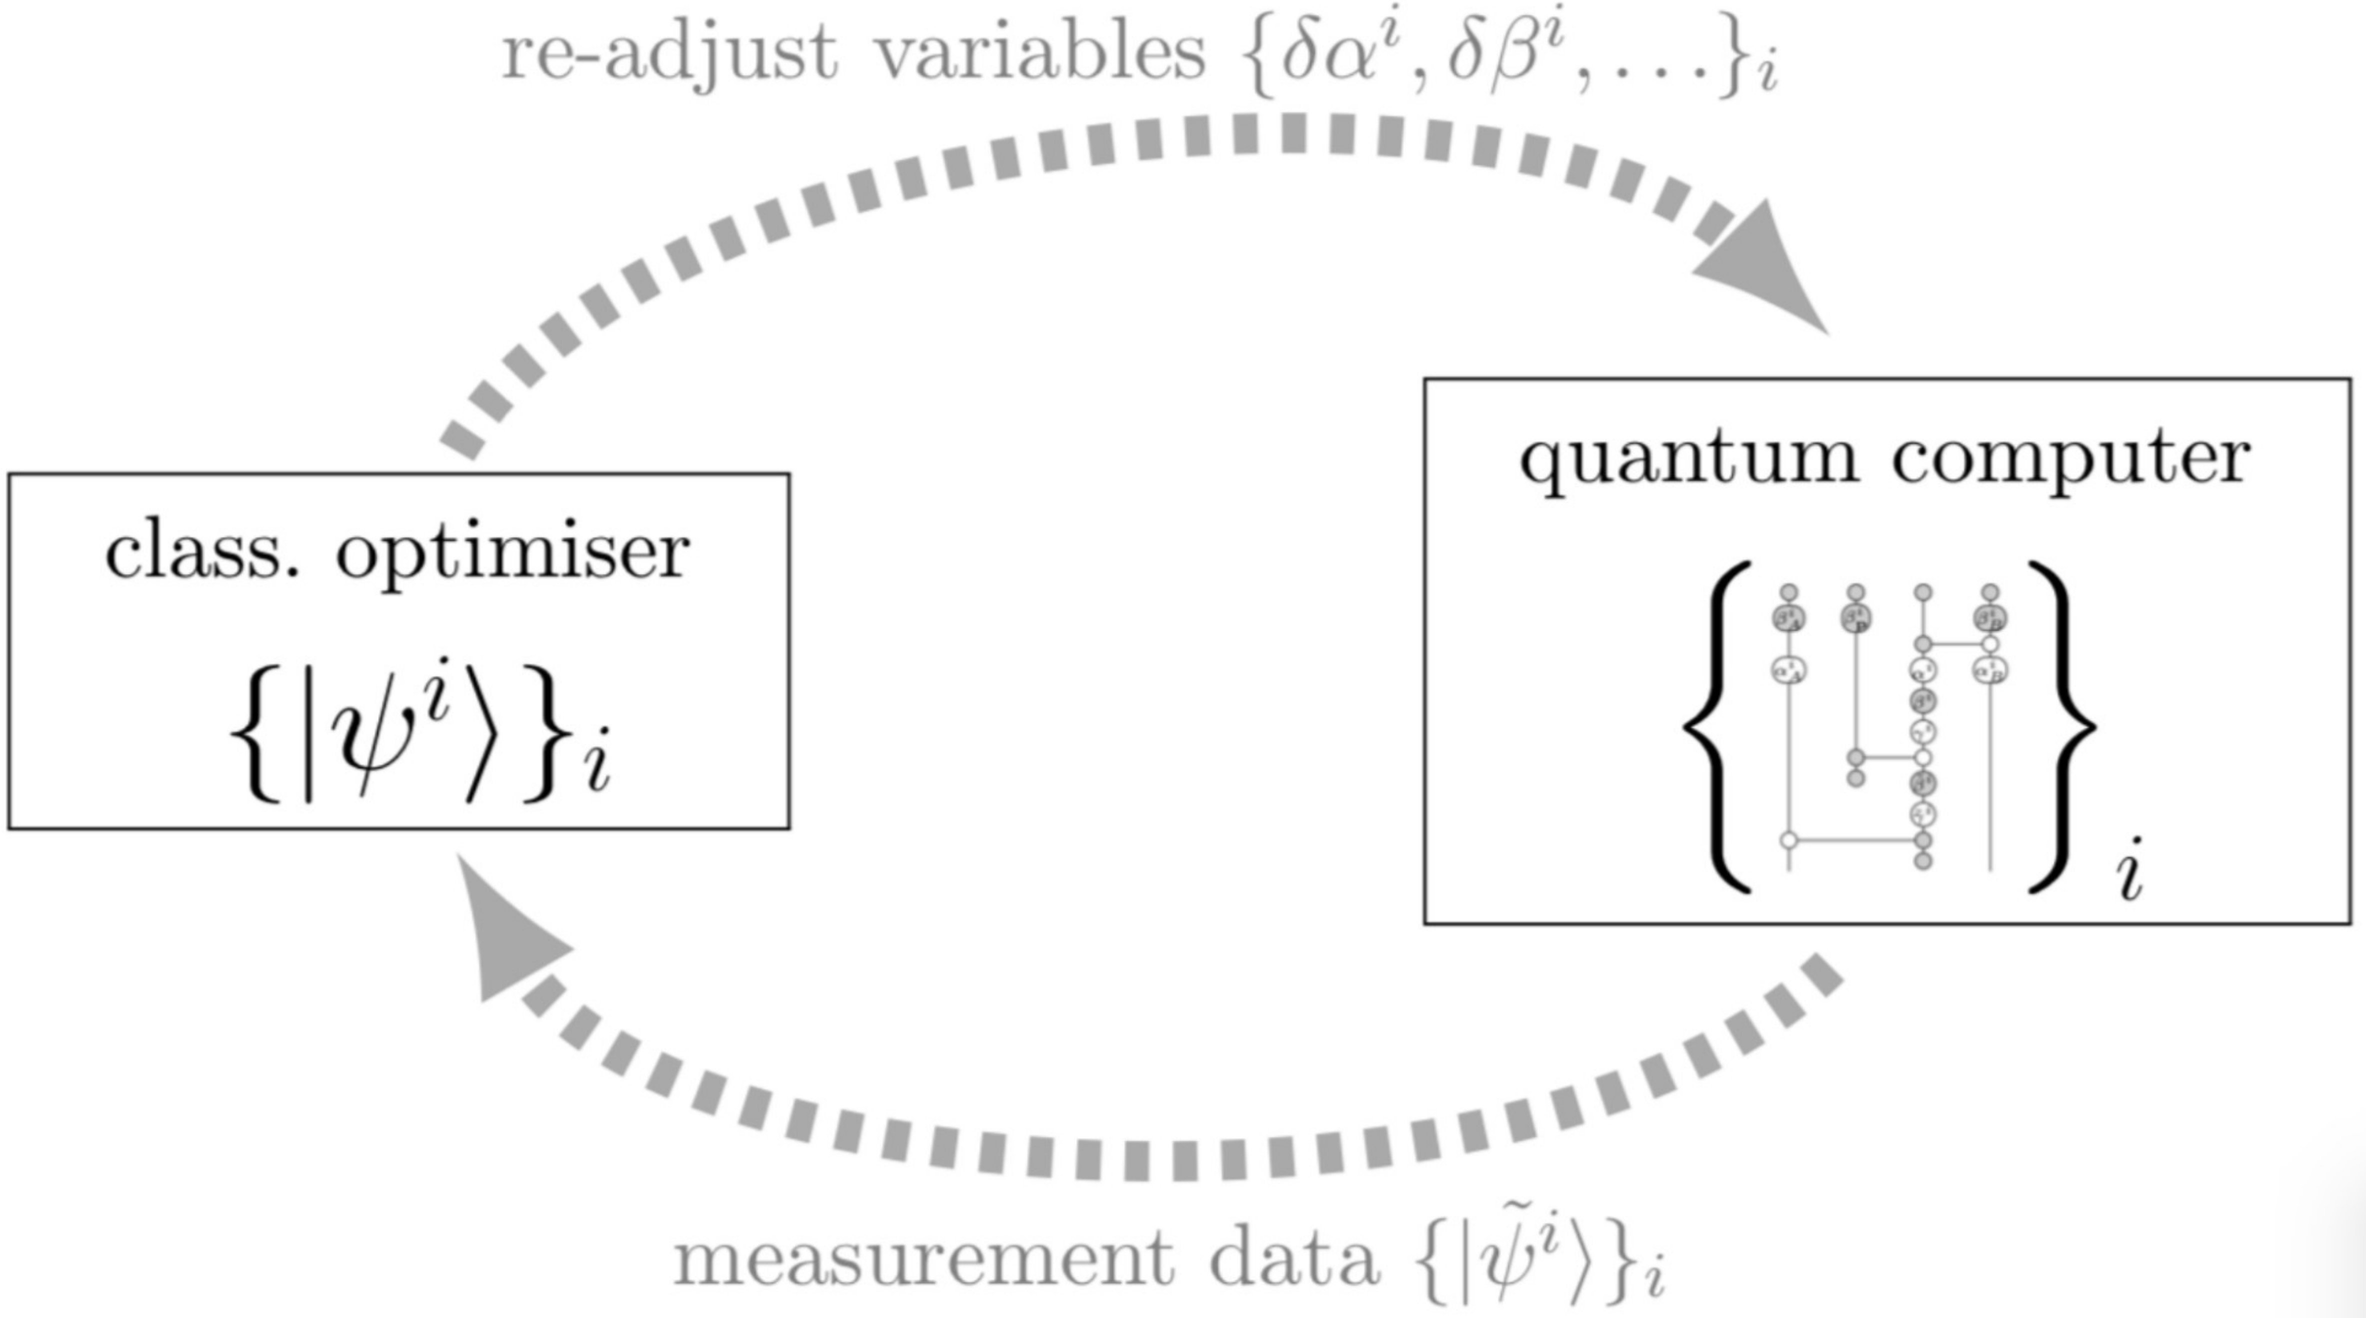

re-adjust variables $\{\delta\alpha^i, \delta\beta^i, \ldots\}_i$
class. optimiser
$\{|\psi^i\rangle\}_i$
quantum computer
measurement data $\{|\tilde{\psi}^i\rangle\}_i$

Just do it!




arXiv.org > quant-ph > arXiv:2012.03756

Search...

Help | Adva

Quantum Physics

[Submitted on 7 Dec 2020]

Grammar-Aware Question-Answering on Quantum Computers

Konstantinos Meichanetzidis, Alexis Toumi, Giovanni de Felice, Bob Coecke

Natural language processing (NLP) is at the forefront of great advances in contemporary AI, and it is arguably one of the most challenging areas of the field. At the same time, with the steady growth of quantum hardware and notable improvements towards implementations of quantum algorithms, we are approaching an era when quantum computers perform tasks that cannot be done on classical computers with a reasonable amount of resources. This provides a new range of opportunities for AI, and for NLP specifically. Earlier work has already demonstrated a potential quantum advantage for NLP in a number of manners: (i) algorithmic speedups for search-related or classification tasks, which are the most dominant tasks within NLP, (ii) exponentially large quantum state spaces allow for accommodating complex linguistic structures, (iii) novel models of meaning employing density matrices naturally model linguistic phenomena such as hyponymy and linguistic ambiguity, among others. In this work, we perform the first implementation of an NLP task on noisy intermediate-scale quantum (NISQ) hardware. Sentences are instantiated as parameterised quantum circuits. We encode word-meanings in quantum states and we explicitly account for grammatical structure, which even in mainstream NLP is not commonplace, by faithfully hard-wiring it as entangling operations. This makes our approach to quantum natural language processing (QNLP) particularly NISQ-friendly. Our novel QNLP model shows concrete promise for scalability as the quality of the quantum hardware improves in the near future.

Subjects: **Quantum Physics (quant-ph)**; Computation and Language (cs.CL)

Just do it BIGGER!

    

[Submitted on 25 Feb 2021]

QNLP in Practice: Running Compositional Models of Meaning on a Quantum Computer

Robin Lorenz, Anna Pearson, Konstantinos Meichanetzidis, Dimitri Kartsaklis, Bob Coecke

Quantum Natural Language Processing (QNLP) deals with the design and implementation of NLP models intended to be run on quantum hardware. In this paper, we present results on the first NLP experiments conducted on Noisy Intermediate-Scale Quantum (NISQ) computers for datasets of size >= 100 sentences. Exploiting the formal similarity of the compositional model of meaning by Coecke et al. (2010) with quantum theory, we create representations for sentences that have a natural mapping to quantum circuits. We use these representations to implement and successfully train two NLP models that solve simple sentence classification tasks on quantum hardware. We describe in detail the main principles, the process and challenges of these experiments, in a way accessible to NLP researchers, thus paving the way for practical Quantum Natural Language Processing.

Subjects: **Computation and Language (cs.CL)**; Artificial Intelligence (cs.AI); Machine Learning (cs.LG); Quantum Physics (quant-ph)

YOU just do it!

arXiv > cs > arXiv:2110.04236

Search...

Help | Advanced

Computer Science > Computation and Language

[Submitted on 8 Oct 2021]

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

We present lambeq, the first high-level Python library for Quantum Natural Language Processing (QNLP). The open-source 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 syntactic 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 describe the most important modules in detail, demonstrating the usage with illustrative examples. Further, we test the toolkit in practice by using it to perform a number of experiments on simple NLP tasks, implementing both classical and quantum pipelines.

Subjects: **Computation and Language (cs.CL)**; Artificial Intelligence (cs.AI); Quantum Physics (quant-ph)

LAMBEQ

Natural Language Processing on Quantum Computers

DisCoCirc on Quantum Computers

Alex meets Beau.
Beau marries Chris.
Alex kicks Beau.

Alex meets Beau.
Beau marries Chris.
Alex kicks Beau.

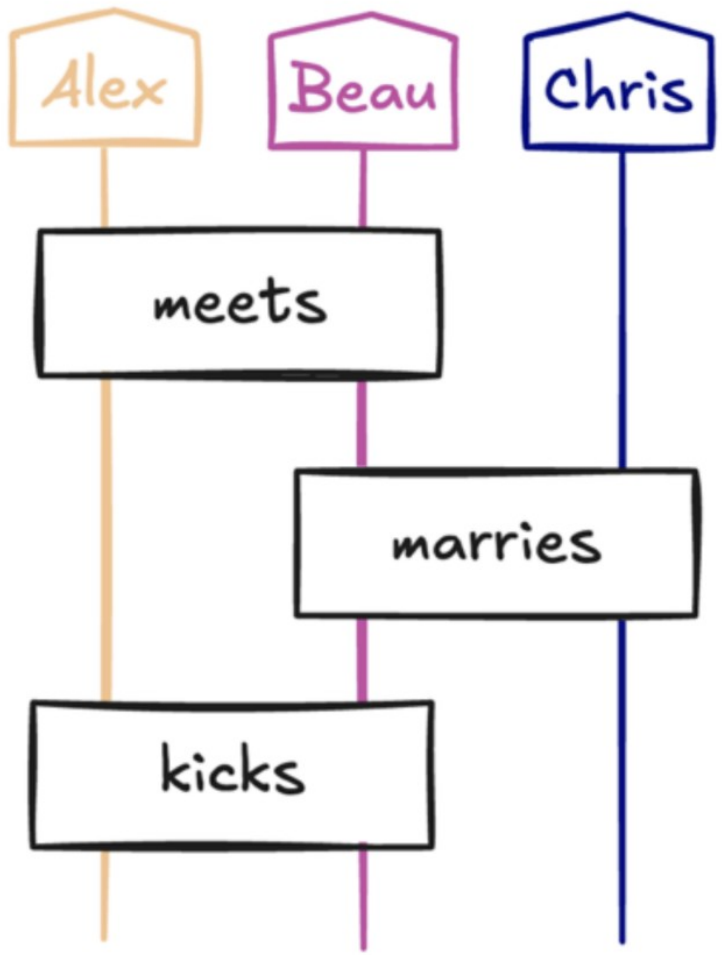

  

Quantum Physics

[Submitted on 12 Aug 2024]

Quantum Algorithms for Compositional Text Processing

Tuomas Laakkonen (Quantinuum), Konstantinos Meichanetzidis (Quantinuum), Bob Coecke (Quantinuum)

Quantum computing and AI have found a fruitful intersection in the field of natural language processing. We focus on the recently proposed DisCoCirc framework for natural language, and propose a quantum adaptation, QDisCoCirc. This is motivated by a compositional approach to rendering AI interpretable: the behavior of the whole can be understood in terms of the behavior of parts, and the way they are put together. For the model-native primitive operation of text similarity, we derive quantum algorithms for fault-tolerant quantum computers to solve the task of question-answering within QDisCoCirc, and show that this is BQP-hard; note that we do not consider the complexity of question-answering in other natural language processing models. Assuming widely-held conjectures, implementing the proposed model classically would require super-polynomial resources. Therefore, it could provide a meaningful demonstration of the power of practical quantum processors. The model construction builds on previous work in compositional quantum natural language processing. Word embeddings are encoded as parameterized quantum circuits, and compositionality here means that the quantum circuits compose according to the linguistic structure of the text. We outline a method for evaluating the model on near-term quantum processors, and elsewhere we report on a recent implementation of this on quantum hardware. In addition, we adapt a quantum algorithm for the closest vector problem to obtain a Grover-like speedup in the fault-tolerant regime for our model. This provides an unconditional quadratic speedup over any classical algorithm in certain circumstances, which we will verify empirically in future work.

Quantum Algorithms for Text Processes

Tuomas Laakkonen, Konstantinos Meichanetzidis, Bob Coecke

Definition 1. *The problem* QUESTION-ANSWERING *is defined as follows: given a set of word embeddings V, a context text T, and a set of k question texts $\{Q_i\}$, determine any j such that*

$$\left| \operatorname{tr}\left(\rho_T(\rho_{Q_j} \otimes I)\right) - \max_i \operatorname{tr}\left(\rho_T(\rho_{Q_i} \otimes I)\right) \right| < \varepsilon$$

where $\rho_T = U_T|0\rangle\langle 0|U_T^\dagger$, $\rho_{Q_i} = U_{Q_i}|0\rangle\langle 0|U_{Q_i}^\dagger$, and U_T, U_{Q_i} are the QDisCoCirc text circuits generated from T and Q_i respectively over V.

Quantum Algorithms for Text Processes

Tuomas Laakkonen, Konstantinos Meichanetzidis, Bob Coecke

Definition 1. *The problem* QUESTION-ANSWERING *is defined as follows: given a set of word embeddings V, a context text T, and a set of k question texts $\{Q_i\}$, determine any j such that*

$$\left|\operatorname{tr}\left(\rho_T(\rho_{Q_j}\otimes I)\right)-\max_i \operatorname{tr}\left(\rho_T(\rho_{Q_i}\otimes I)\right)\right|<\varepsilon$$

where $\rho_T = U_T|0\rangle\langle 0|U_T^\dagger$, $\rho_{Q_i} = U_{Q_i}|0\rangle\langle 0|U_{Q_i}^\dagger$, and U_T, U_{Q_i} are the QDisCoCirc text circuits generated from T and Q_i respectively over V.

Theorem 4. *Suppose that a set of word embeddings V satisfies the following:*

1. *The operations of V use one qubit for each input wire,*
2. *V contains arbitrarily many proper nouns,*
3. *V contains at least two adjectives that generate a dense subset of $SU(2)$,*
4. *V contains at least one transitive verb that is entangling*

then for any fixed $\varepsilon < \frac{1}{7}$, QUESTION-ANSWERING *is BQP-hard.*


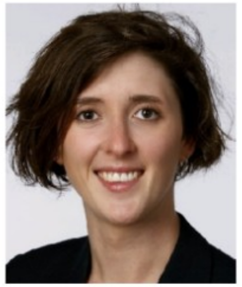





arXiv > quant-ph > arXiv:2409.08777

Search

Help

Quantum Physics

[Submitted on 13 Sep 2024]

Scalable and interpretable quantum natural language processing: an implementation on trapped ions

Tiffany Duneau, Saskia Bruhn, Gabriel Matos, Tuomas Laakkonen, Katerina Saiti, Anna Pearson, Konstantinos Meichanetzidis, Bob Coecke

We present the first implementation of text-level quantum natural language processing, a field where quantum computing and AI have found a fruitful intersection. We focus on the QDisCoCirc model, which is underpinned by a compositional approach to rendering AI interpretable: the behaviour of the whole can be understood in terms of the behaviour of parts, and the way they are put together. Interpretability is crucial for understanding the unwanted behaviours of AI. By leveraging the compositional structure in the model's architecture, we introduce a novel setup which enables 'compositional generalisation': we classically train components which are then composed to generate larger test instances, the evaluation of which asymptotically requires a quantum computer. Another key advantage of our approach is that it bypasses the trainability challenges arising in quantum machine learning. The main task that we consider is the model-native task of question-answering, and we handcraft toy scale data that serves as a proving ground. We demonstrate an experiment on Quantinuum's H1-1 trapped-ion quantum processor, which constitutes the first proof of concept implementation of scalable compositional QNLP. We also provide resource estimates for classically simulating the model. The compositional structure allows us to inspect and interpret the word embeddings the model learns for each word, as well as the way in which they interact. This improves our understanding of how it tackles the question-answering task. As an initial comparison with classical baselines, we considered transformer and LSTM models, as well as GPT-4, none of which succeeded at compositional generalisation.

Alice walks north.
Bob walks south.
Bob follows Alice.

⋮

Is Bob going in the same direction as Alice?

QUANTINUUM
H1-1

YES!

Compositional Model

DisCoCirc

arXiv: 2301.10595

DisCoCirc

Alice walks north.
Bob walks south.
Bob follows Alice.
⋮
Is Bob going in the same direction as Alice?

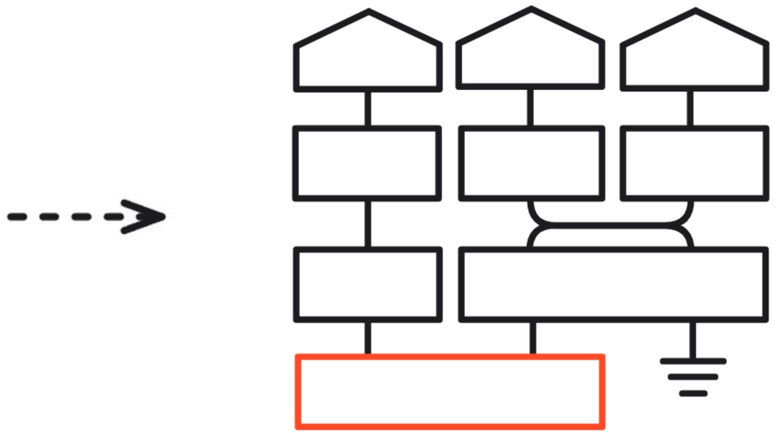

We convert *sentences* into DisCoCirc *diagrams*.

The *sentences* with shared *nouns* compose sequentially.

Unrelated *phrases* are separated.

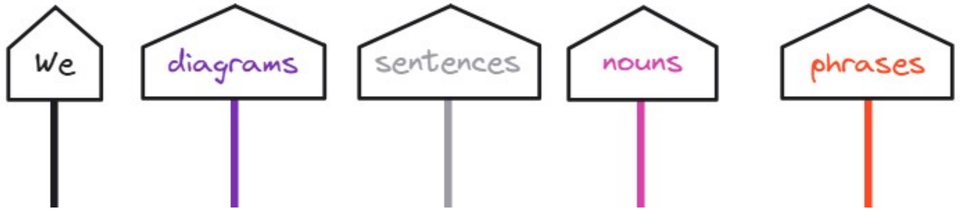

We convert sentences into **DisCoCirc** diagrams.

The sentences with **shared** nouns **compose** sequentially.

Unrelated phrases are **separated**.

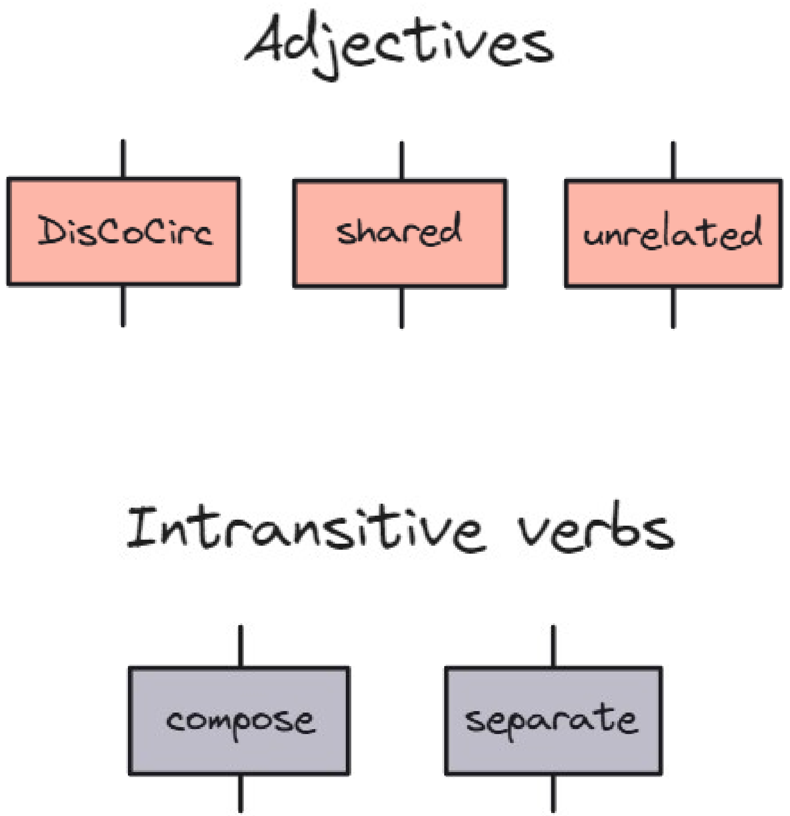

We **convert** sentences **into** DisCoCirc diagrams.

