Quantum Theory in Knowledge Representation

Reasoning with a Quantum Model of Concepts

Ward Gauderis
Prof. Dr. Dr. Geraint Wiggins

November 19, 2024

Computational Creativity Lab, Artificial Intelligence Research Group, Vrije Universiteit Brussel

Quantum Theory

Quantum theory

Mathematical framework at the core of quantum mechanics, independent of physical interpretation

Since inception in early 20th century:

- Successfully describes and predicts behaviour of subatomic particles
- Predicts **Quantum weirdness** contrary to classical theories and everyday physical experience
  - Non-locality, contextuality, entanglement, superposition, incompatible measurements, ...

Attitude shift since 1980s:

- **Features, not bugs!** How can we use them?
  - Quantum information theory, computing, cryptography, machine learning, ...

General Probabilistic Theory

Quantum theory is a **general probabilistic theory**

- Slightly different axiomatisation from classical probability theory (Hardy 2001)
- Quantum two-norm vs. classical one-norm probability
- Geometric interpretation of probability as the length of projections onto subspaces

Utilitarian modelling beyond the domain of physics

- Underlying processes are not inherently quantum, but share mathematical structure
    - Non-determinism, non-separability, invasive measurements, contextuality, superposition, ...
- Quantum **modelling advantage** $\rightarrow$ quantum **computational advantage**

Linear algebra and probability theory are widespread in artificial intelligence

Quantum game theory

- Quantum foundations
- Reinforcement learning

Generalised satisfiability

- Relaxed SAT
- Hamiltonian complexity

Tensor networks

- Numerical simulation
- Machine learning

Quantum NLP

- Language modelling
- Information retrieval

Quantum cognition

- Cognitive science
- Cognitive modelling

Quantum Picturialism

> ***Quantum picturialism*** *refers to the use of diagrams to represent and reason about essential features of quantum theory. It aims to describe the logic of interacting quantum processes, such that diagrammatic equations become the very foundation of quantum theory.* - Coecke and Kissinger (2018)

Hilbert space formalism:

- Low-level and reductionist
- Isolated systems and their state
- Highlights deviations from classical theory

Diagrammatic language:

- High-level and constructivist
- Composite processes and their interaction
- Highlights features of quantum theory

Founded in a **categorical quantum mechanics**:

- Rigorous mathematical foundation in **symmetric monoidal categories**
- Emphasises connection to other types of systems and processes

String Diagrams

Symmetric Monoidal Category Process Theory	Objects System-types	Morphisms Processes
Relations	Sets	Relations
Linear maps	Vector spaces	Linear maps
Classical probability	Measurable spaces	Markov kernels
Quantum maps	(operators on) Hilbert spaces	Completely positive maps

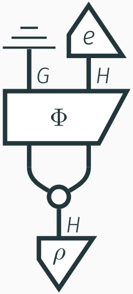

$$= \quad (\mathrm{Tr}_G \otimes e) \circ \Phi \circ \delta_H \circ \rho \quad \in \mathbb{C}$$

$$\rho \in \mathcal{B}(H)$$
$$\delta(\rho) = \sum_{ij} |ii\rangle \langle i| \rho |j\rangle \langle jj|$$
$$\Phi : \mathcal{B}(H \otimes H) \to \mathcal{B}(G \otimes H)$$
$$\mathrm{Tr}_G : \mathcal{B}(G) \to \mathbb{C}$$
$$e : \mathcal{B}(H) \to \mathbb{C}$$

String diagram interpreted in the category of quantum maps

Quantum Model of Concepts (Tull et al. 2023)

Cognitive science:

*"**Concepts** are the glue that holds our mental world together."* - G. Murphy

- Essential to cognitive processes such as reasoning, decision-making, perception, language, ...

But how to represent concepts?

Artificial intelligence:

- Create AI agents that reason and act more effectively, similar to how humans use concepts
- Ameliorate negative consequence of black-box connectionist models

How to automatically learn and reason with concepts?

Conceptual Space Theory

How to model cognitive representations?

Symbolic approach: High-level

- Representations express propositional relations between discrete objects
- Cognition is computation at the level of symbols

\+ Compositional aspects of cognition

Conceptual spaces: Intermediate (Gärdenfors 2000)

- Instrumentalist level of representation

\+ Bridge between symbolic and subsymbolic approaches

Subsymbolic approach: Low-level

- Associations between types information elements are the centre of representation
- Computation is a consequence of developing representations

\+ Fine-grained similarity between representations

Quantum Conceptual Model

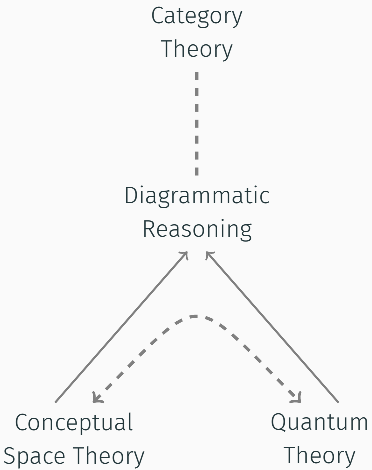

Convex conceptual spaces ⟶ Diagrammatic conceptual models ⟶ Quantum conceptual models

(Tull et al. 2023)

Quantum Conceptual Model

Hilbert space $H \subseteq H_1 \otimes \cdots \otimes H_n$ and isometry U s.t.

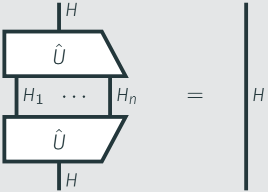

Quantum Instance

Pure normalised quantum state $\langle h|$

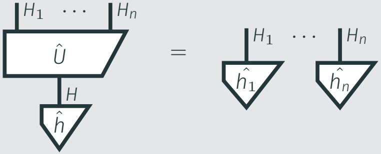

Quantum Concept

Quantum effect c

Concept testing = Quantum measurement (Born rule)

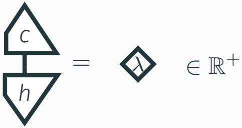

Semantic conceptual properties ↔ measurable quantum properties

- **Partial order** on concepts ↔ partial order on quantum processes
- **Pure concepts** at the bottom of the order ↔ pure quantum states
- **Prototypical instances** of concepts ↔ eigenstates of a quantum measurement
- Fuzzy, crisp, product, separable, … concepts

Entangled Quantum Concepts

Given a set of pure concepts $\langle c_1| , \ldots , \langle c_n|$, how can they be combined?

Classical combinations

$$c = |c_1\rangle \langle c_1| + \cdots + |c_1\rangle \langle c_n|$$

- Separable → No generalisation: c compares to each c_i individually

Quantum combinations

$$c = (|c_1\rangle + \cdots + |c_n\rangle) \otimes (\langle c_1| + \cdots + \langle c_n|)$$

- Entangled → Generalisation: c captures **structural relations between domains**
  - Any quantum map $f : H \to G$ can be captured by a quantum concept c

Reasoning with a Quantum Model of Concepts

Symbolic and subsymbolic representations are complementary

Symbolic models:

+ Compositional
+ Human-interpretable
+ Generalise through reuse
– Hand-crafted
– Grounding problem
– Exhaustive combinatorial search

Subsymbolic models:

– Binding problem
– Uninterpretable
– Limited generalisation
+ Learnable from raw data
+ Grounded in data
+ Robust to noise

Can quantum conceptual models serve as practical intermediate representations for agents that use both symbolic and subsymbolic reasoning?