The sentences **with** shared nouns compose **sequentially**.

Unrelated phrases are separated.

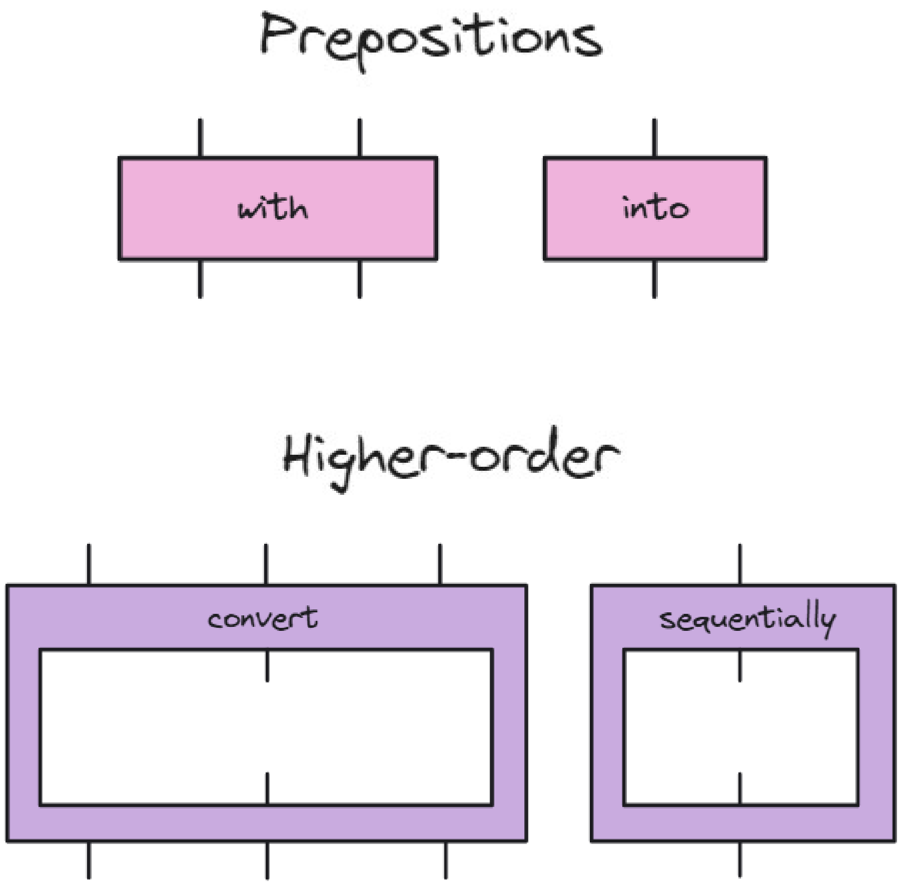

We convert *sentences* into DisCoCirc *diagrams*.

The *sentences* with shared *nouns* compose sequentially.

Unrelated *phrases* are separated.

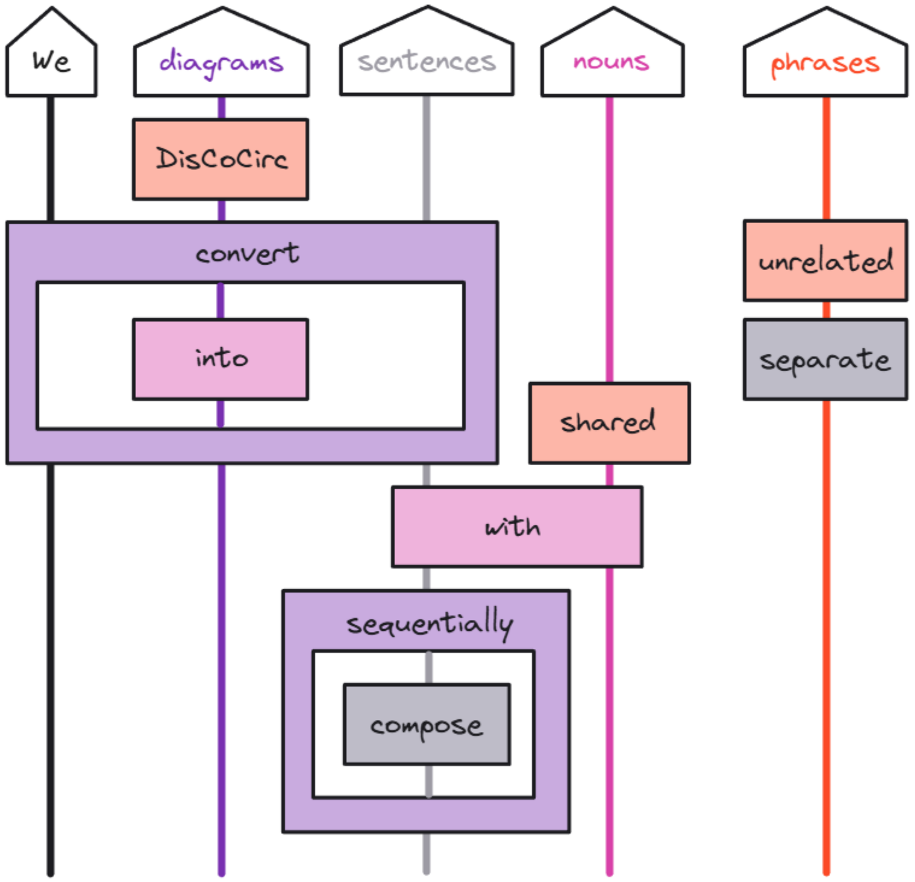

We convert sentences into DisCoCirc diagrams.

The sentences with shared nouns compose sequentially.

Unrelated phrases are separated.

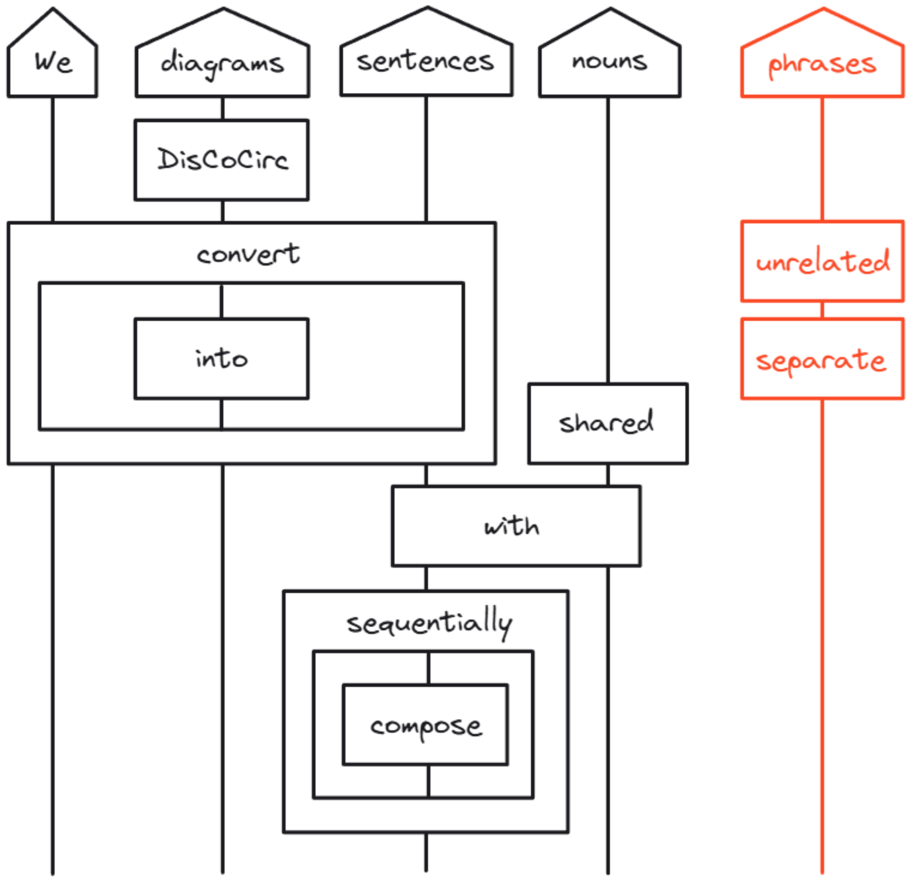

Quantum DisCoCirc

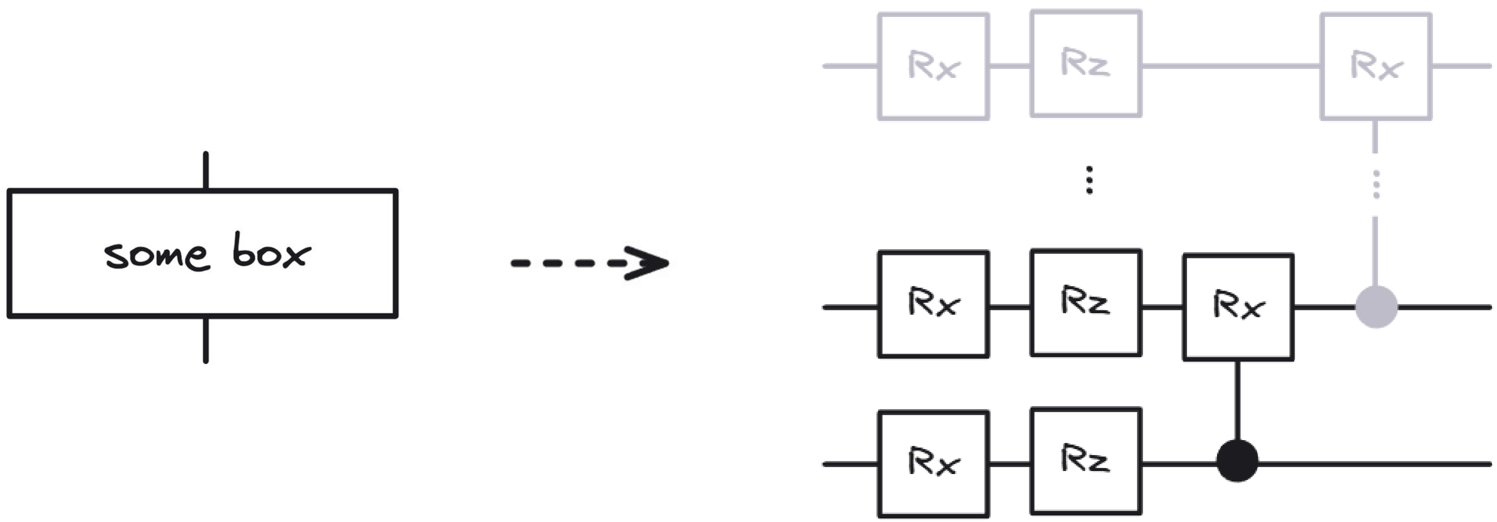
some box
Rx
Rz
Rx
Rx
Rz
Rx
Rx
Rz

Ansatze

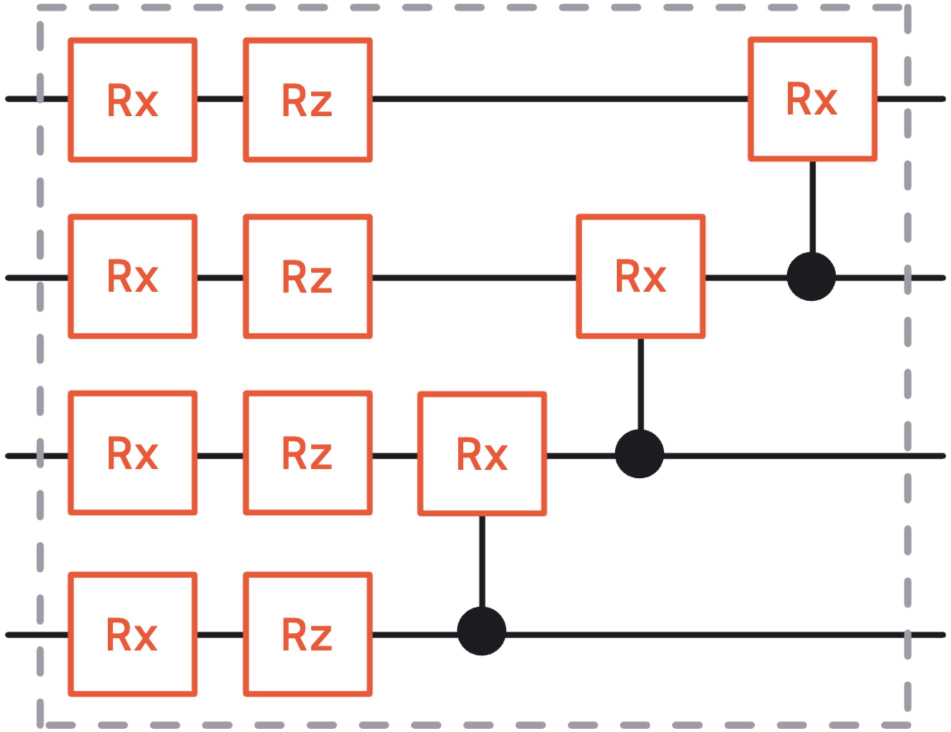

Sim4

S. Sim et al. Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms, 2019.

Euler Decomposition

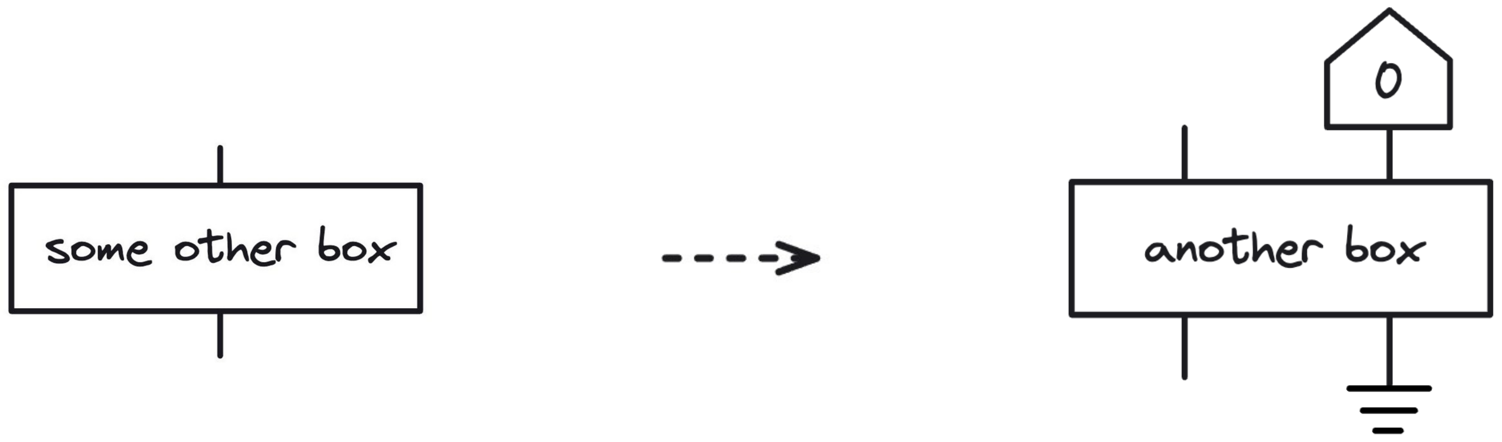
some other box
0
another box

Sandwich

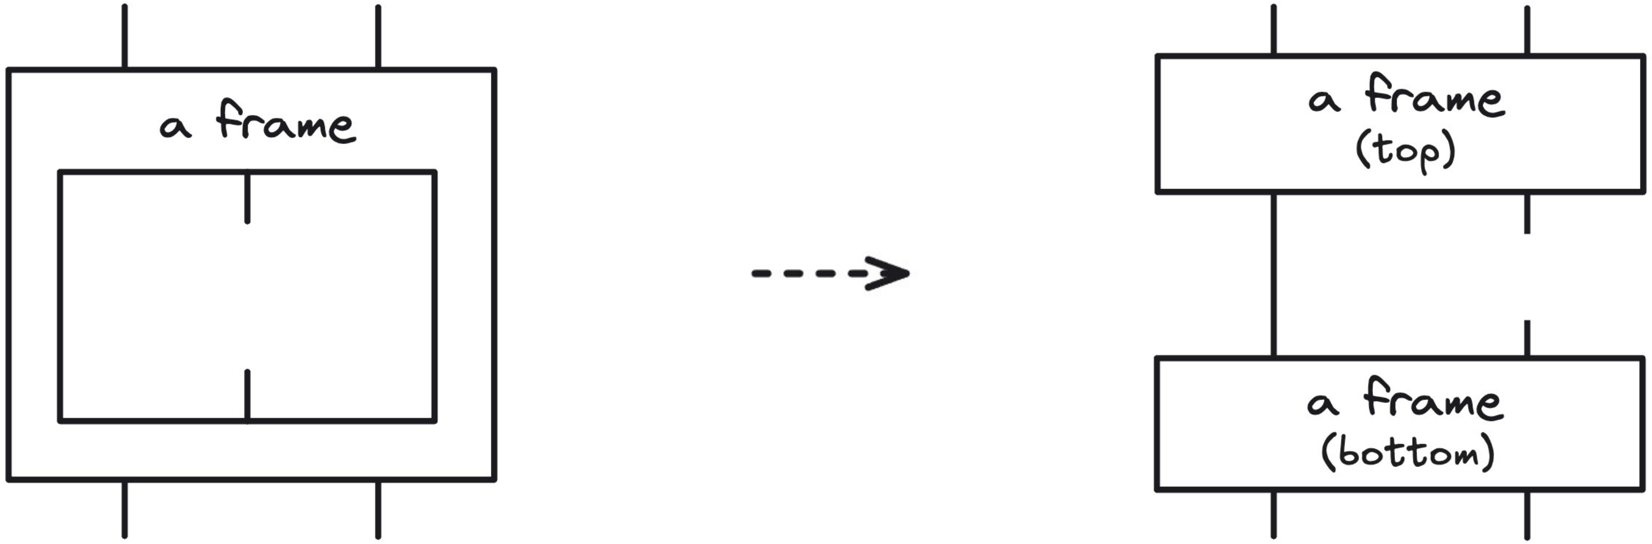

Questions

A question

Is Alice going in the same direction as Bob?

Assertions

Is Alice going in the same direction as Bob?

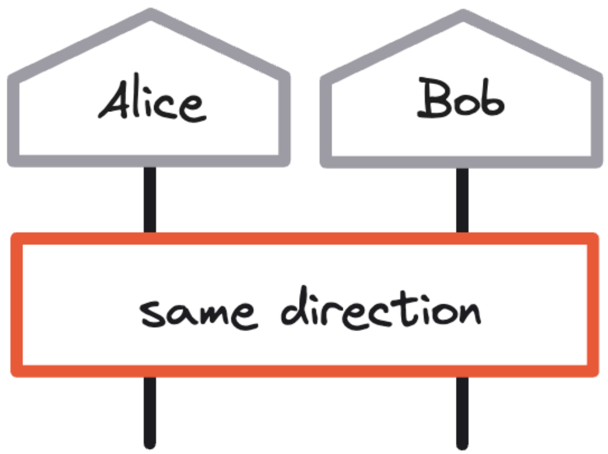

Yes

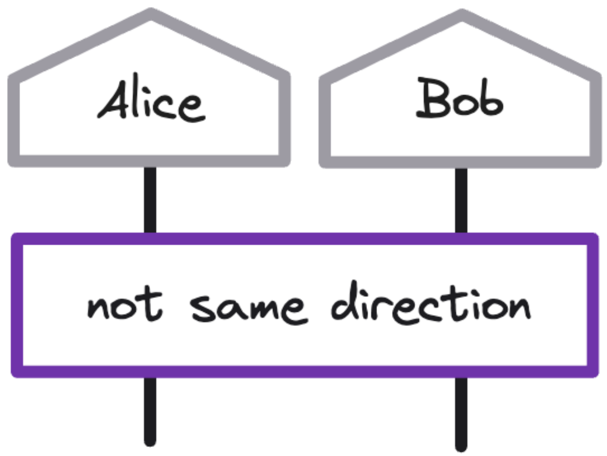

No

Is Alice going in the same direction as Bob?

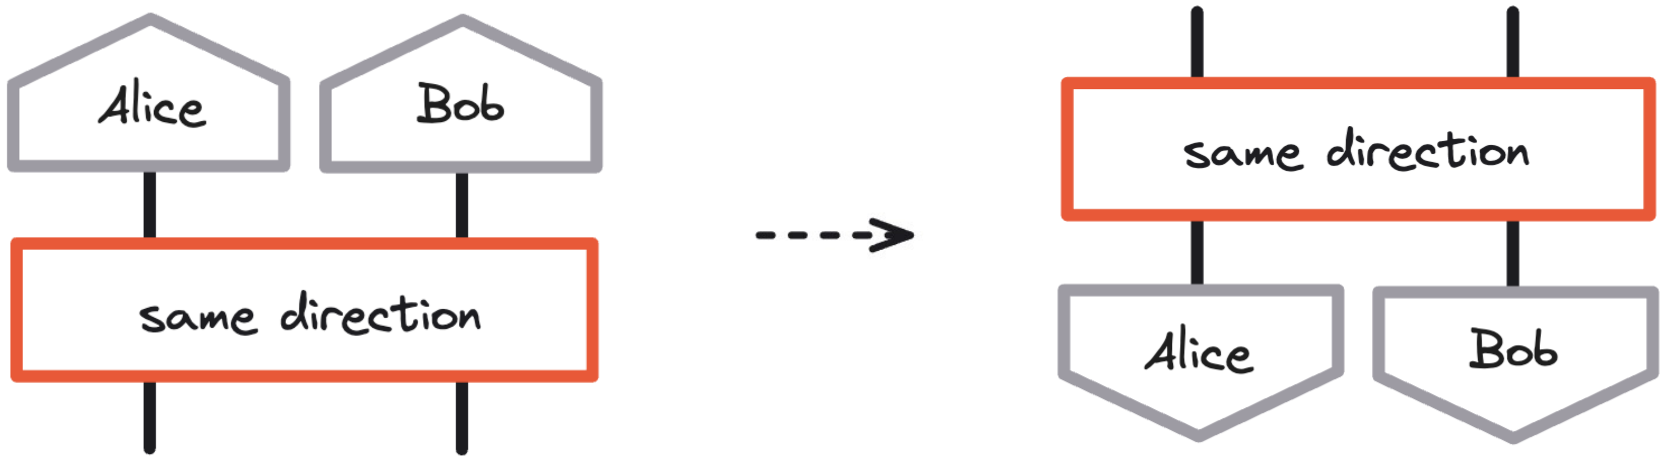

Yes

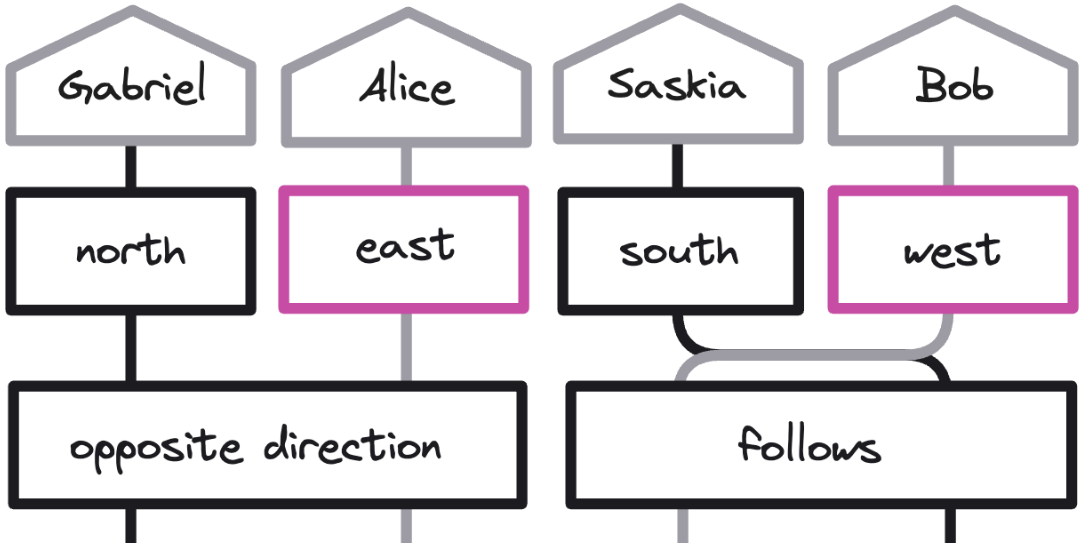
Gabriel
Alice
Saskia
Bob
north
east
south
west
opposite direction
follows

Yes

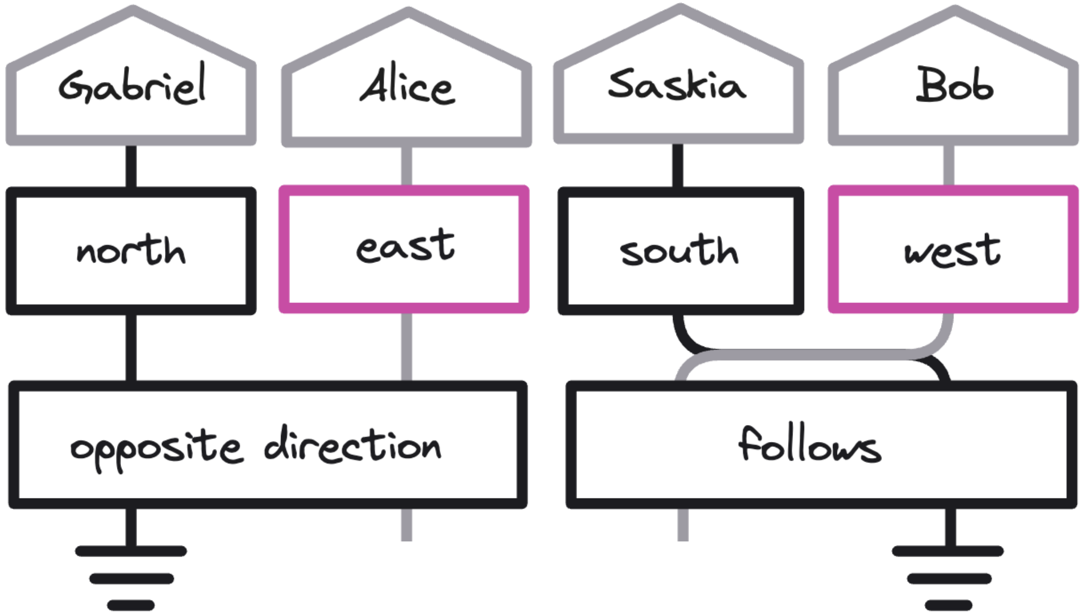
Gabriel
Alice
Saskia
Bob
north
east
south
west
opposite direction
follows