Hybrid Quantum-Classical Variational Circuit

1. Classical preprocessing
2. Quantum state preparation
3. Measurement and post-processing
4. Classical optimisation

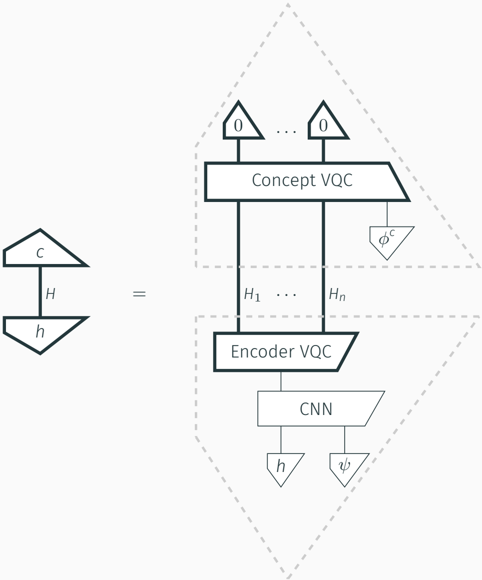

Practical Blueprint for Quantum Concepts

Problem

On a quantum computer, only sub-causal effects can be realised as branches of causal non-deterministic processes implemented by pure unitary maps

Solution

1. Scale c with $r \in \mathbb{R}^+$ such that $c' = rc$ is sub-causal
2. Embed c' as branch 0 of a demolition POVM measurement
3. Apply the Ozawa dilation to the process to obtain an ONB measurement
4. Transform the resulting isometry U into a unitary U'
5. Postselect on the outcome 0 in the ONB measurement

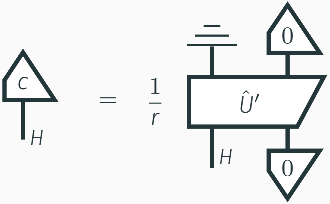

Experiment 1: Shapes Dataset

$$H \subseteq Colour \otimes Position \otimes Shape \otimes Size$$

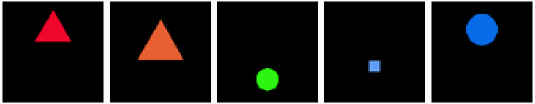

Learn instance and property representations with meaningful similarity

- Contrastive self-supervised learning of cognitively separable domains using BCE Loss

$$\mathcal{L}(\phi,\psi) = -\frac{1}{d}\sum_{i=1}^{d}\sum_{j=1}^{n}\left[y_{ij}\log c_{ij}(h_i) + (1-y_{ij})\log(1-c_{ij}(h_i))\right]$$

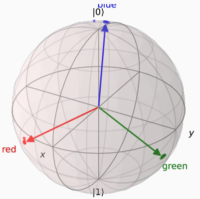

Colour

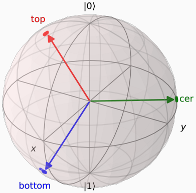

Position

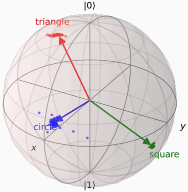

Shape

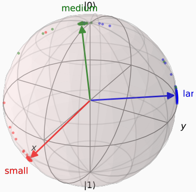

Size

Experiment 2: Rainbow Dataset

Property packing density

- On n qubits, 2^n orthogonal properties can be distinguished
  - Single ONB measurement
- More than 2^n properties cannot be orthogonal
  - Repeated POVM measurements

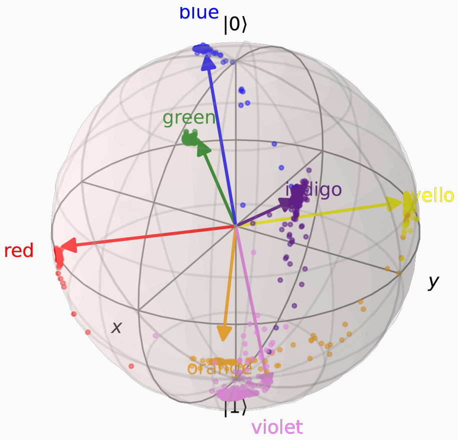

Colour

Experiment 3: Decoder Loss

Retain variational information within representations

- Decoder network reconstructs instances with an unsupervised penalty

$$\mathcal{L}(\phi, \psi, \pi) = \mathcal{L}(\phi, \psi) + \frac{\lambda}{d \cdot 3 \cdot 64^2} \sum_{i=1}^{d} ||h_i - \text{TransCNN}_{\pi}(\text{CNN}_{\psi}(x_i))||^2$$

Reconstructions of prototypical instances

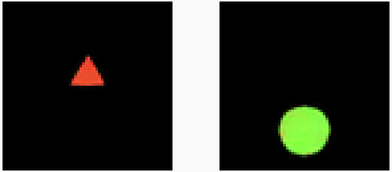

(red, centre, triangle, small) (green, bottom, circle, large)

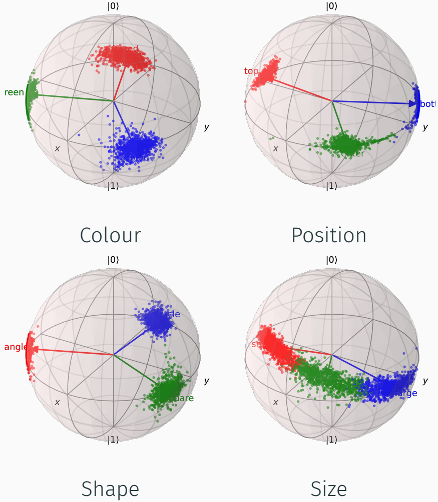

Experiments 4-5-6: Learning Concepts

Learn concept representations with meaningful similarity on (frozen) domains:

- Supervised learning with class imbalance and BCE loss
- Interpretable conceptual properties ↔ quantum circuit properties

Experiment 4: correlated concepts

- 100% accuracy with entangled concepts

Experiment 5: general concepts

- 100% accuracy with mixed concepts and discarding (partial trace)

Experiment 6: logic operators in concepts

- 100% accuracy on conjunction and disjunction within and across domains

General concept:

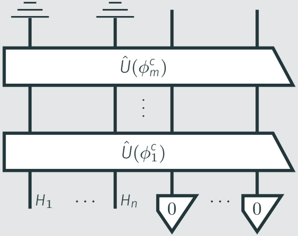

Beyond Concept Recognition

Can the compositional features of the quantum concepts be used to solve abstract reasoning problems with perceptual uncertainty?

Blackbird datasets

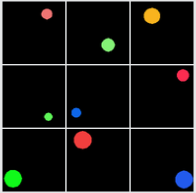

- Synthetic puzzles inspired by Raven's Progressive Matrices
  - Used by Hersche et al. (2023) to demonstrate vector-symbolic reasoning
- Complete missing panels in a 3x3 grid of abstract shapes
  - Noisy variation in 2 continuous domains
  - textitcolumns and *row* constraints

Experiment 7: Quantum Conceptual Model of Puzzles

$$G = \bigotimes_{i=1}^{3} \bigotimes_{j=1}^{3} H_{ij} \quad \text{with } H \subseteq \textit{colour} \otimes \textit{position}$$

Learn factorised models

1. Learn domain representations of H
2. Learn row and column concepts
    - 100% accuracy with general concepts

Learning from prototypical instances

- Replace training set with prototypes
- Mimicks human learning from idealised cases

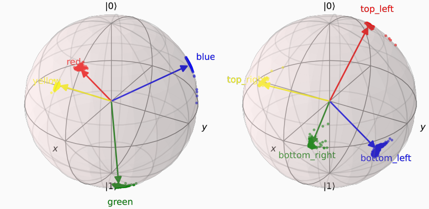

Without decoder loss

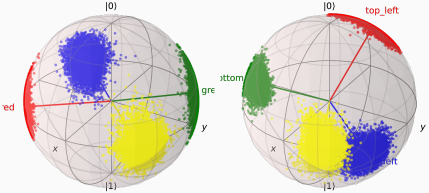

With decoder loss

Experiment 8: Composition of Quantum Concepts

Composition of quantum concepts

- String diagrams capture shared structures in Boolean relations and quantum processes
- Similar to logic programs, complex concepts are composed by reusing sub-concepts