Yes

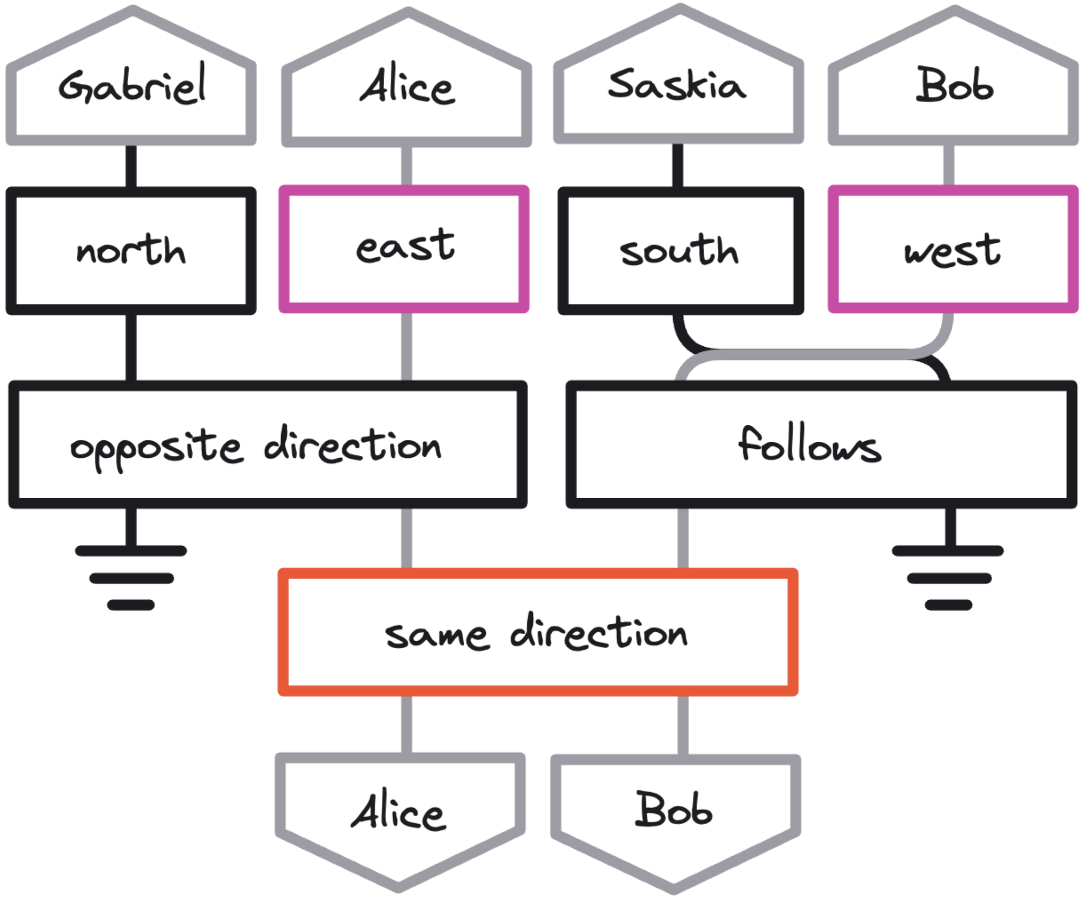
Gabriel
Alice
Saskia
Bob
north
east
south
west
opposite direction
follows
same direction
Alice
Bob

Yes

Task

“Following” Dataset

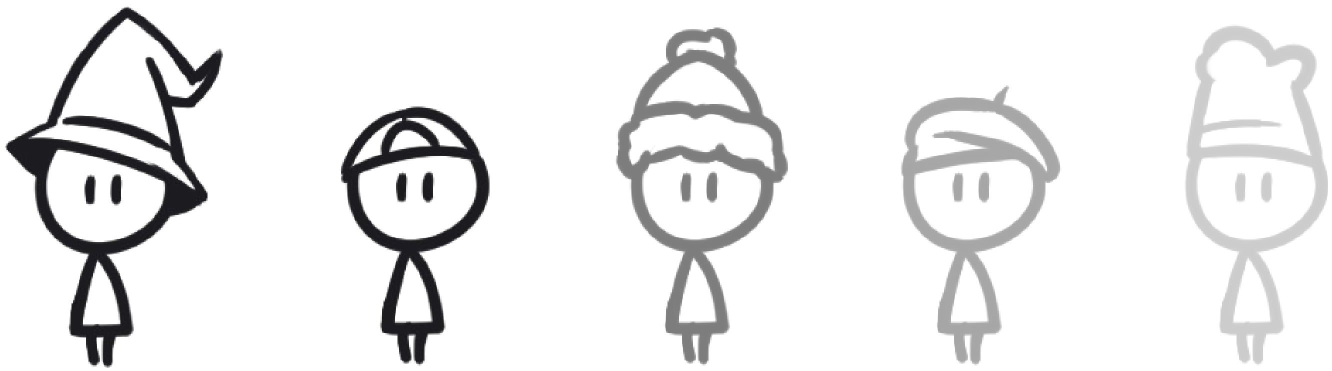

Alice Bob …

2

directions

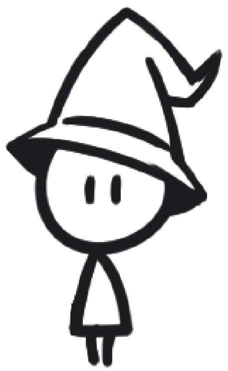

4

directions

Alice walks south

north

east

west

2

direction
s

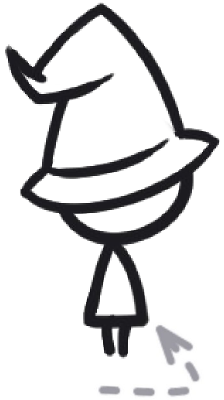

Alice walks south
north
east
west

Alice turns around
left
right

4

direction
s

2 directions

4 directions

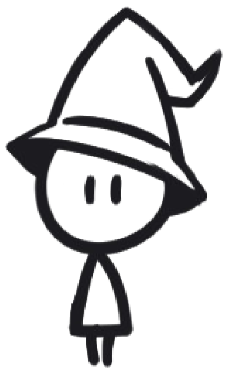

Alice walks south

north

east

west

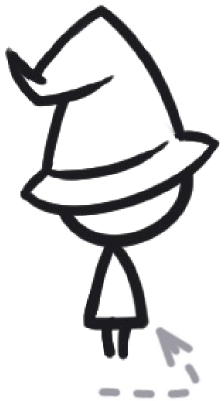

Alice turns around

left

right

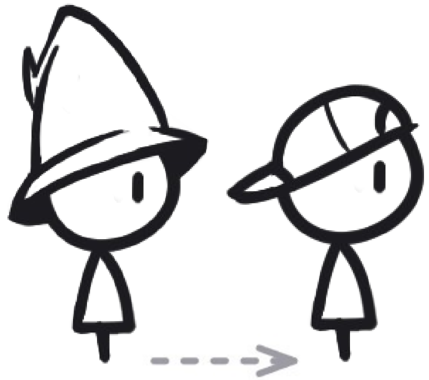

Alice follows Bob.

goes in the opposite direction of

Productivity

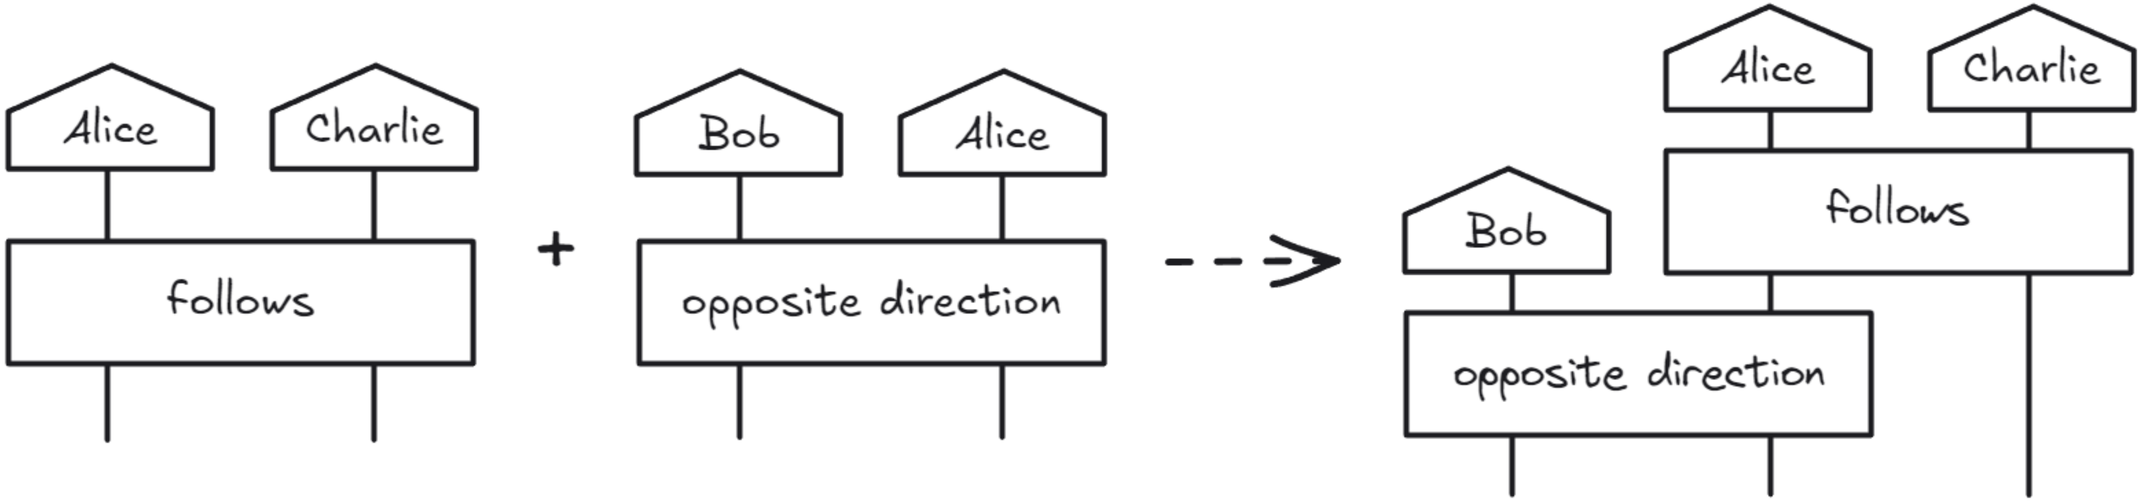

Dataset Split

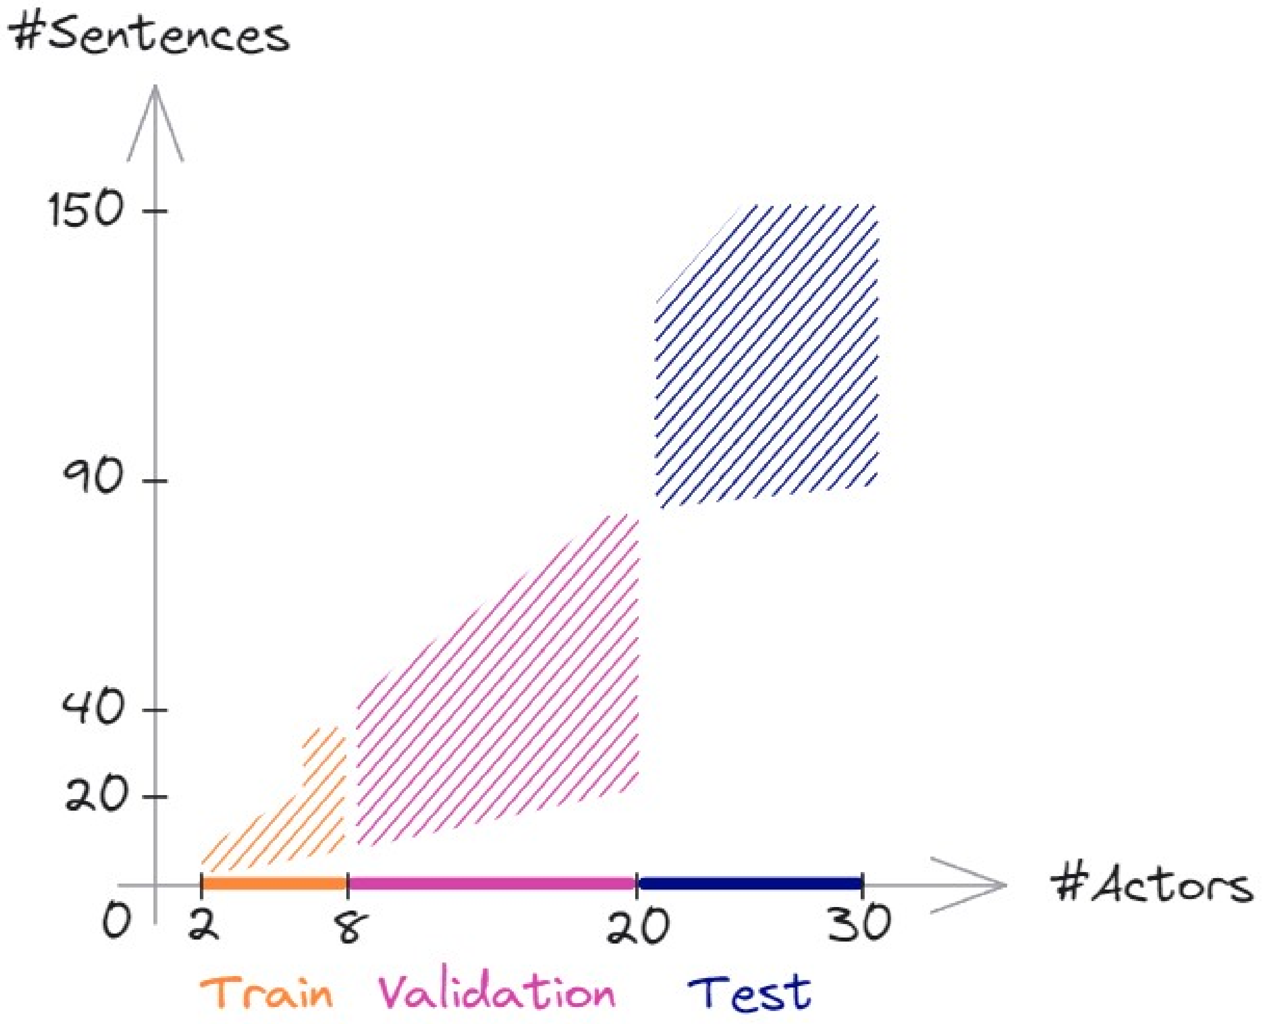

Pipeline

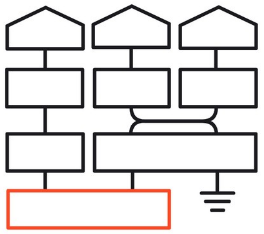

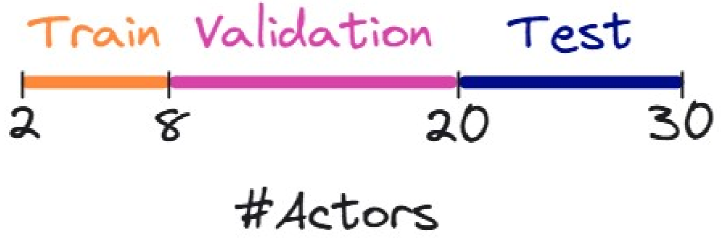

Pipeline

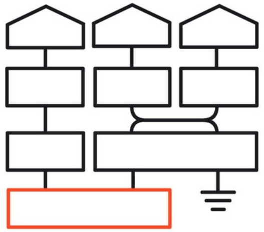

Tensor Network

PyTorch

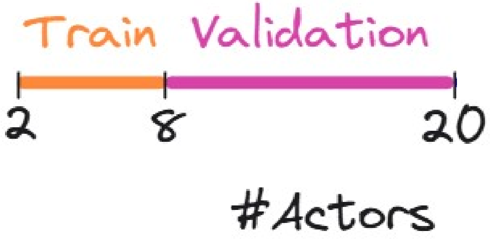

Pipeline

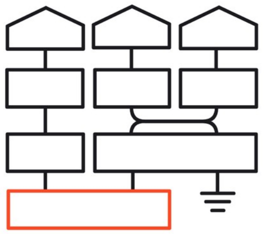

TKET

QUANTINUUM H1-1

qujax

via

QUANTINUUM NEXUS

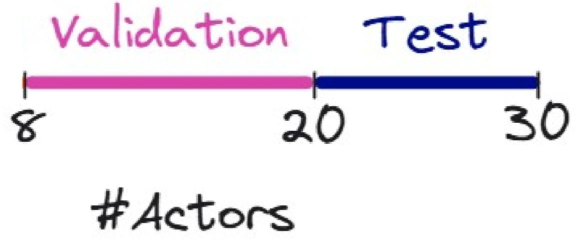

2

direction

Evaluating Test dataset

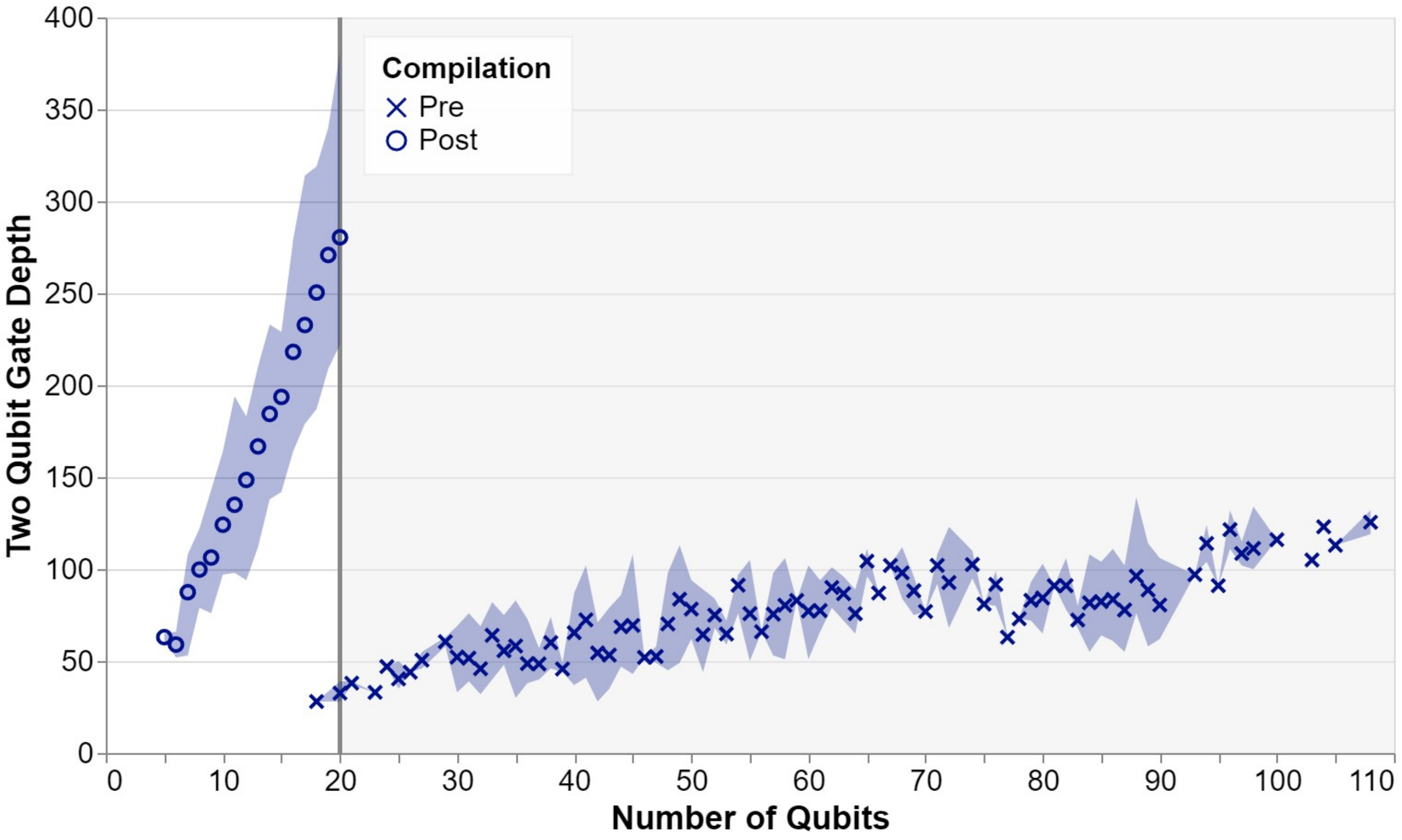

TKET

Compositional Generalisation

"train small, test big"

2

direction
s

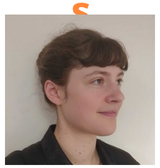

Compositional generalisation

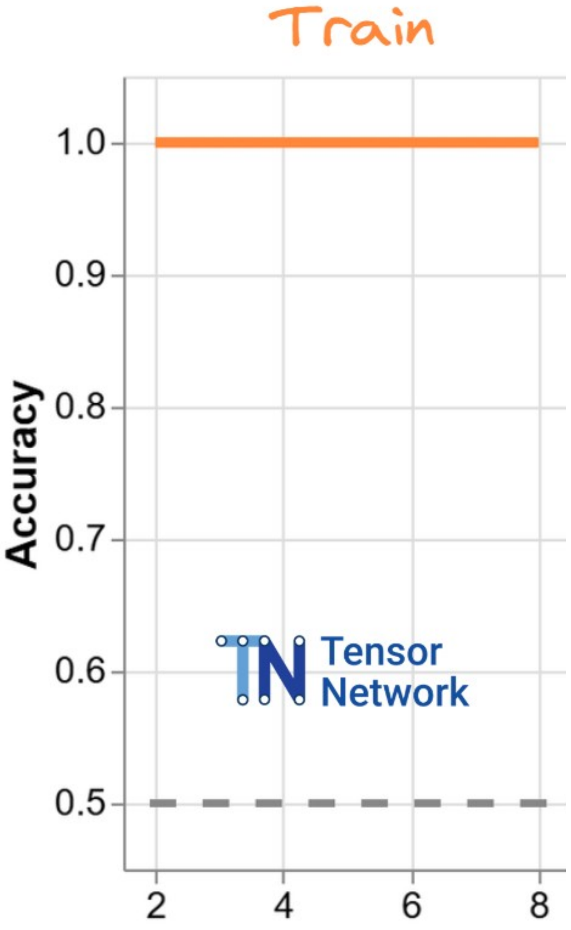

2

direction
s

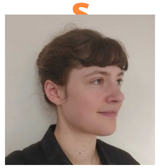

Compositional generalisation

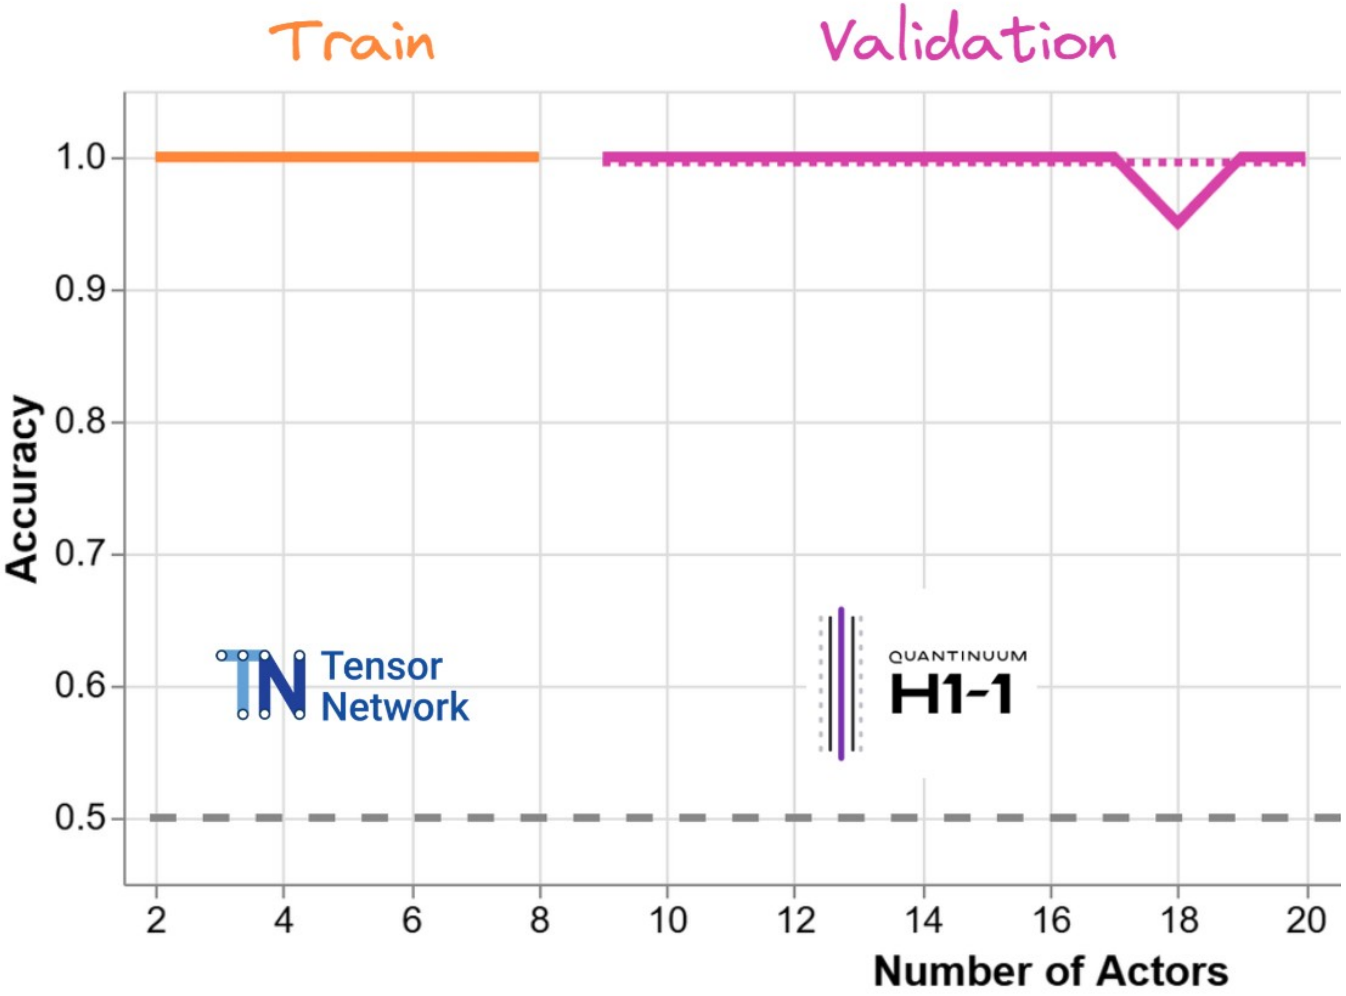

2

directions

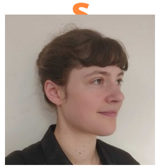

Compositional generalisation

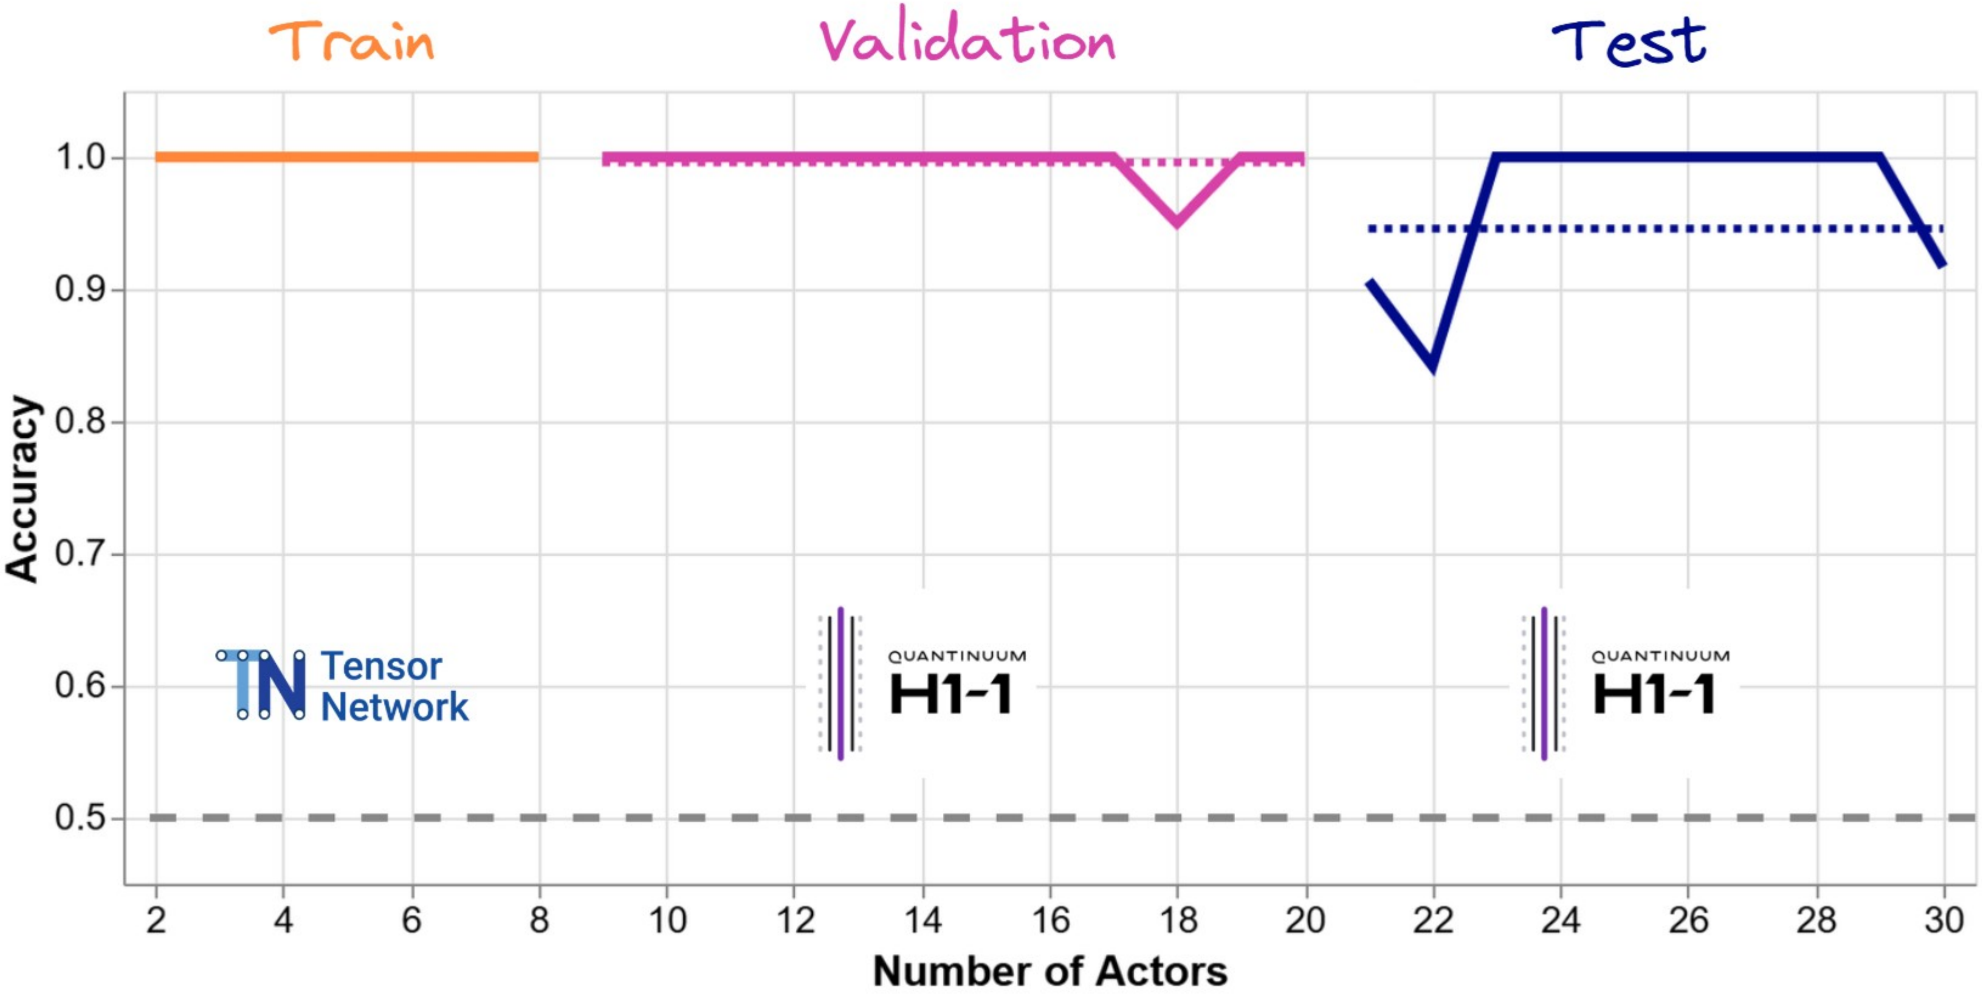

2

directions

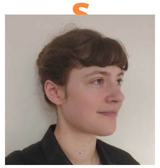

Compositional generalisation

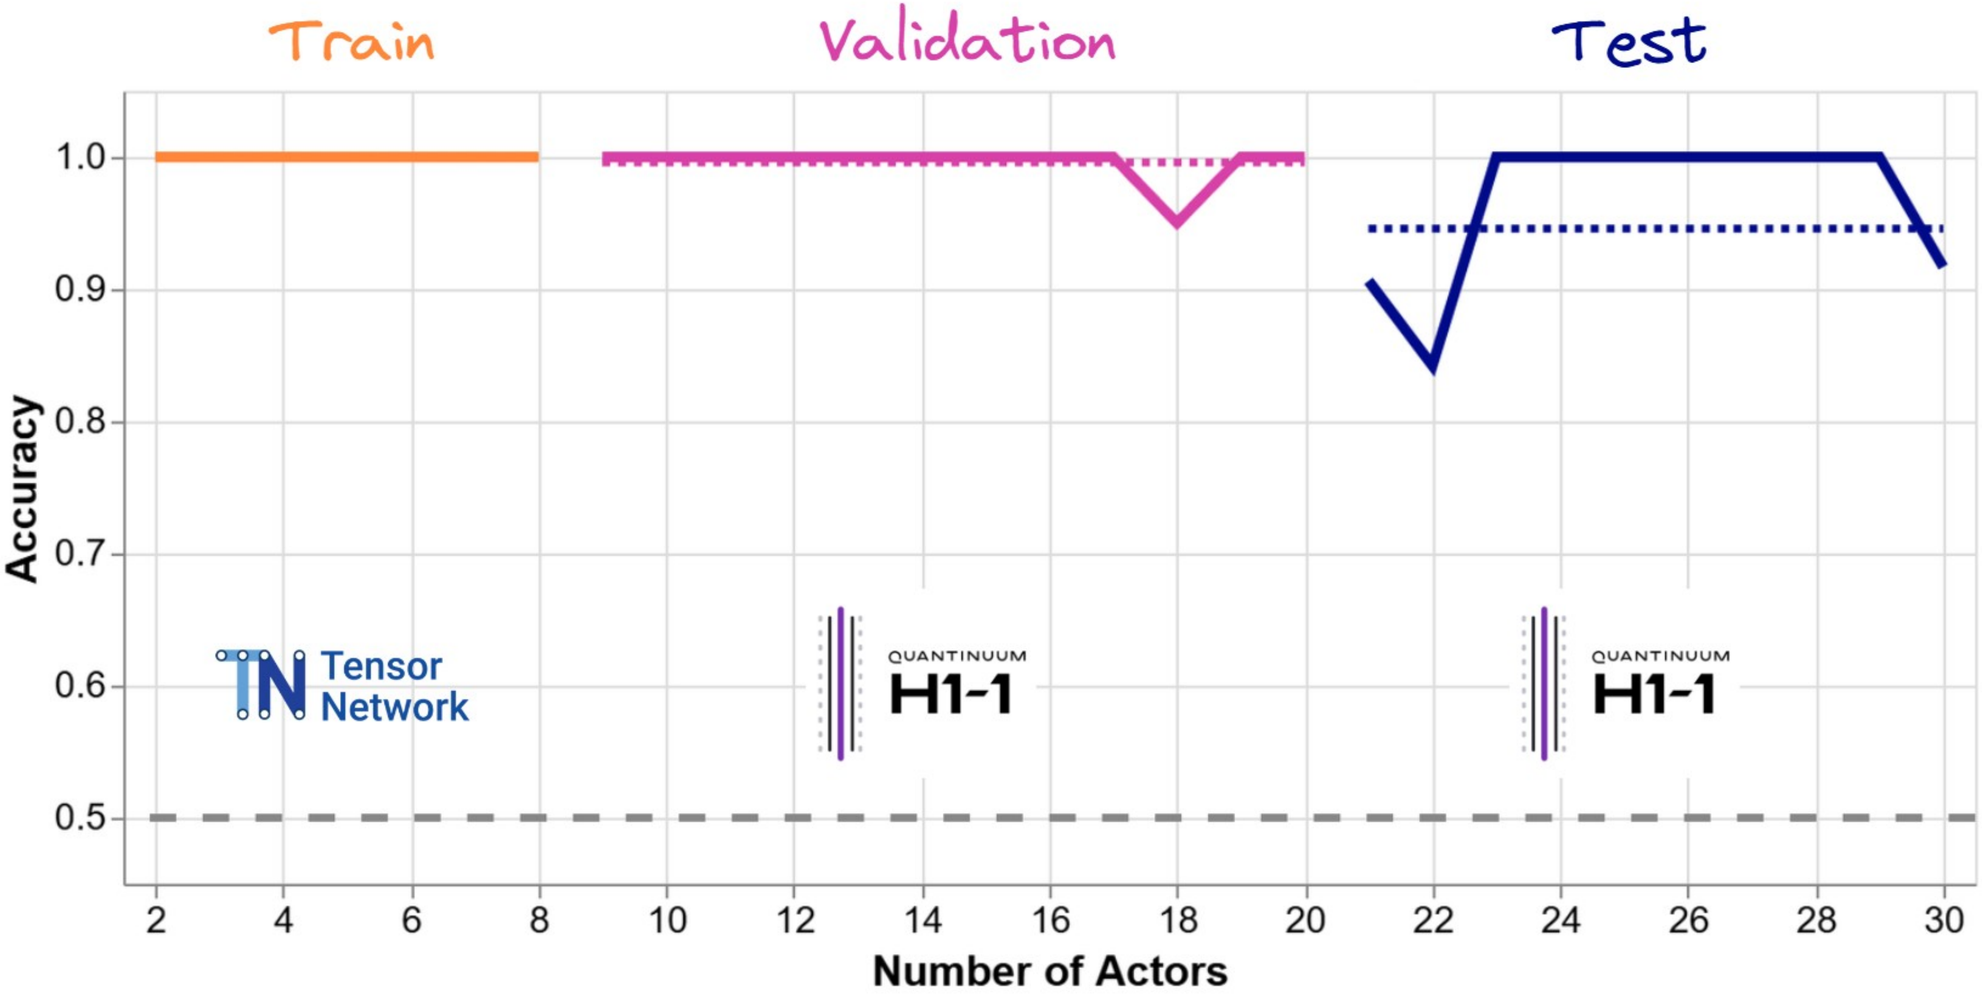

2

directions

Noise Simulations on Validation dataset

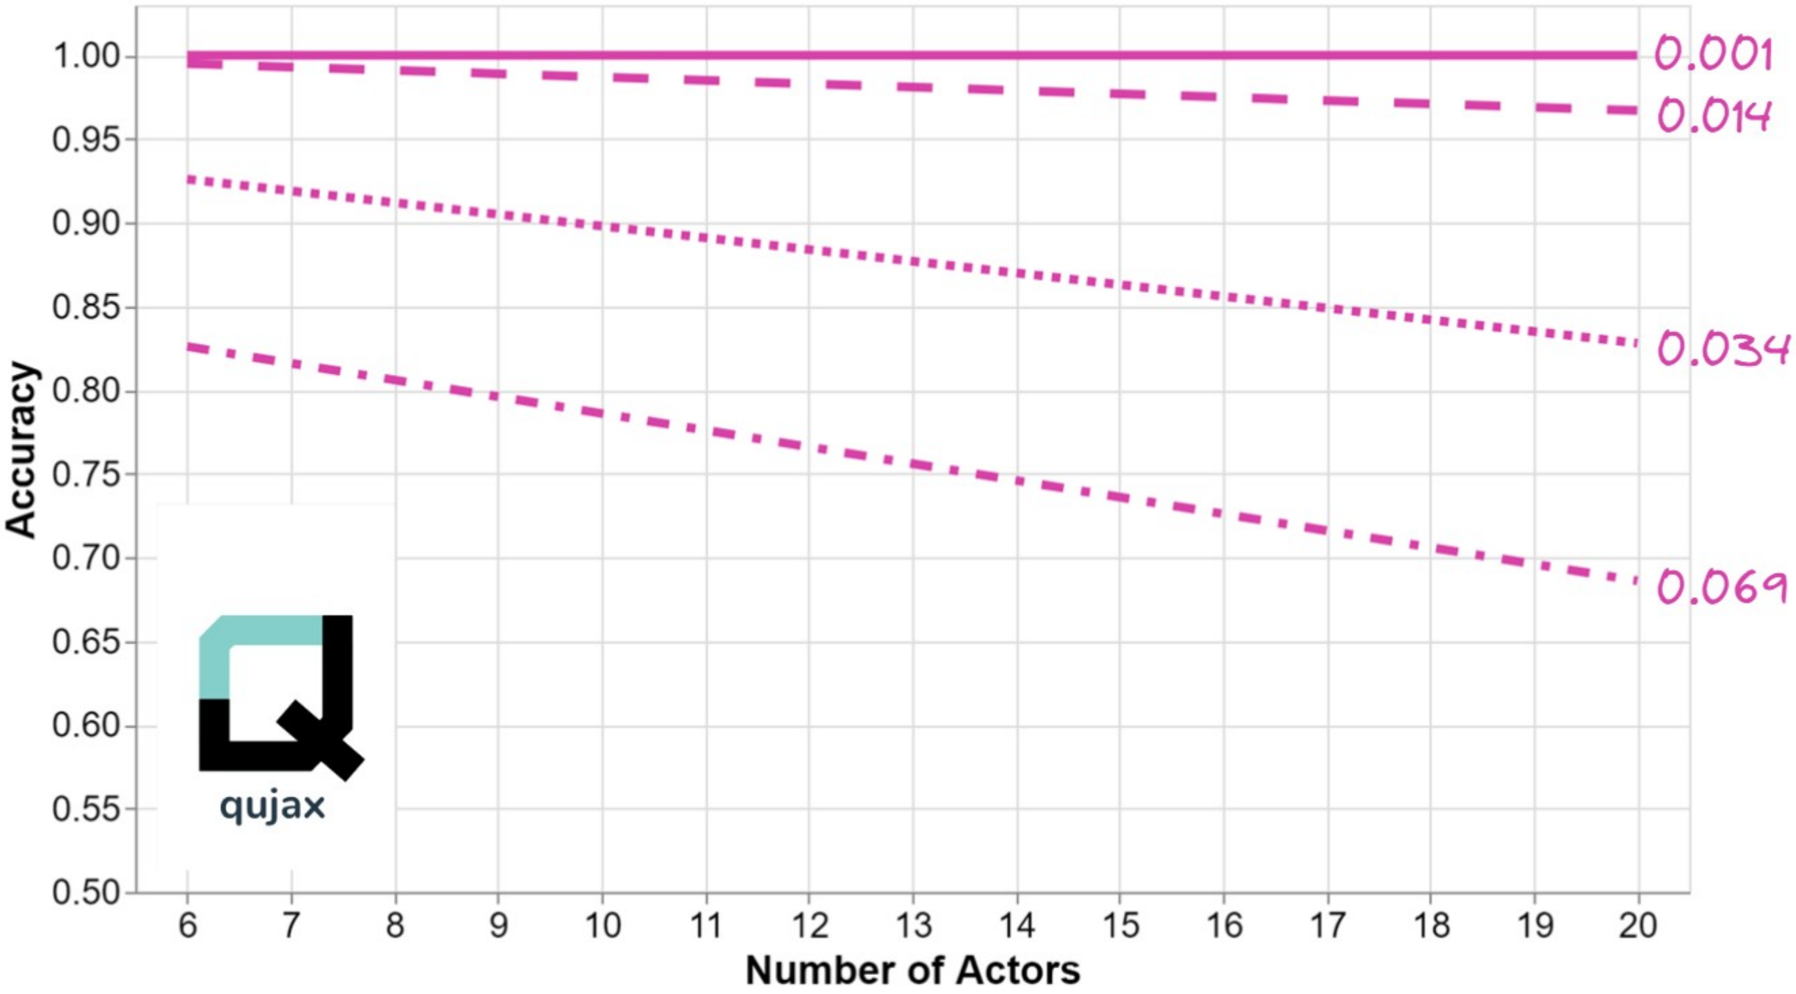

2
directions

Noise Simulations on Validation dataset