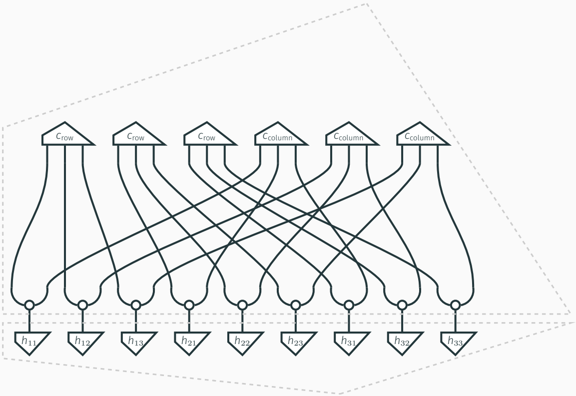

Compose *puzzle* concept from *row* and *column* concepts

- 100% concept classification accuracy

Experiment 9: Quantum Concepts as Generative Processes

Quantum concepts as generative processes

- Quantum concepts encode joint probability distributions
  - Conditioning by process-state duality
  - Marginalisation by discarding (partial trace)
- Quantum conceptual processes enable generative instance sampling

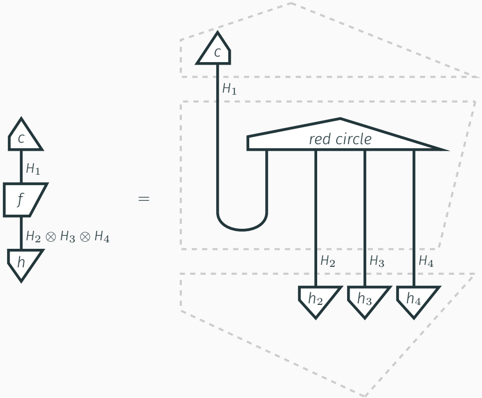

Predicting the colour of an instance from the *red circle* concept

- 100% prediction accuracy
- Marginal probability ∼ concept frequency
- Conditional probability ∼ structural relations

	$P(\text{red circle})$	$P(\text{not red circle})$	
$P(\text{red})$	0.31	0.19	0.50
$P(\text{not red})$	0.01	0.49	0.50
	0.32	0.68	1.00

$P(\text{red} \mid \text{red circle}) = 0.98$

Experiment 10: Reasoning with Quantum Concepts

Reasoning with quantum concepts

1. Automatically compose and apply a quantum conceptual process to an incomplete puzzle
2. The prepared quantum state encodes a joint probability distribution over missing panels
3. Sample and predict the most likely panels