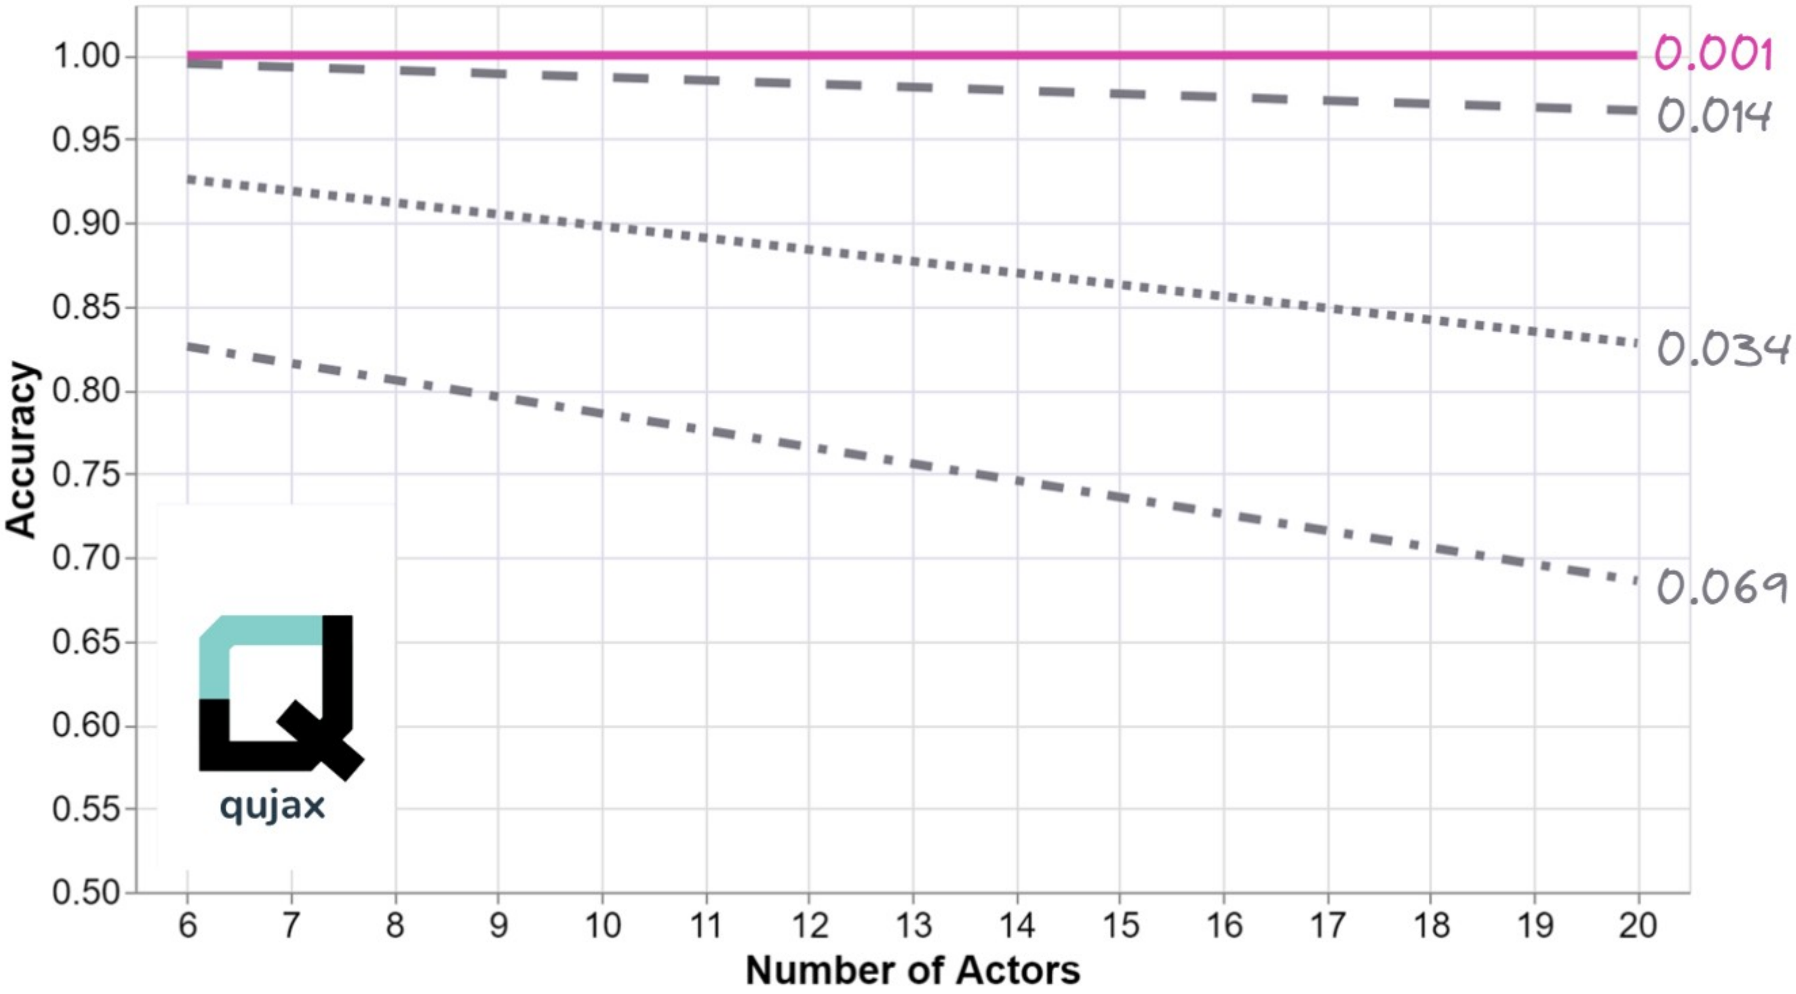

4

directions

Compositional generalisation

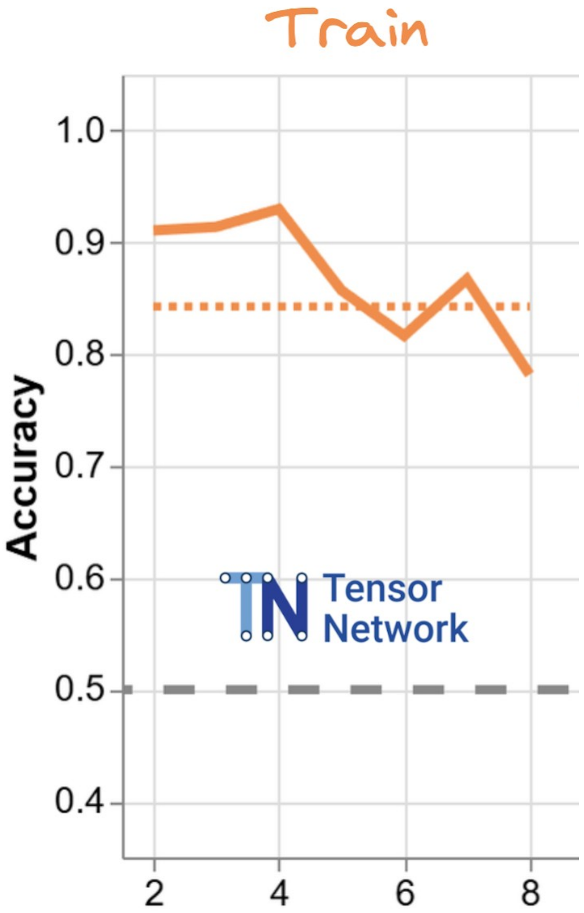

4

directions

Compositional generalisation

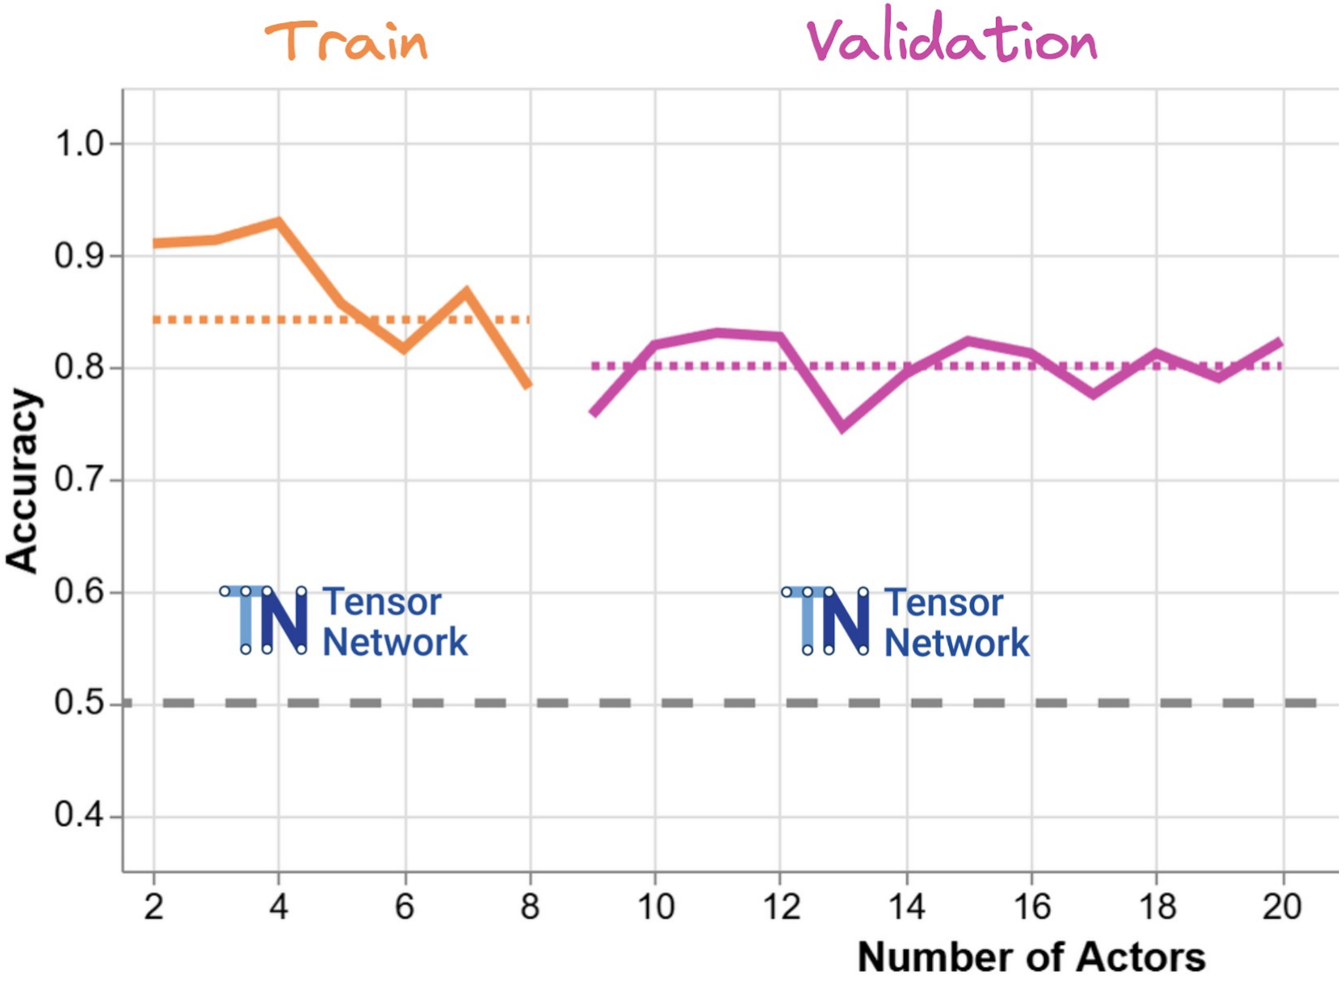

4

directions

Compositional generalisation

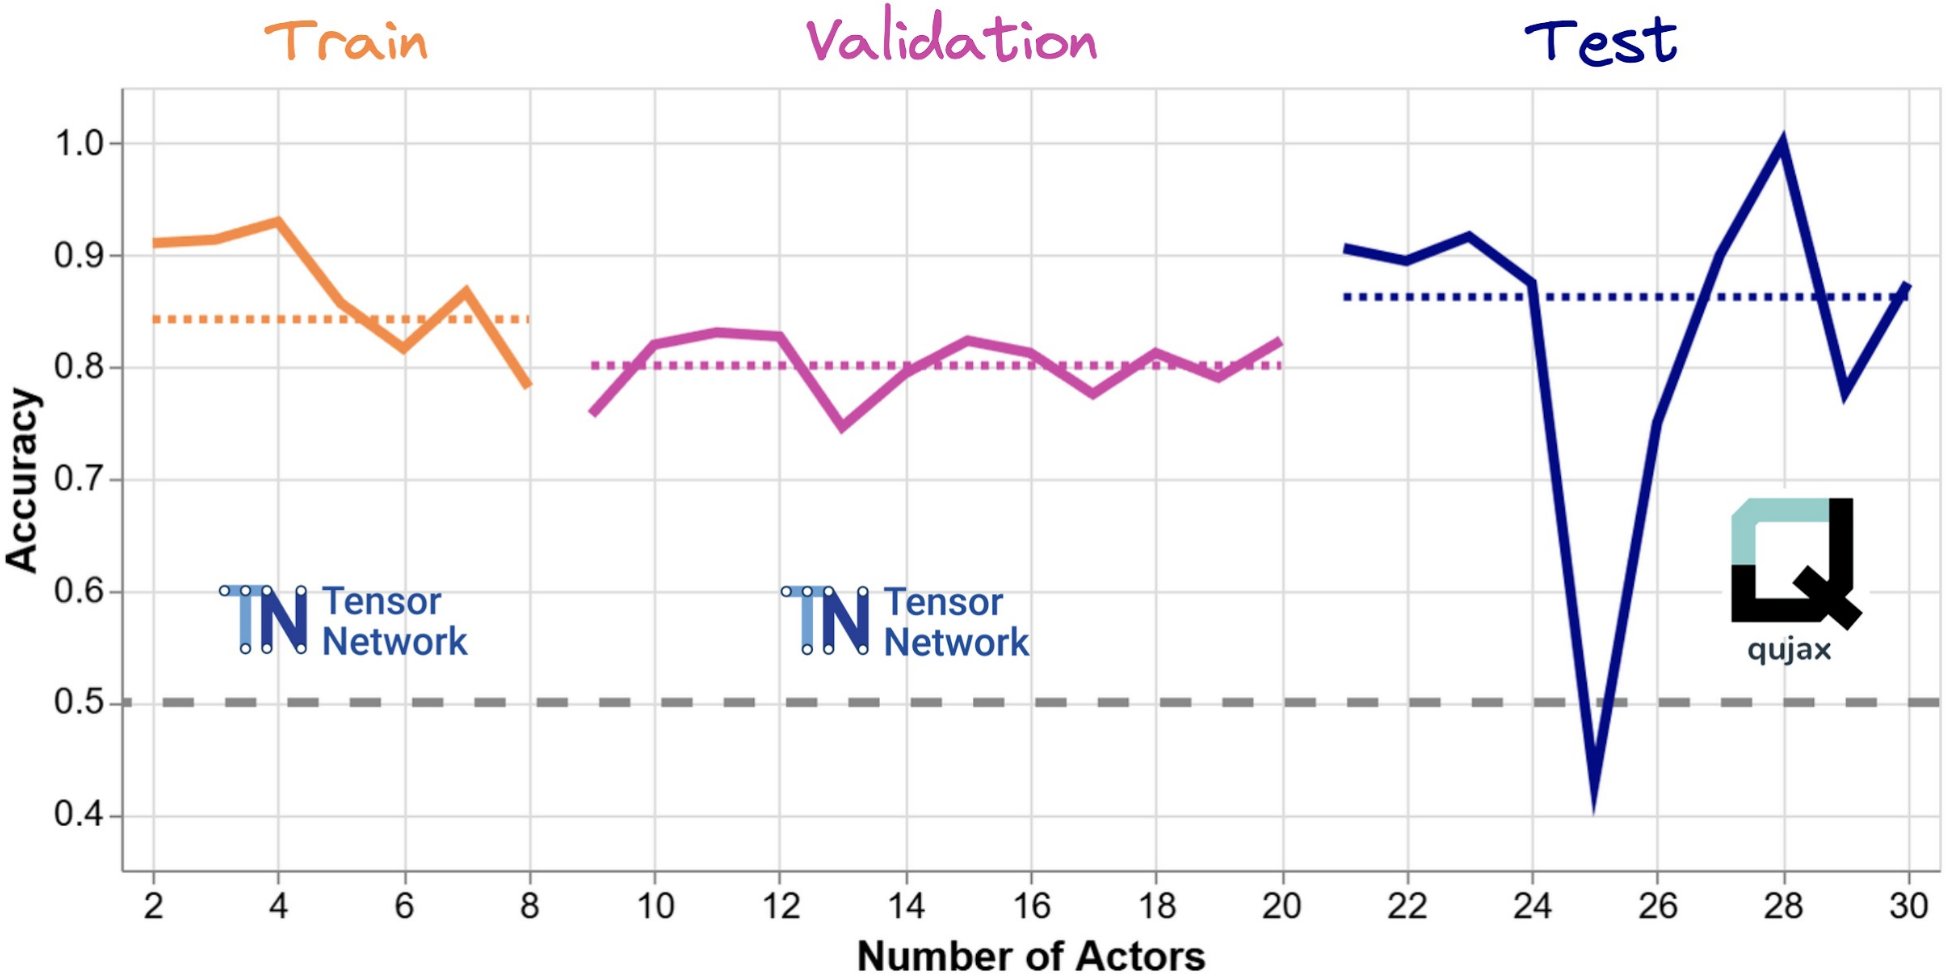

4
directions

Compositional generalisation

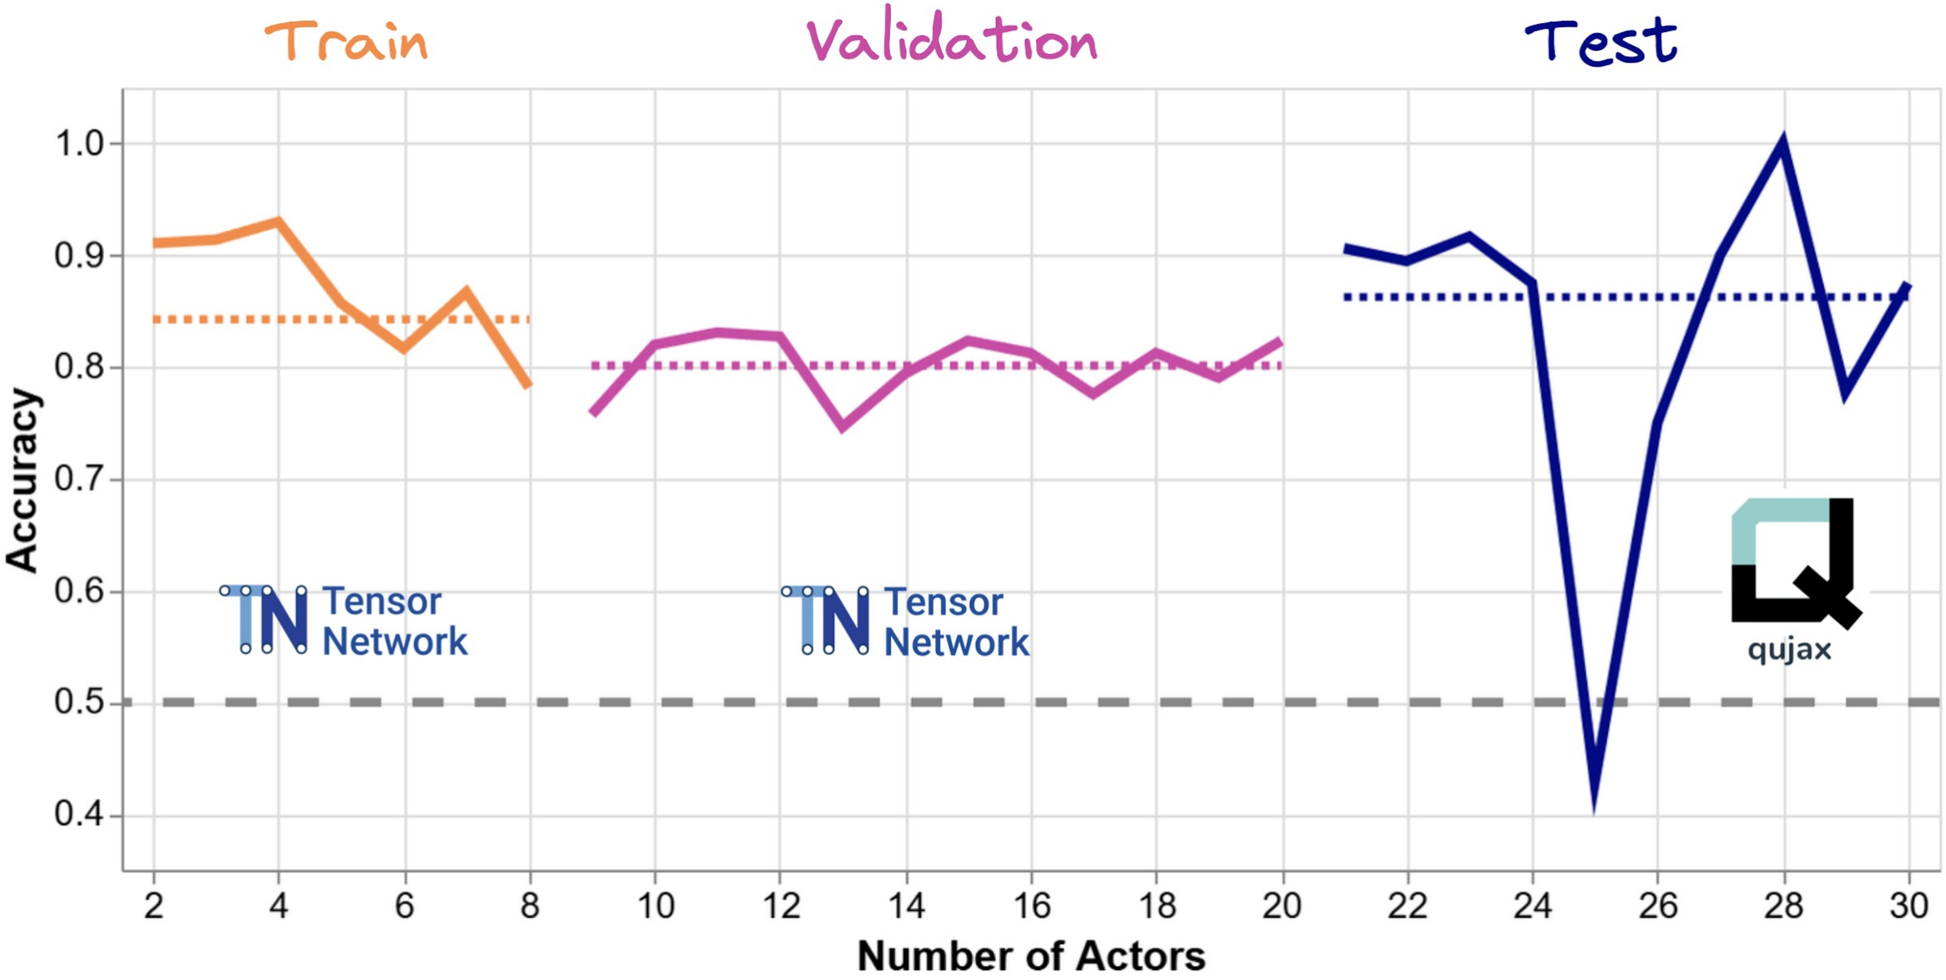

Compositional Interpretability

arXiv: 2406.17583

Black Box

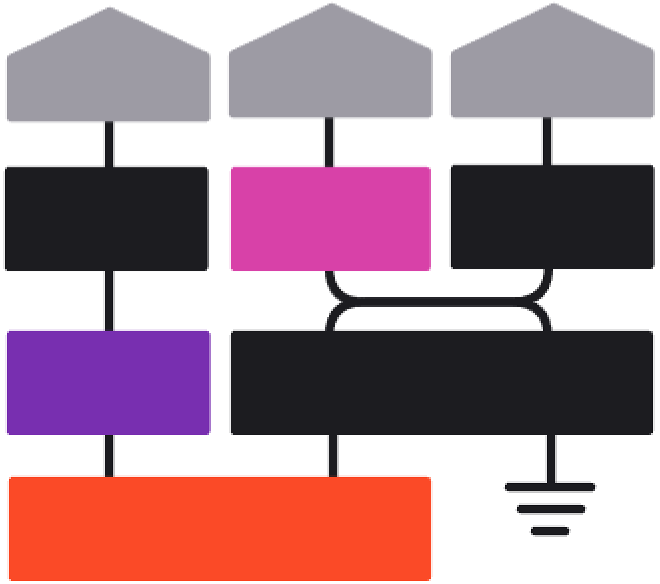

Many Black Boxes

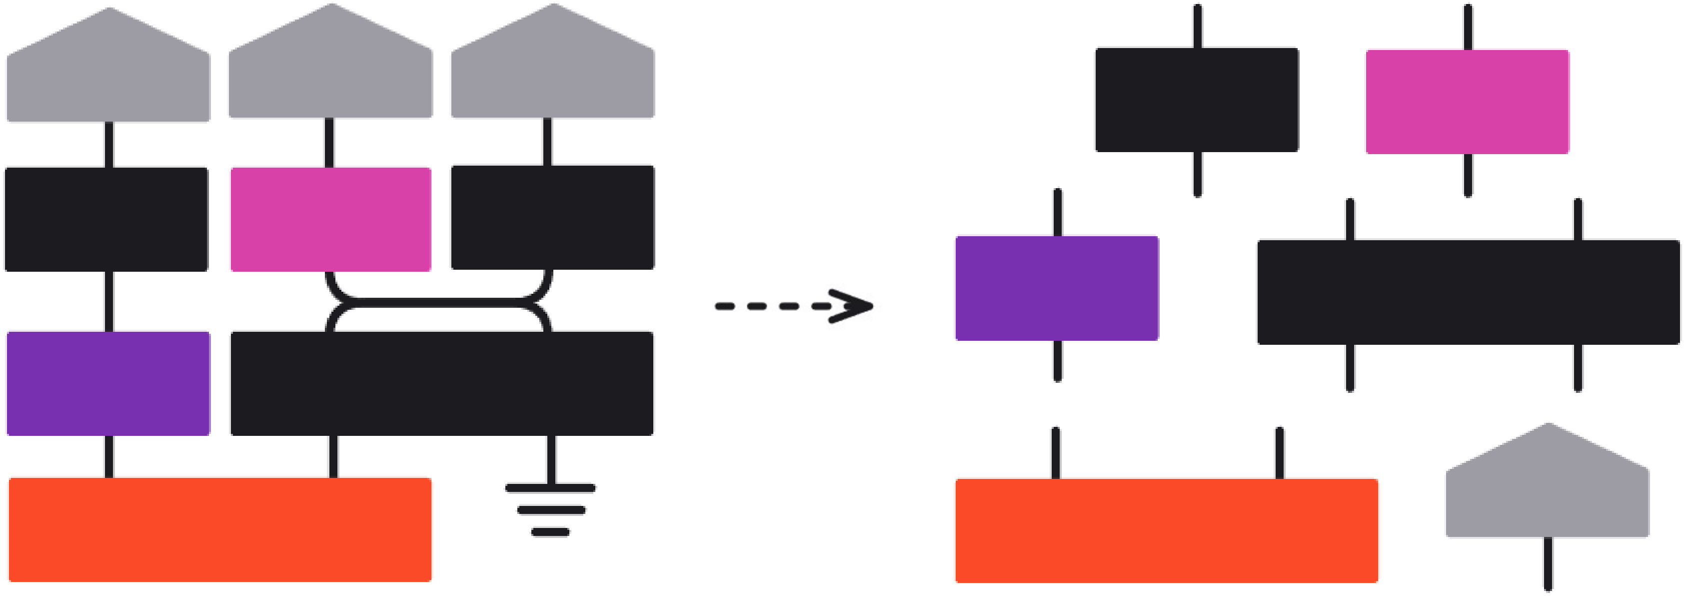

Many Black Boxes ...decomposed!

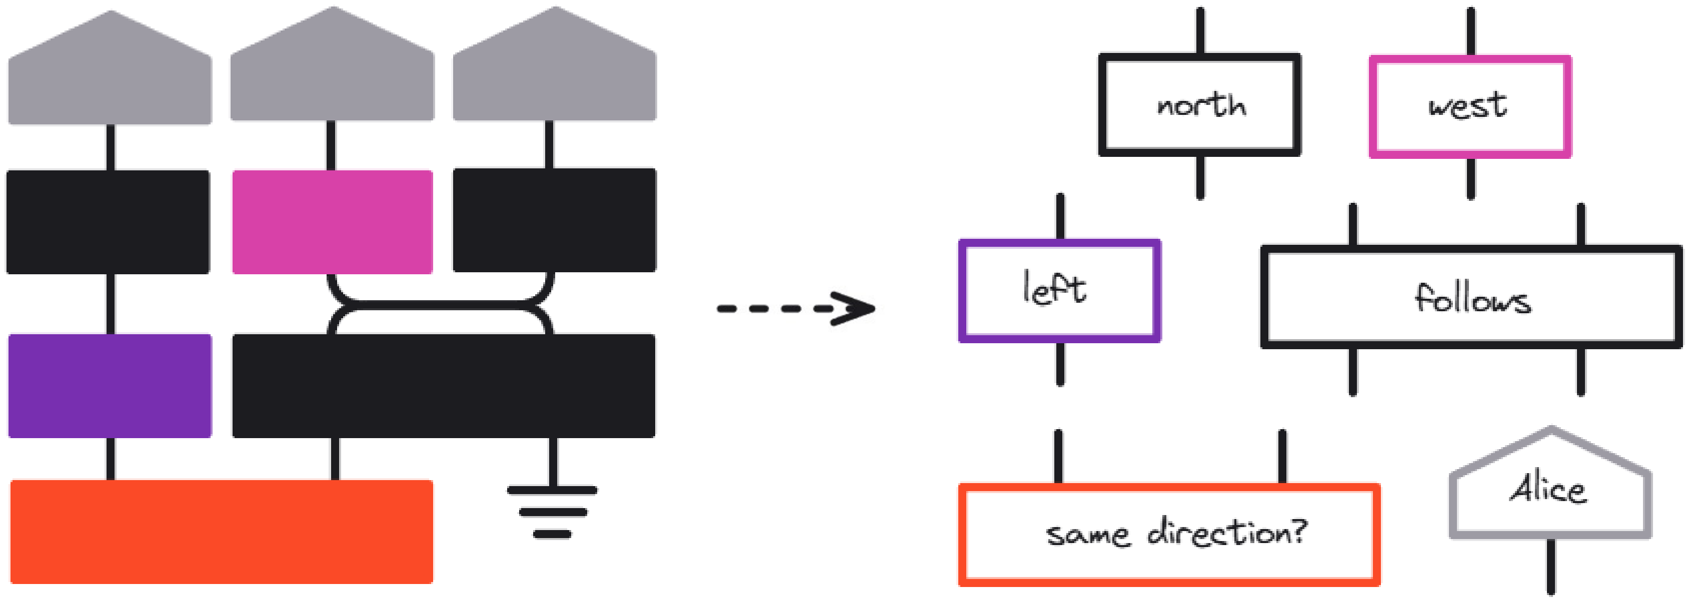

Many Black Boxes ...decomposed!

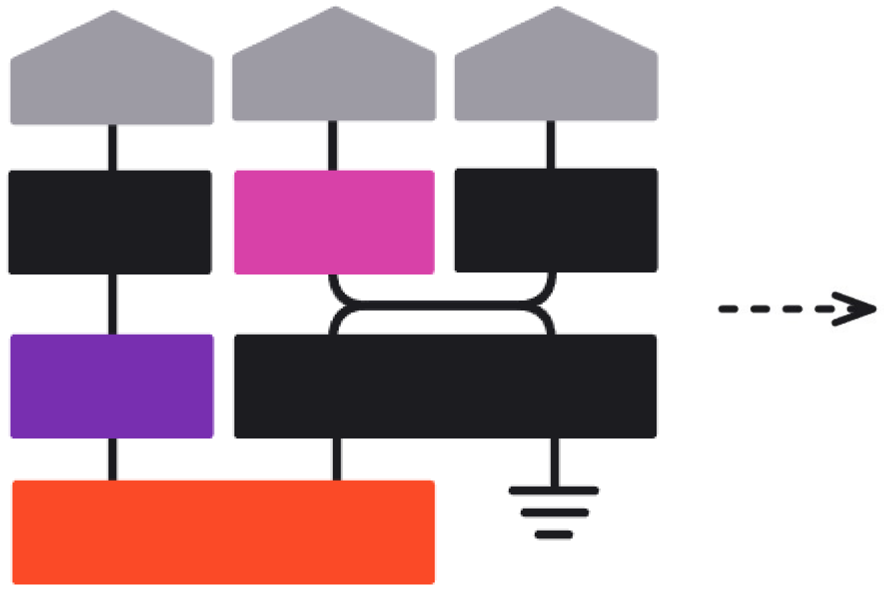

Many Black Boxes

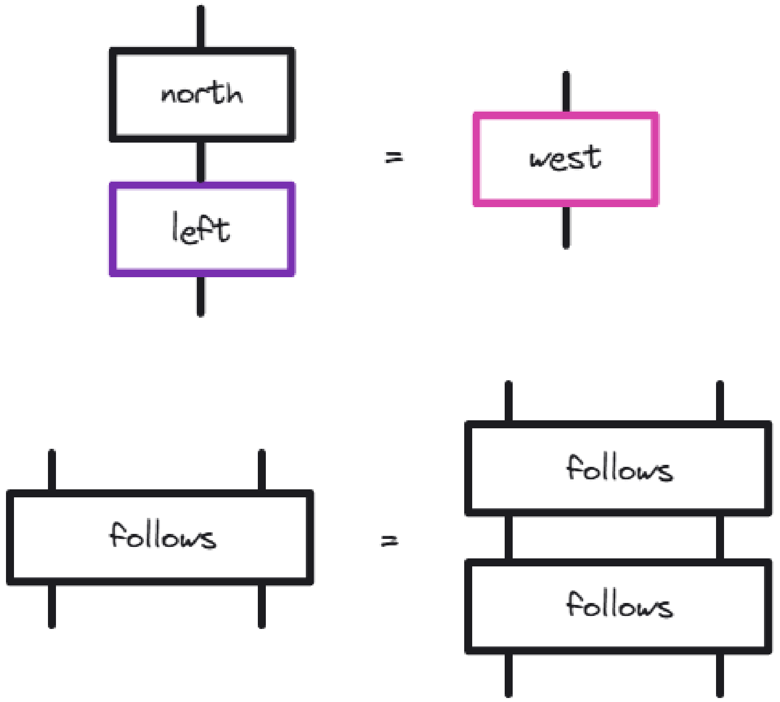

...recomposed!

2

directions

Built-in axioms

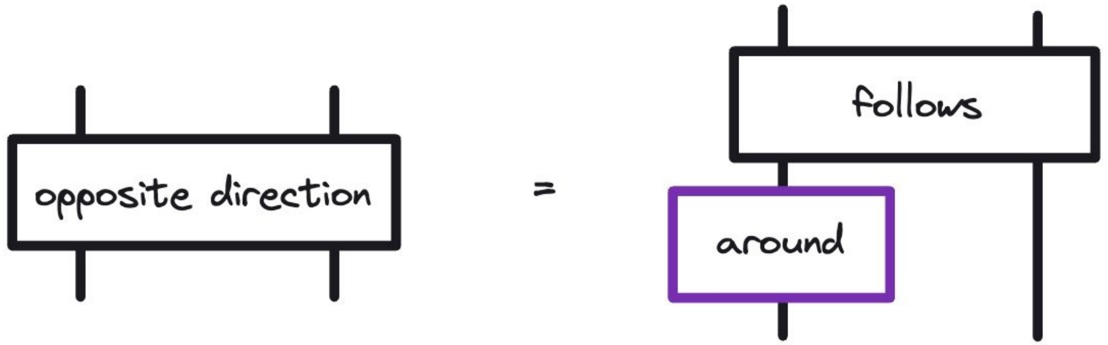

Built-in axioms

Interpreting Quantum

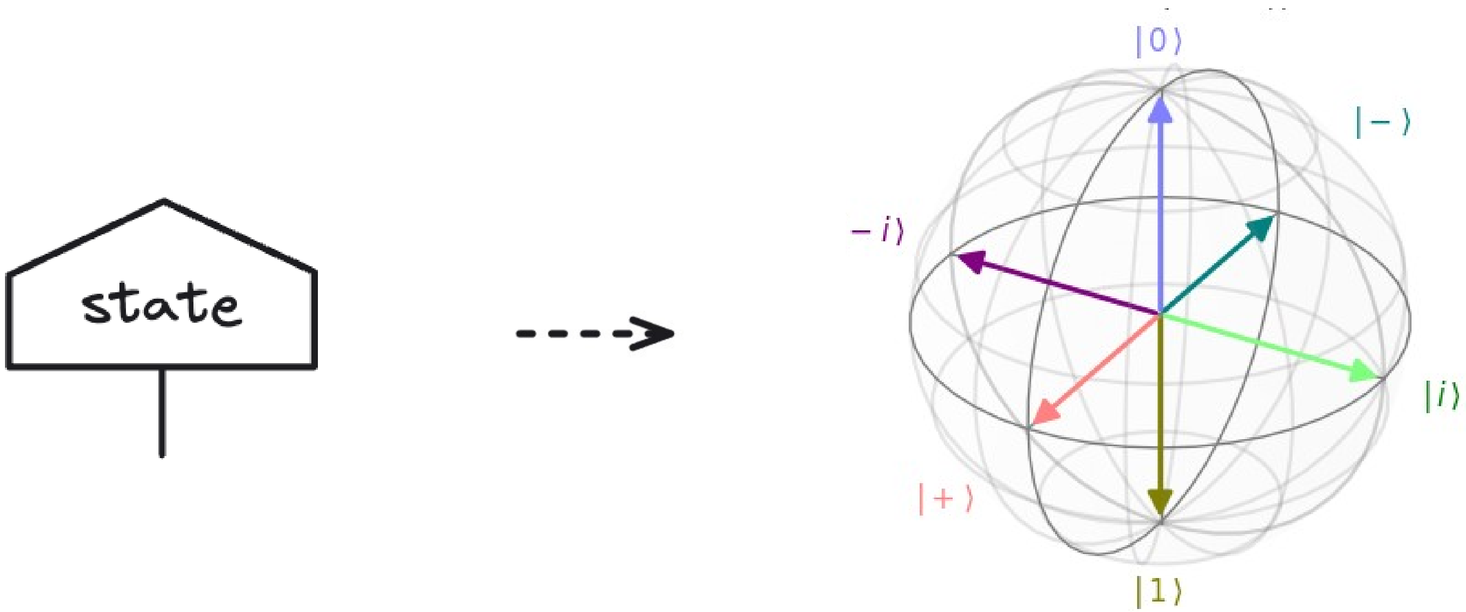
state
|0⟩
|−⟩
−i⟩
|i⟩
|+⟩
|1⟩

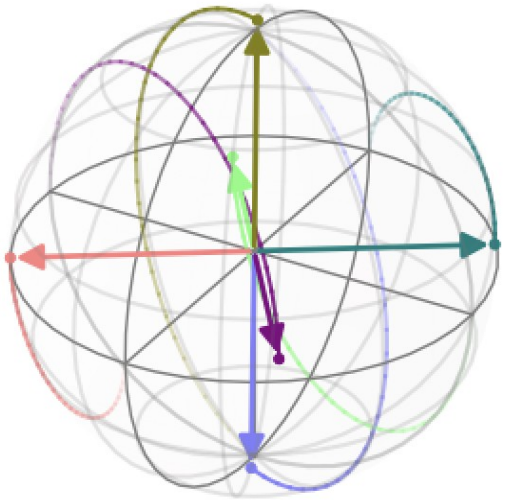

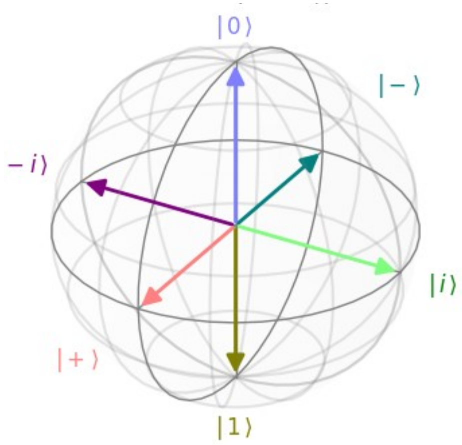

Reference

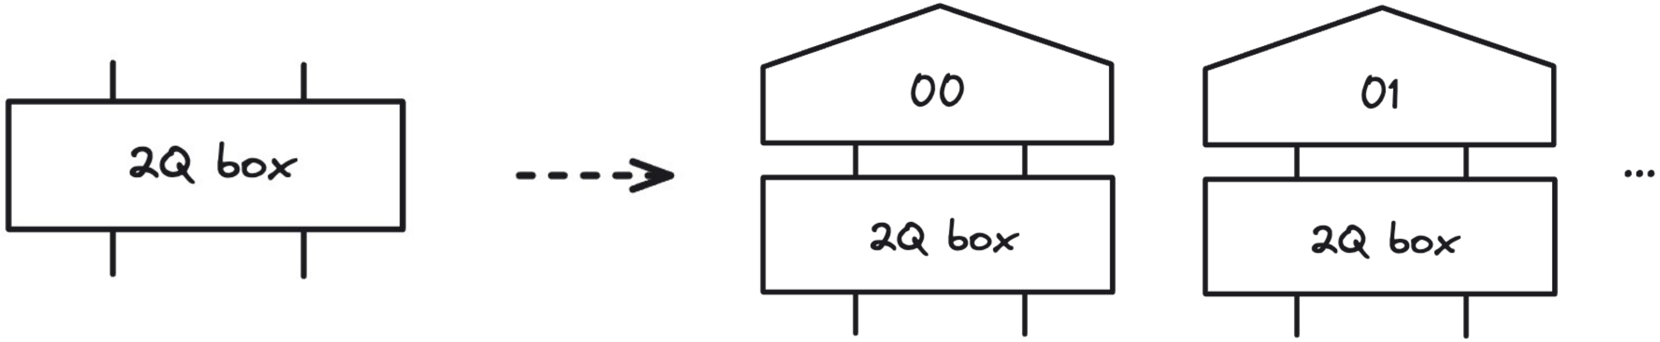
2Q box
00
2Q box
01
2Q box
...