Solve `blackbird` puzzles with quantum concepts

- 100% prediction accuracy
- Tested on NISQ `ibm_kyiv` hardware

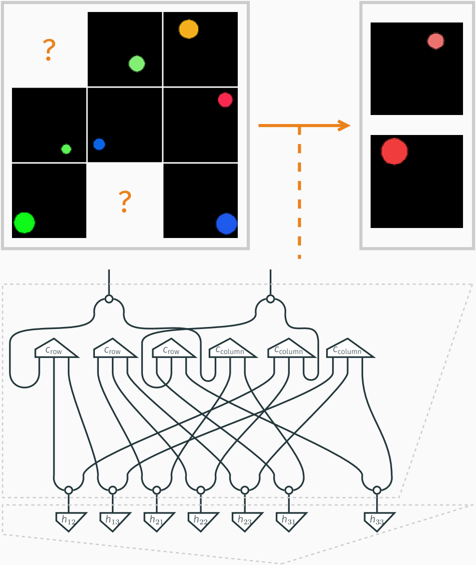

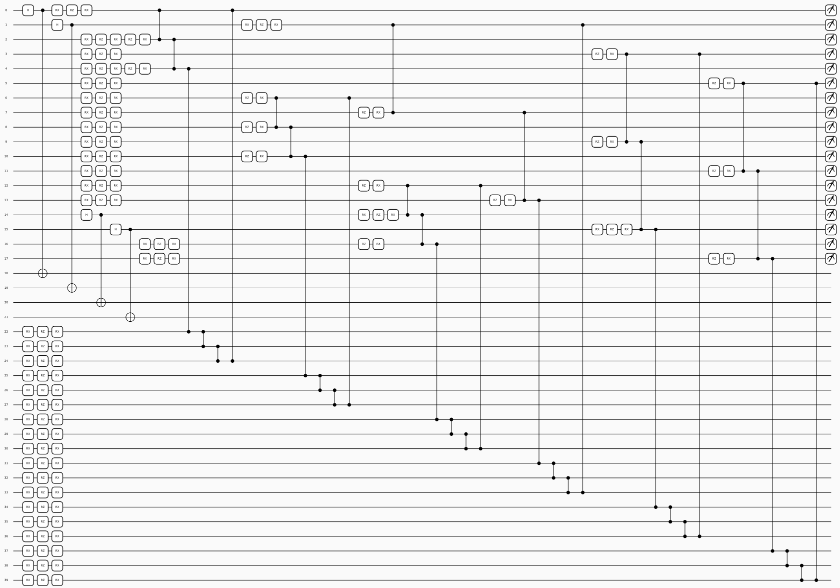

Simplified compiled generative concept circuit of the *puzzle* concept

Conclusion

Quantum theory is a **general probabilistic theory** beyond physics

- Quantum picturalism emphasises its compositional features and relates them to other theories, leading to applications in cognitive science and AI

Quantum conceptual models unite quantum theory and conceptual space theory

Quantum concepts are generative intermediate representations capable of solving abstract reasoning problems

Symbolic

- Compositional grounding
- Human-interpretable
- Generalise through reuse

Subsymbolic

- Grounded in perceptual data
- Learnable from raw data
- Robust to variation and uncertainty

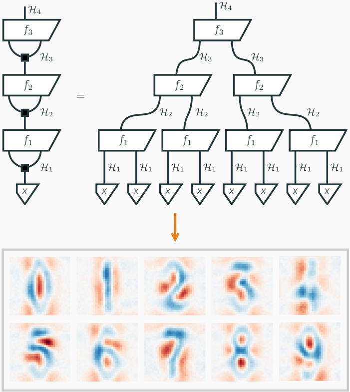

Ongoing work with Thomas Dooms

- Extending the study of compositionality to ML
- Compositionally-Interpretable Tensor Neural Networks
    - Linear tensor networks ∩ non-linear neural networks
    - Quantum-compositional ∩ mechanistic interpretability

Find out more

Quantum conceptual models + datasets

- `github.com/WardGauderis/Quantum-Conceptual-Model`

Ward Gauderis

- `ward.gauderis@vub.be`
- Artificial Intelligence Research Group
  Vrije Universiteit Brussel

References

- Coecke, Bob and Aleks Kissinger (2018). **‘Picturing Quantum Processes’.** In: *Diagrammatic Representation and Inference*. Ed. by Peter Chapman et al. Lecture Notes in Computer Science. Cham: Springer International Publishing, pp. 28–31. ISBN: 978-3-319-91376-6. DOI: 10.1007/978-3-319-91376-6_6.
- Gärdenfors, Peter (2000). ***Conceptual Spaces: The Geometry of Thought.*** Cambridge, Mass: MIT Press. 307 pp. ISBN: 978-0-262-07199-4.
- Hardy, Lucien (Sept. 25, 2001). ***Quantum Theory From Five Reasonable Axioms.*** arXiv: quant-ph/0101012. URL: http://arxiv.org/abs/quant-ph/0101012 (visited on 03/28/2023). preprint.
- Hersche, Michael et al. (Mar. 3, 2023). ***A Neuro-vector-symbolic Architecture for Solving Raven’s Progressive Matrices.*** DOI: 10.48550/arXiv.2203.04571. arXiv: 2203.04571 [cs]. URL: http://arxiv.org/abs/2203.04571 (visited on 03/31/2023). preprint.
- Tull, Sean et al. (Feb. 7, 2023). ***Formalising and Learning a Quantum Model of Concepts.*** DOI: 10.48550/arXiv.2302.14822. arXiv: 2302.14822 [quant-ph, q-bio]. URL: http://arxiv.org/abs/2302.14822 (visited on 03/07/2023). preprint.