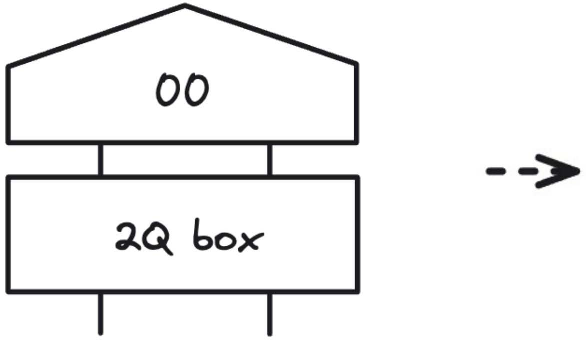

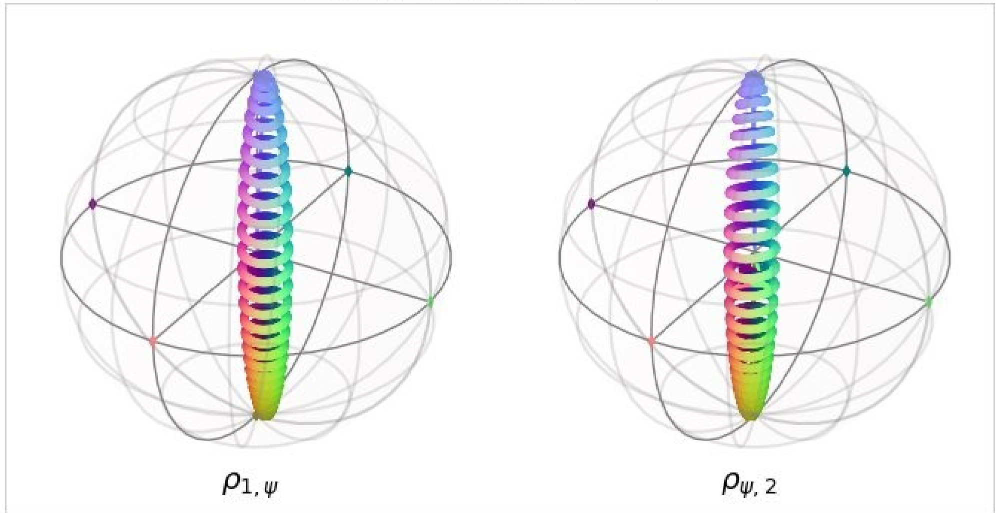

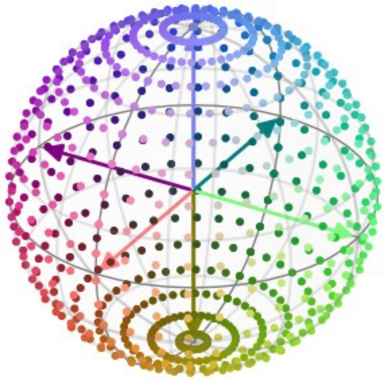

Reference

J. Altepeter et al. Multiple-qubit quantum state visualization, 2009.

2

directions

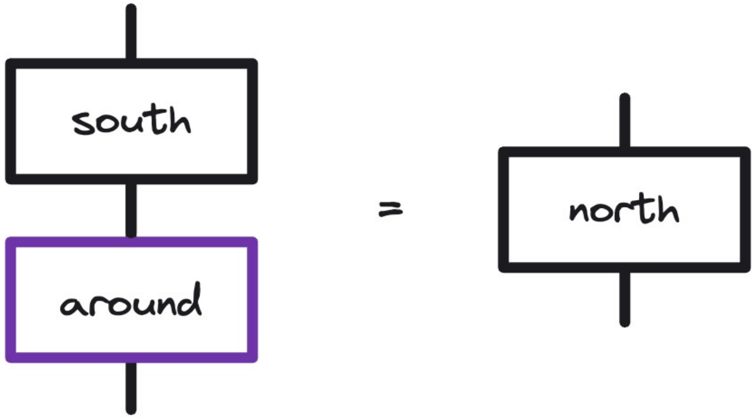

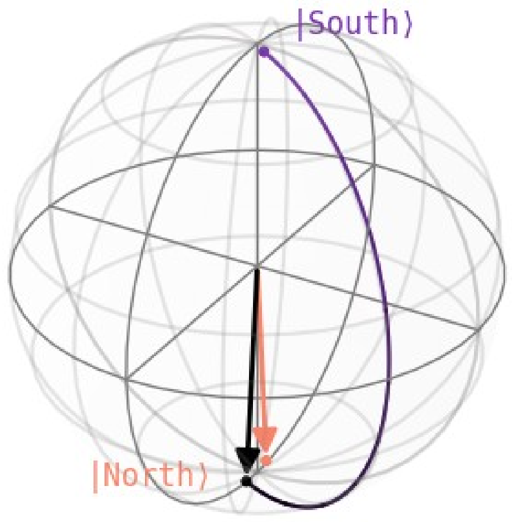

2

directions

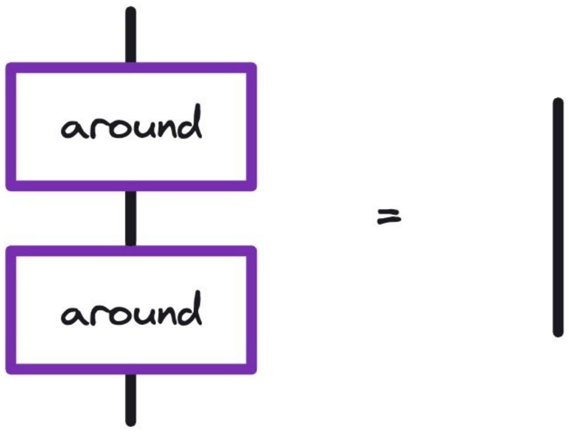

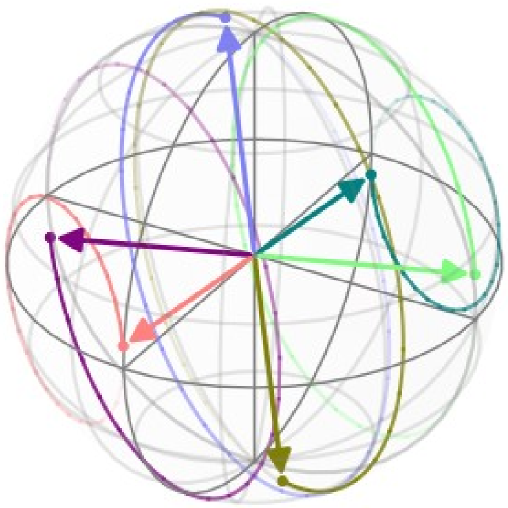

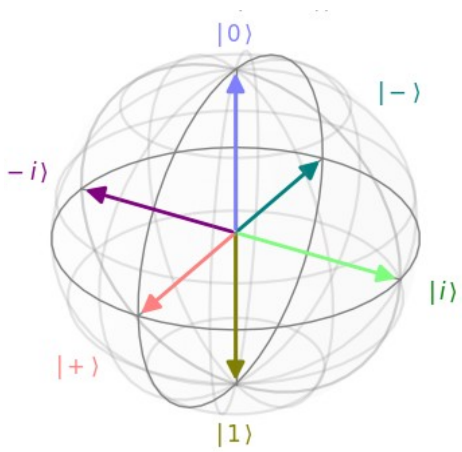

Reference

2

directions

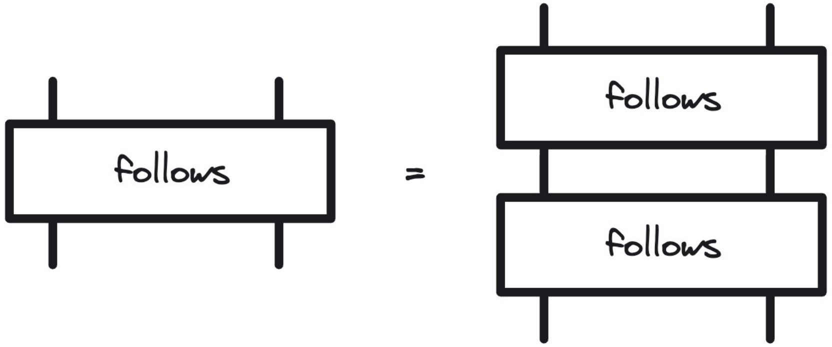

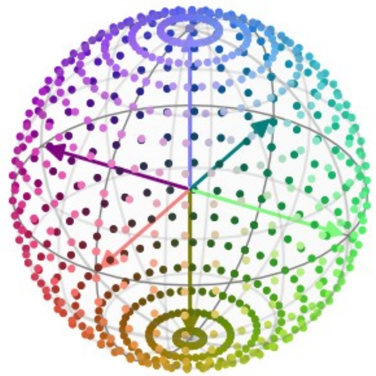

Reference

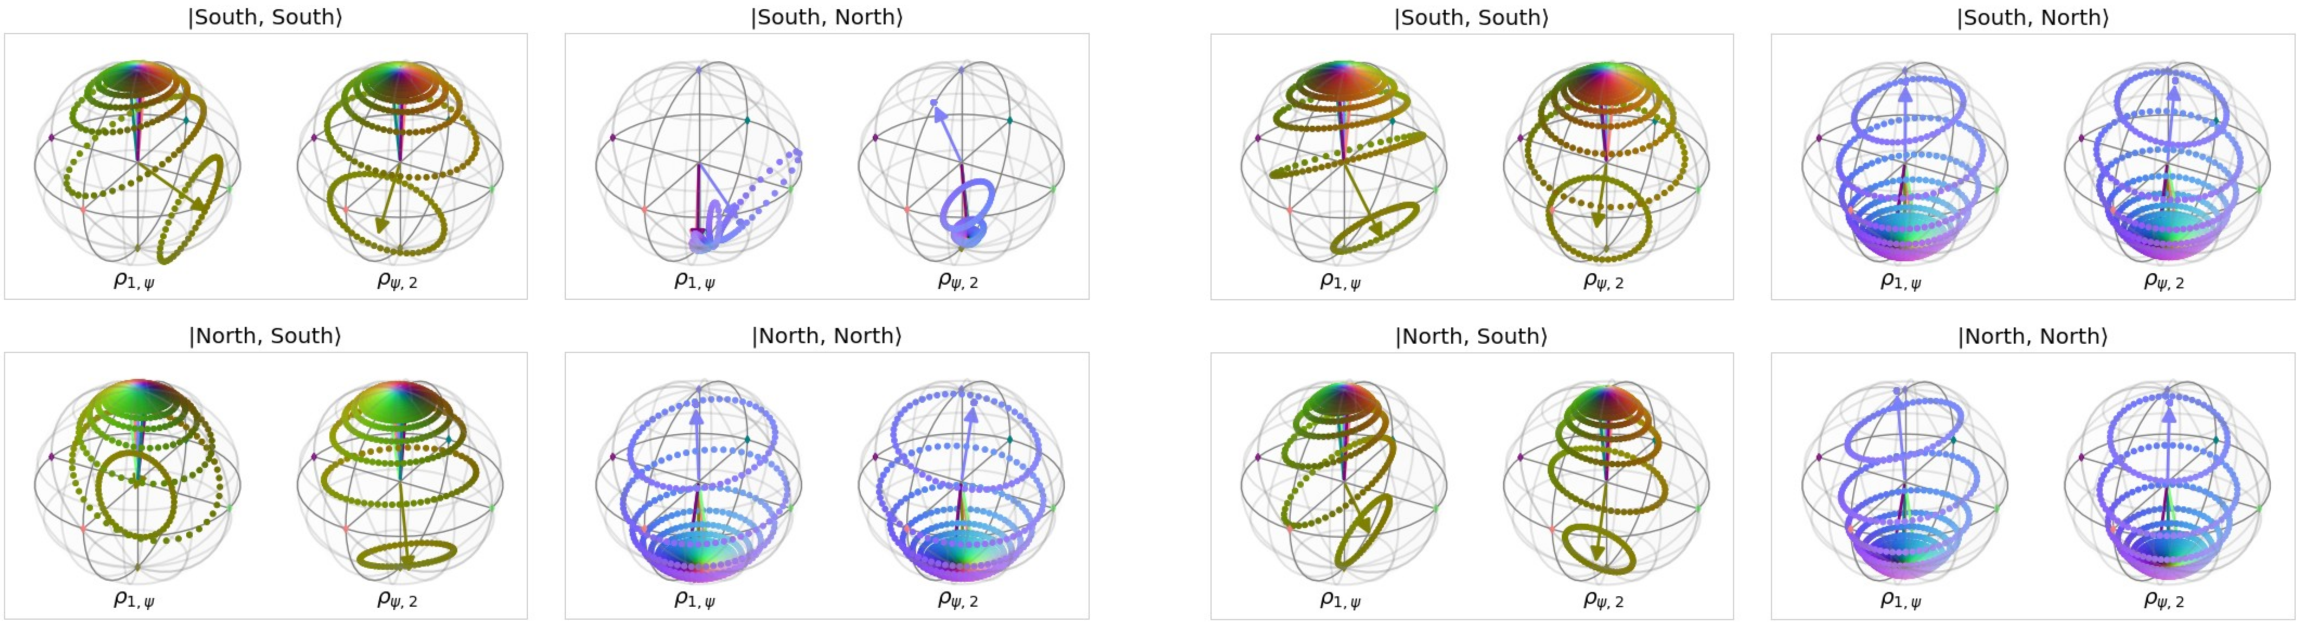

4

directions

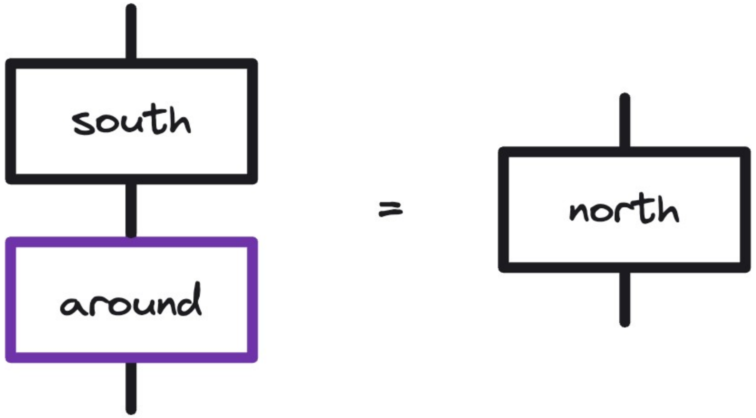

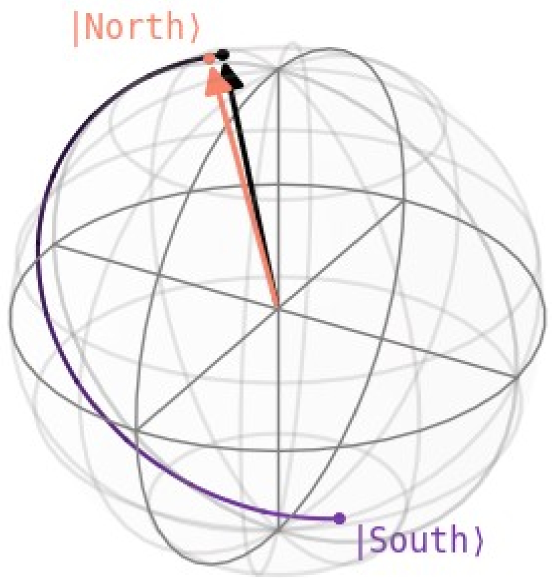

4

**direction
s**

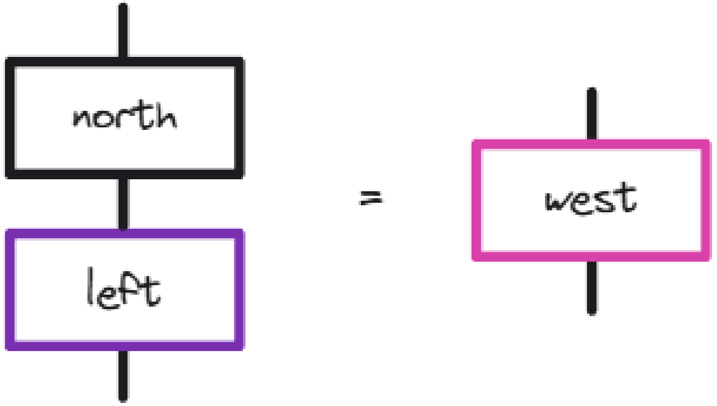

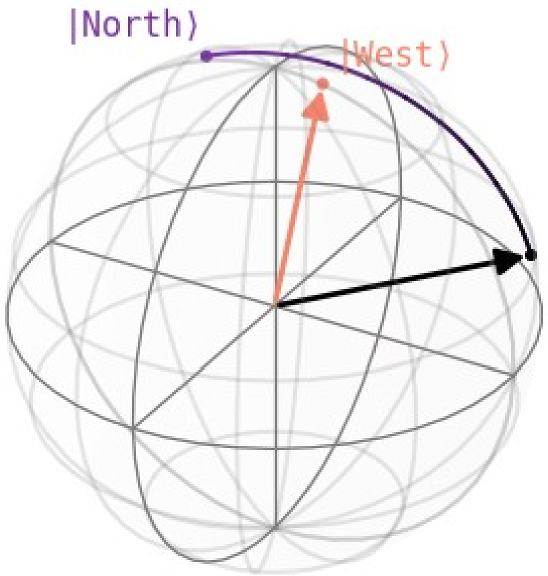

4

**direction
s**

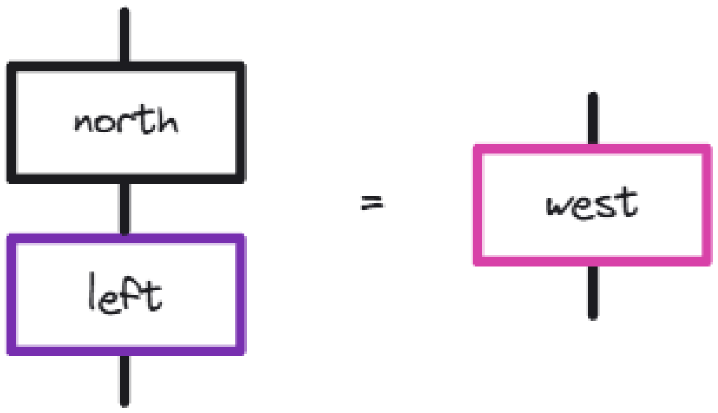

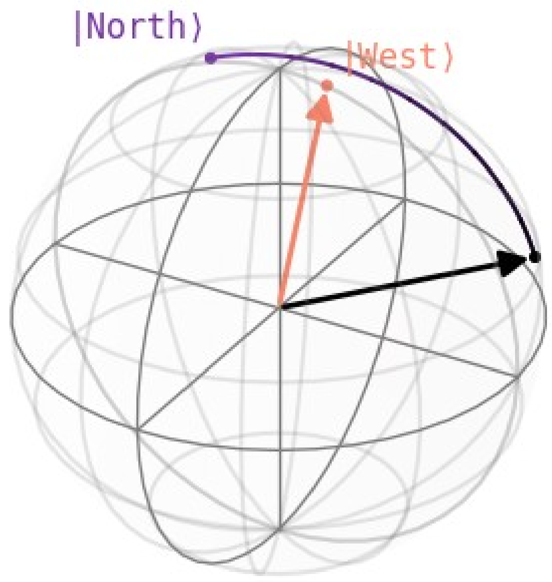

4 directions

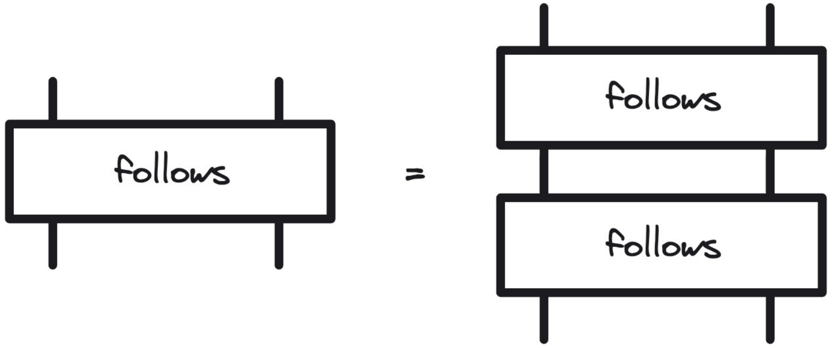

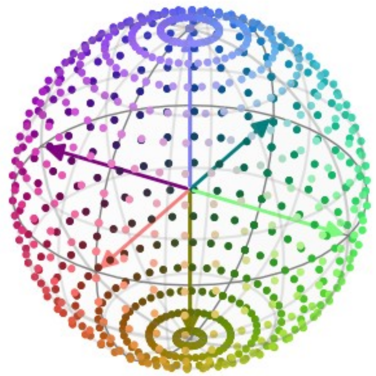

Reference

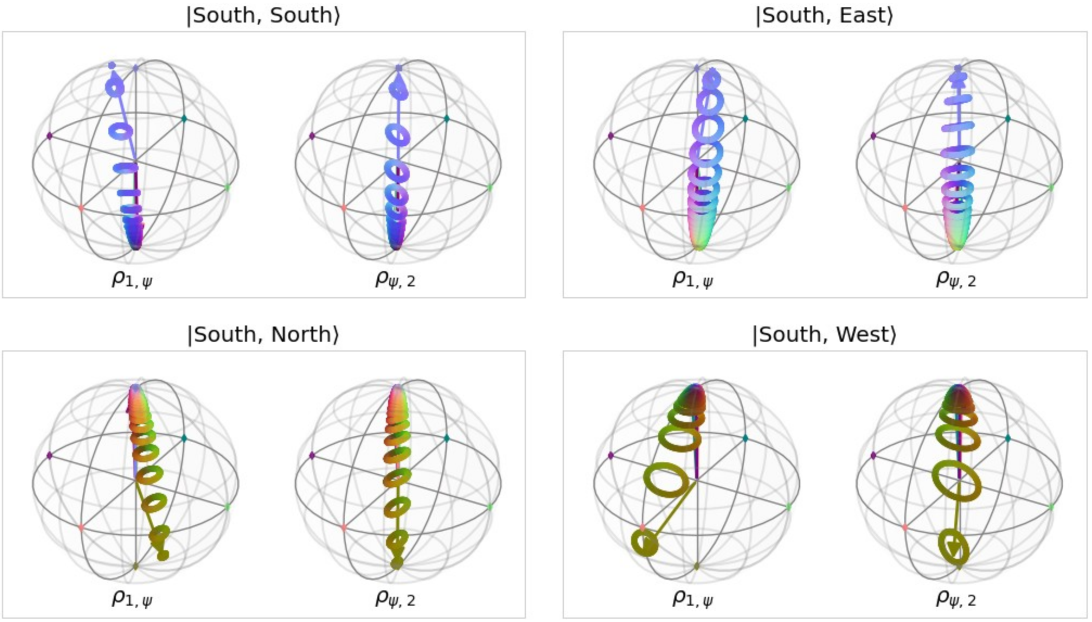

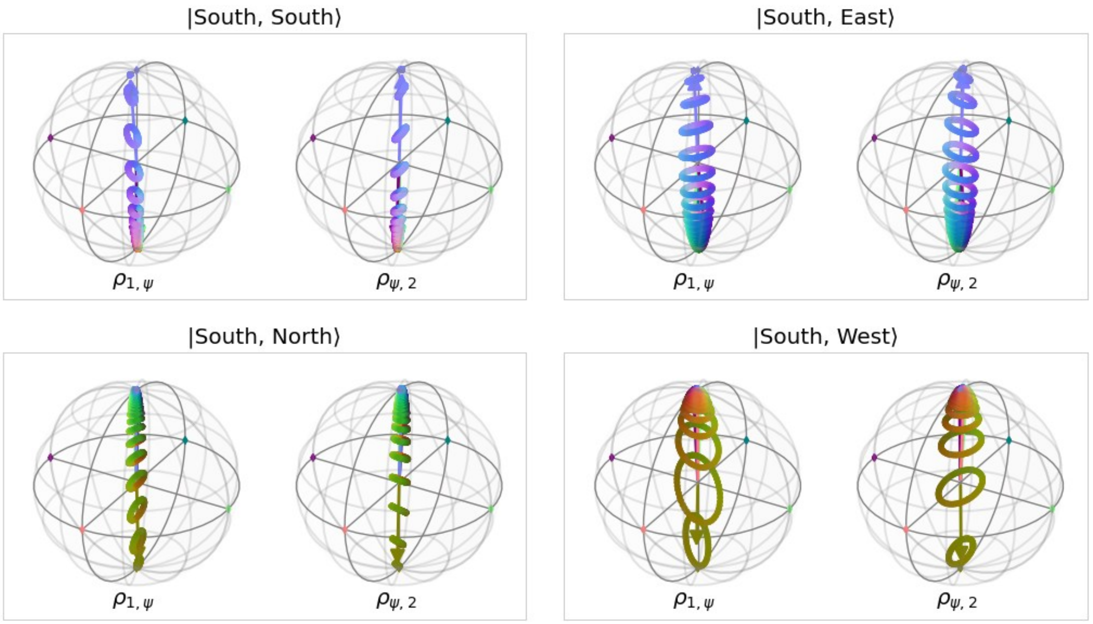

4

directions

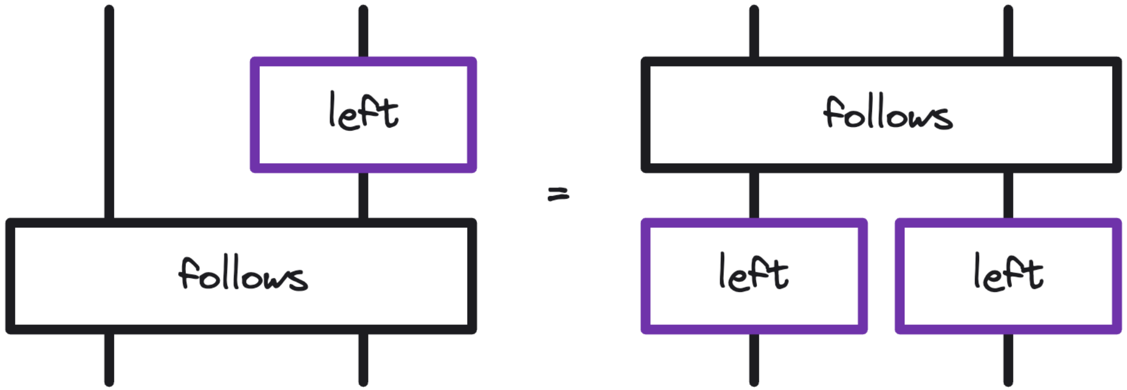

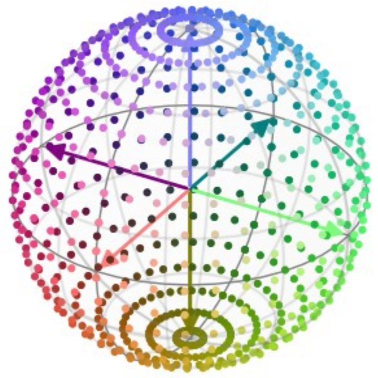

Reference

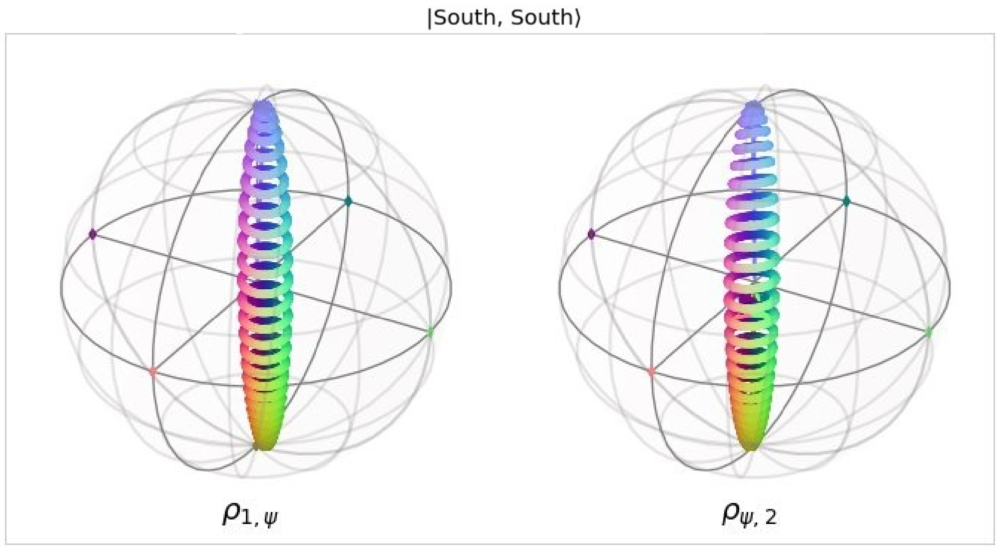

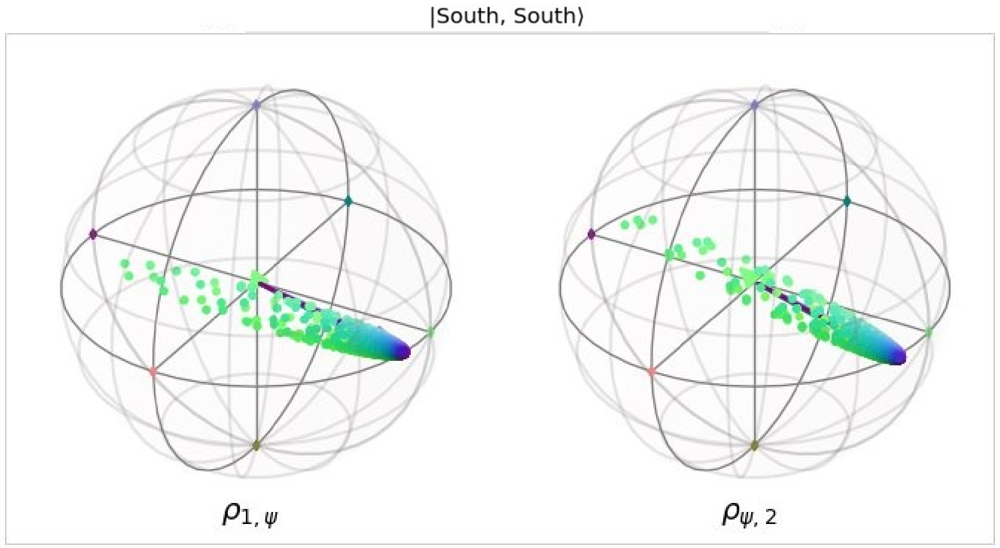

2
directions

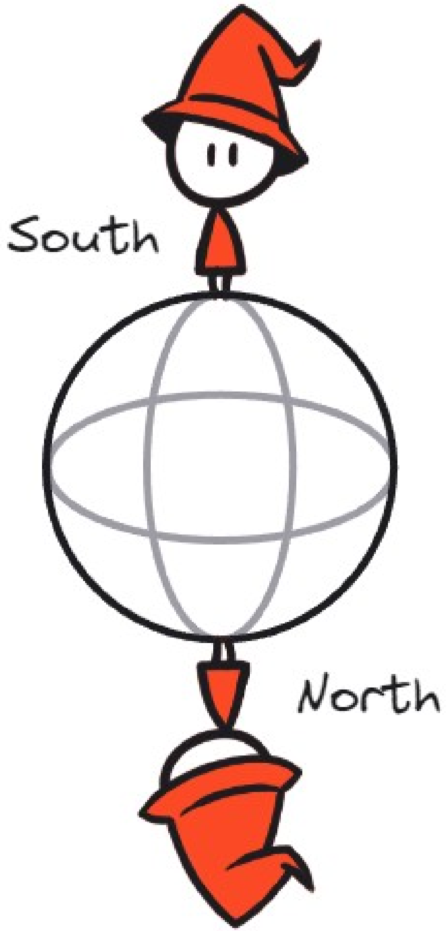

4
directions

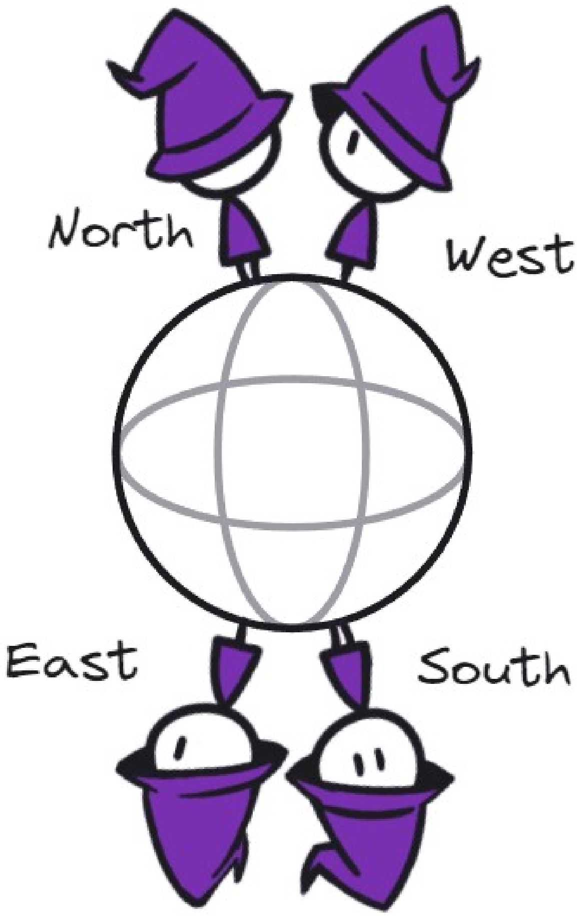

4
directions

Compositional generalisation

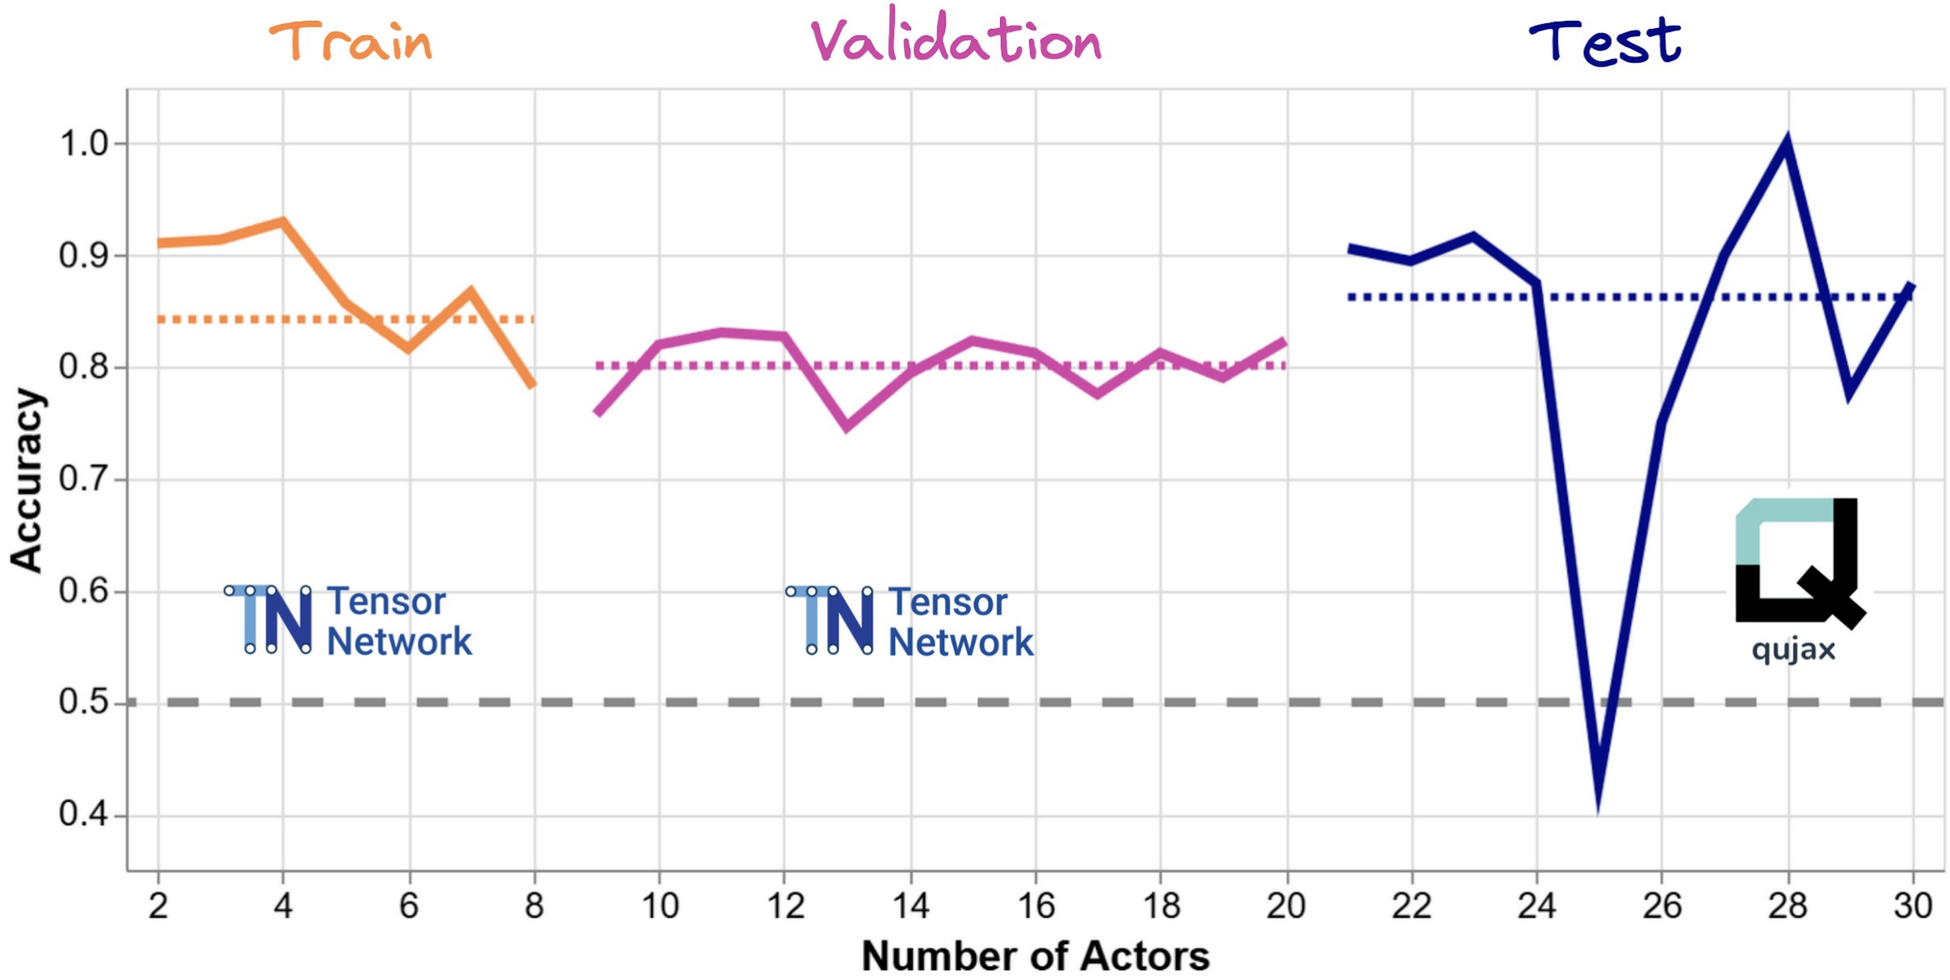

Classical Baselines

Accuracy Summaries

2

directions

4

directions

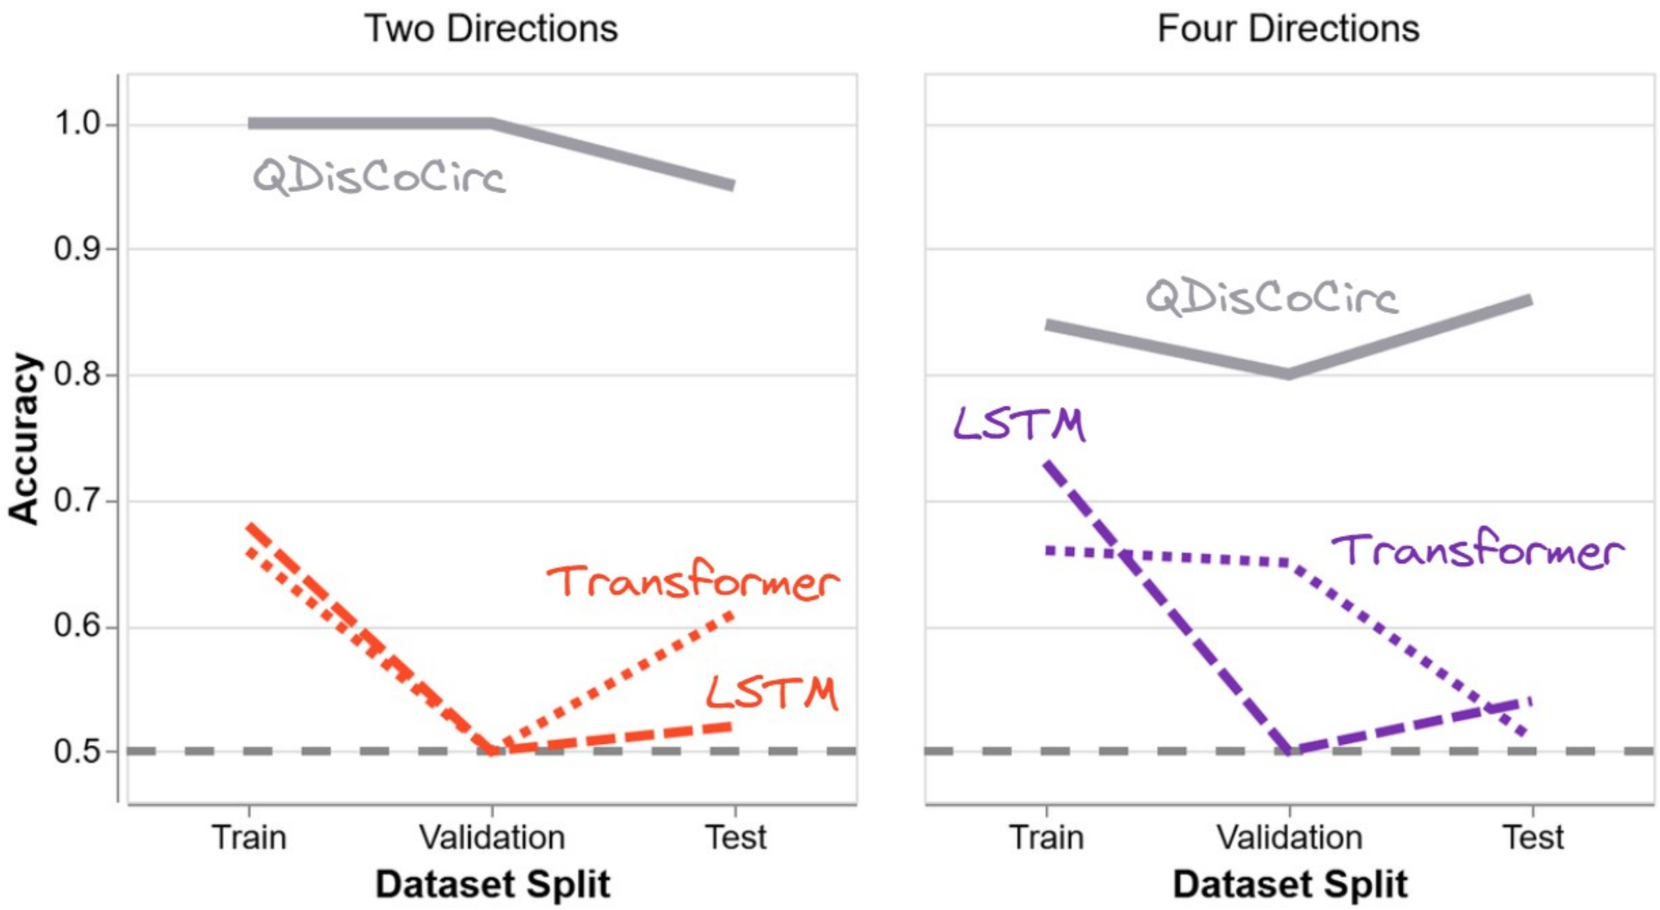

GPT-4

2
directions

4
directions

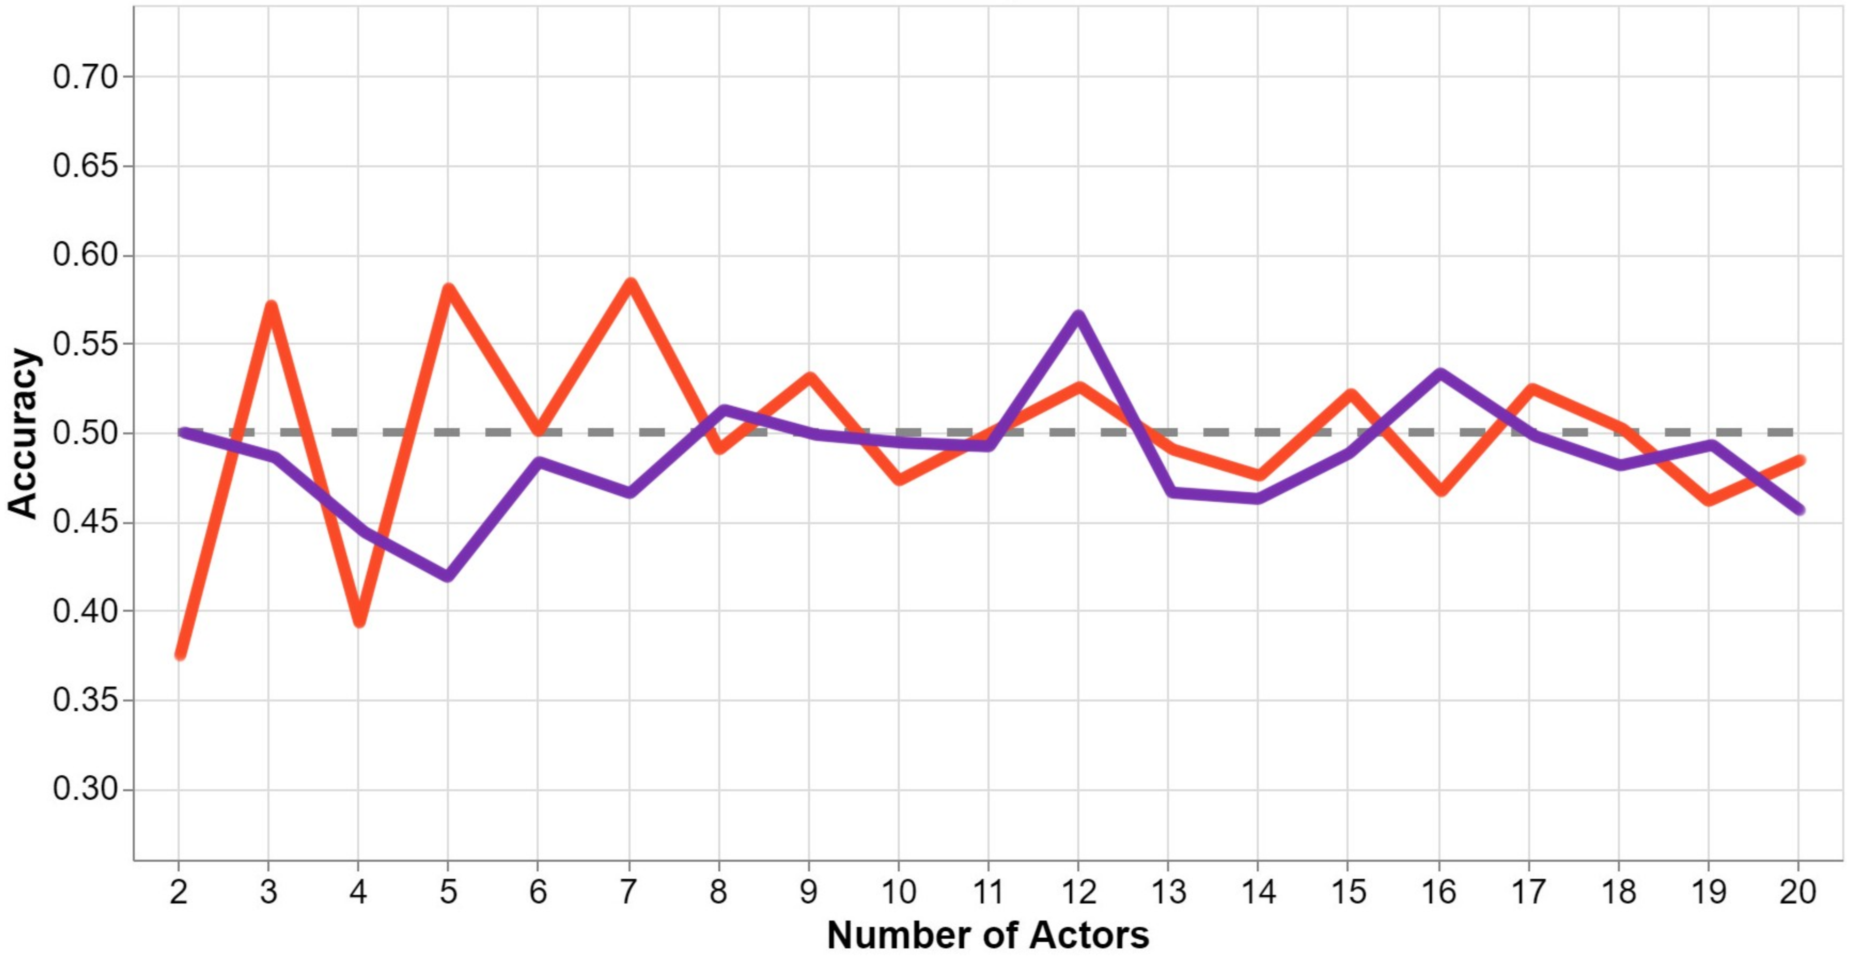

Accessed on 25th April 2024.

Summary

Compositional...

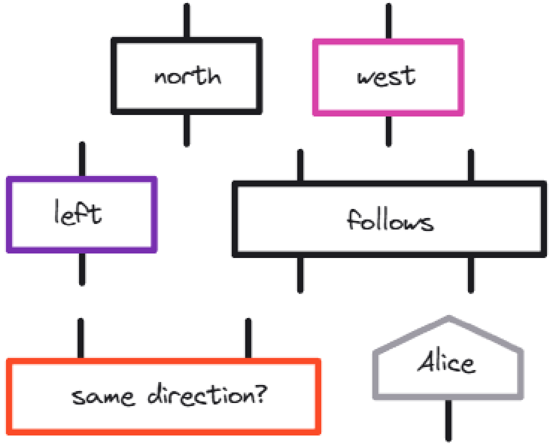

Model

Compositional...

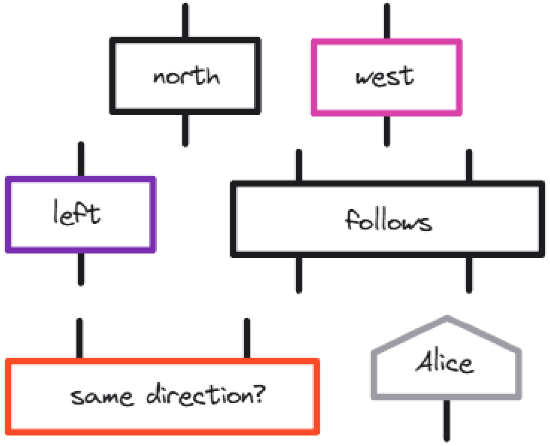

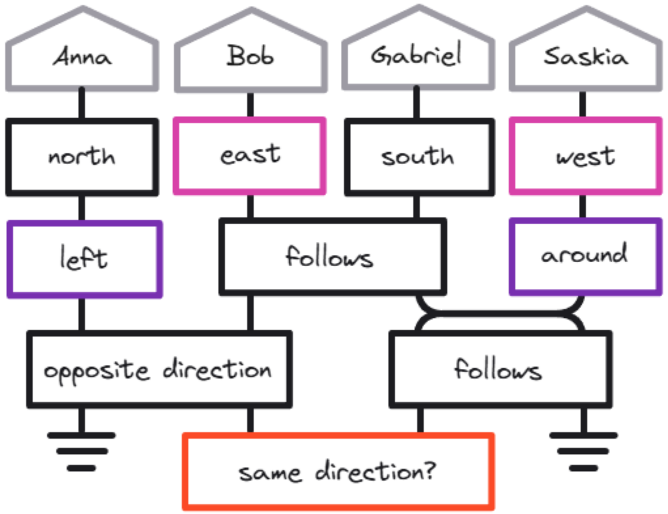

Model ----> Generalisation

Compositional...

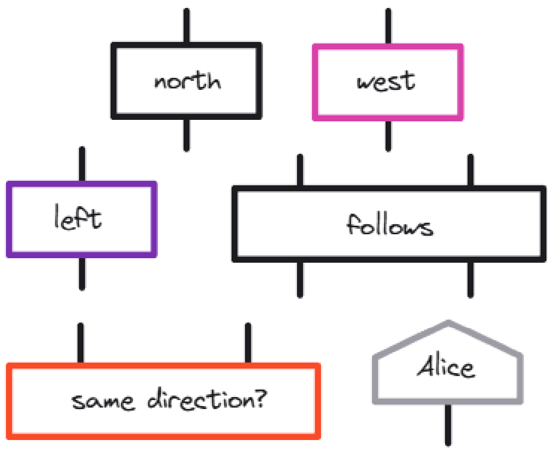

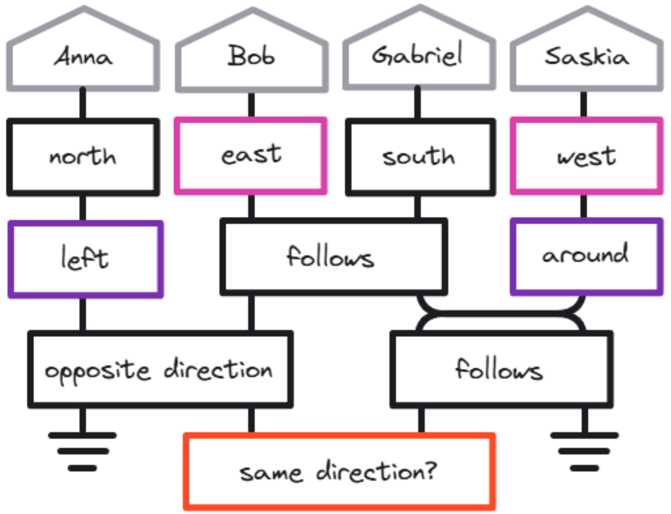

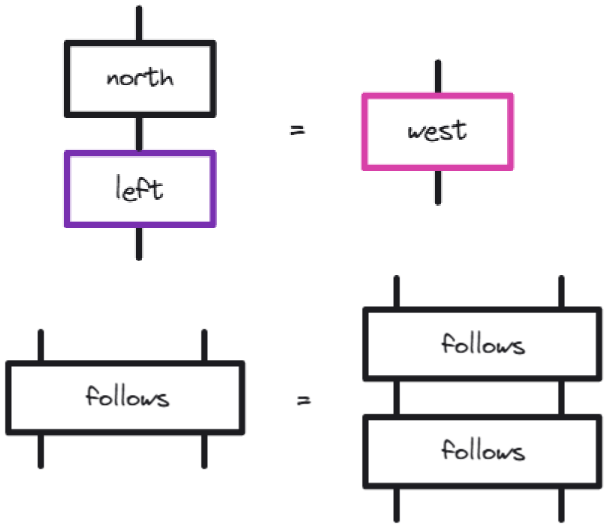

Model ----> Generalisation + Interpretability

LAMBEQ

Natural Language Processing on Quantum Computers

λambeq Gen II:

A Quantum-Enhanced Interpretable and Scalable Text-based NLP Software Package

QUANTINUUM

arXiv > quant-ph > arXiv:2505.13208

Search..

Help | Ac

Quantum Physics

[Submitted on 19 May 2025]

Efficient Generation of Parameterised Quantum Circuits from Large Texts

Colin Krawchuk, Nikhil Khatri, Neil John Ortega, Dimitri Kartsaklis

Quantum approaches to natural language processing (NLP) are redefining how linguistic information is represented and processed. While traditional hybrid quantum-classical models rely heavily on classical neural networks, recent advancements propose a novel framework, DisCoCirc, capable of directly encoding entire documents as parameterised quantum circuits (PQCs), besides enjoying some additional interpretability and compositionality benefits. Following these ideas, this paper introduces an efficient methodology for converting large-scale texts into quantum circuits using tree-like representations of pregroup diagrams. Exploiting the compositional parallels between language and quantum mechanics, grounded in symmetric monoidal categories, our approach enables faithful and efficient encoding of syntactic and discourse relationships in long and complex texts (up to 6410 words in our experiments) to quantum circuits. The developed system is provided to the community as part of the augmented open-source quantum NLP package lambeq Gen II.

Subjects: **Quantum Physics (quant-ph)**; Artificial Intelligence (cs.AI); Computation and Language (cs.CL)

Cite as: arXiv:2505.13208 **[quant-ph]**

(or arXiv:2505.13208v1 **[quant-ph]** for this version)

https://doi.org/10.48550/arXiv.2505.13208

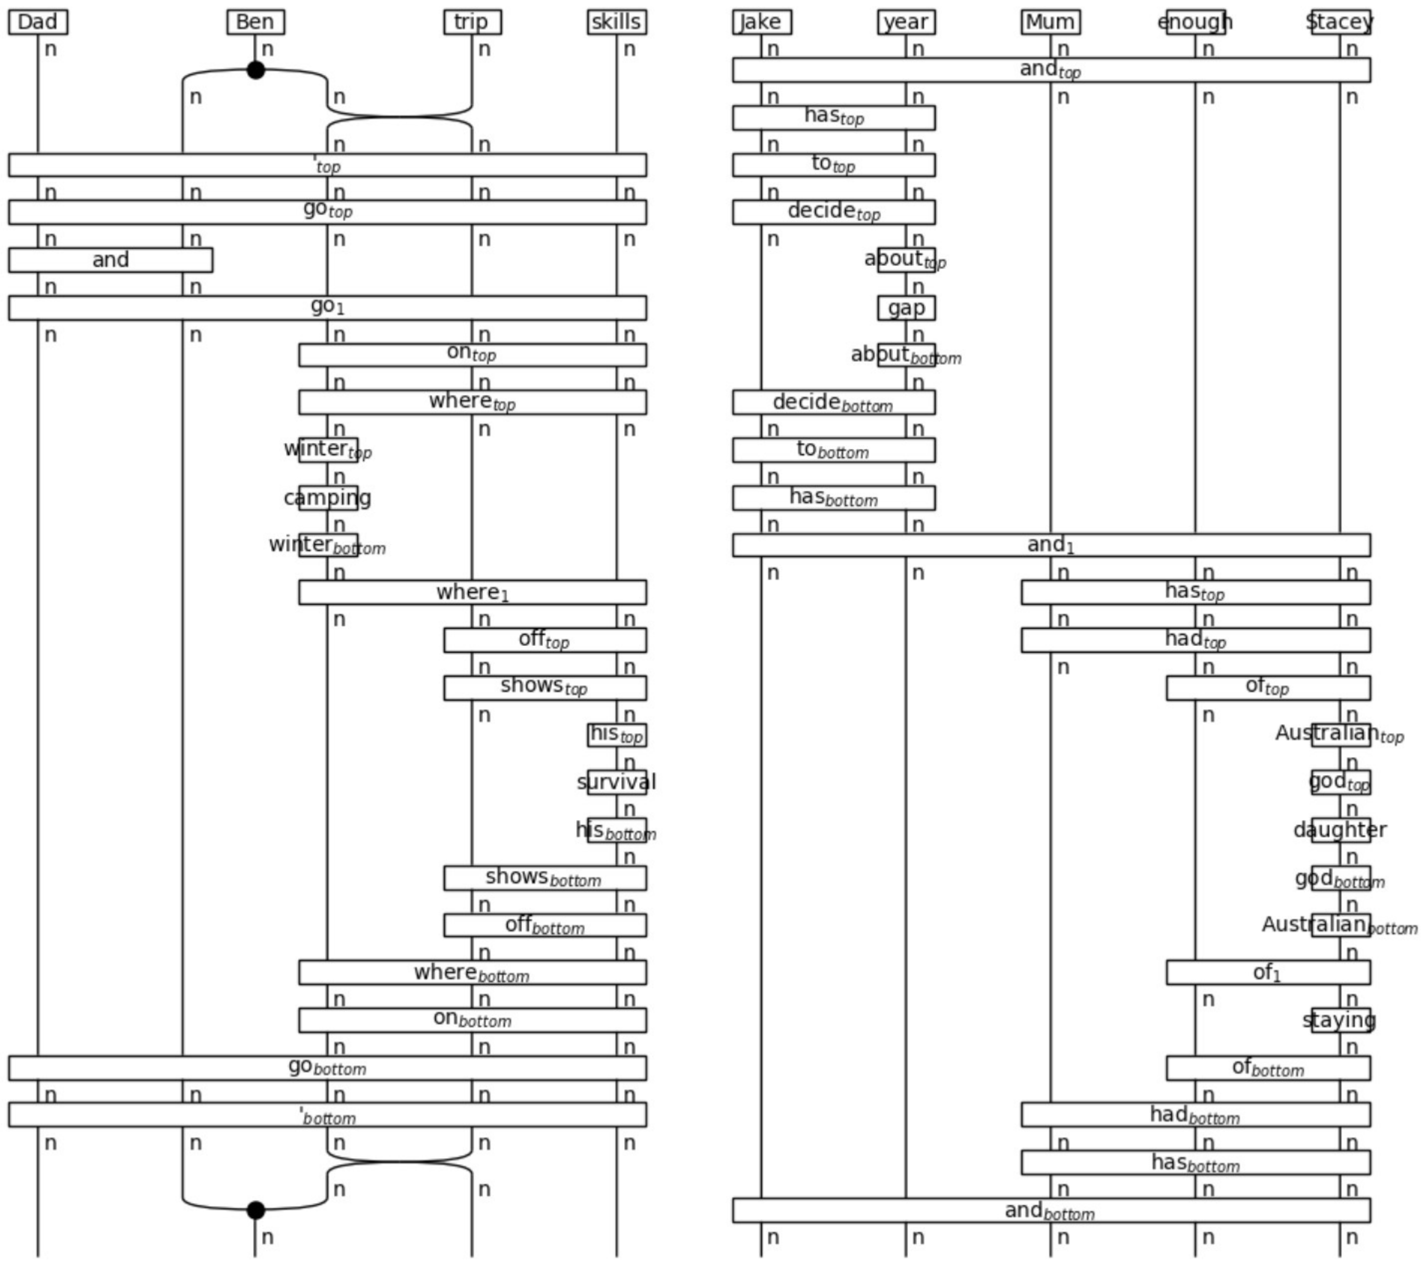

Real Data

"Dad and Ben go on a winter camping trip, where Ben shows off his survival skills. Jake has to decide about a gap year, and Mum has had enough of Australian god-daughter Stacey staying."

- BBC Synopsis Dataset

Try it yourself!

```
> pip install lambeq
```

DisCoCirc
2301.10595

Interpretability
2406.17583

I. Theory
2408.06061

II. Experiment
2409.08777

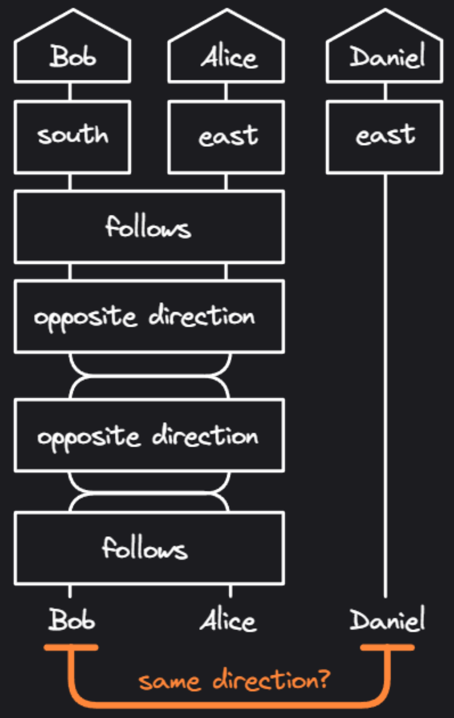

Blog post
https://www.quantinuum.com/blog/talking-quantum-circuits

References

- DisCoCirc Background

https://arxiv.org/abs/2301.10595

- Theory paper

https://arxiv.org/abs/2408.06061

- Interpretability paper

https://arxiv.org/abs/2406.17583

- Experimental paper (this work)

https://arxiv.org/abs/2409.08777

- Blog post

https://www.quantinuum.com/blog/talking-quantum-circuits

- Lambeq release - DisCoCirc parser out now!

https://docs.quantinuum.com/lambeq/

- H-series hardware

https://www.quantinuum.com/products-solutions/quantinuum-systems

- Qujax

https://github.com/CQCL/qujax

- Nexus

https://docs.quantinuum.com/nexus/

Links

Training DisCoCirc

https://github.com/CQCL/lambeq

Circuit Compilation

https://docs.quantinuum.com/tket/

Circuit Simulation

https://github.com/CQCL/qujax

Thank you for your attention!

Muhammad Hamza Waseem
DPhil Physics, University of Oxford
Research Scientist at Quantinuum

Mathematics and Physics of Quantum Computing and Quantum Learning
Porquerolles, France

May 23-28, 2025

QUANTINUUM

UNIVERSITY OF OXFORD