Contextual Semantics: From Quantum Mechanics to Logic, Databases, Constraints, Complexity and Natural Language Semantics

Samson Abramsky

Department of Computer Science, University of Oxford

Johan

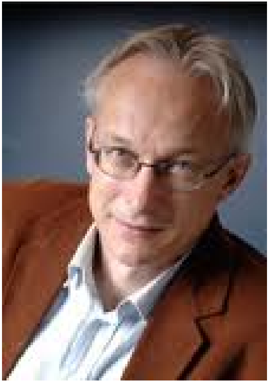

Johan

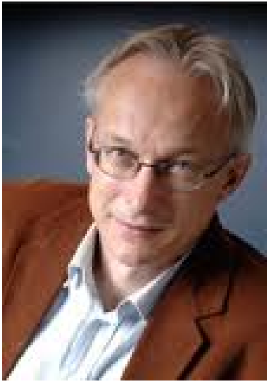

Thank you for encouraging me to broaden my intellectual horizons and to address wider conceptual issues.

Johan

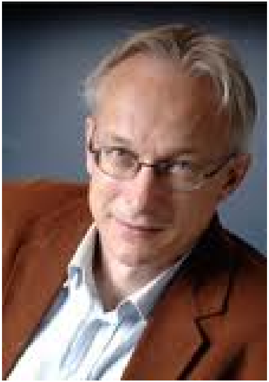

Thank you for encouraging me to broaden my intellectual horizons and to address wider conceptual issues.

I thought I was out, but Johan pulled me back in . . .

Beginnings . . .

Beginnings . . .

The first axiom I learnt in Computer Science:

Beginnings . . .

The first axiom I learnt in Computer Science:

Computers might as well be made of green cheese

Beginnings . . .

The first axiom I learnt in Computer Science:

Computers might as well be made of green cheese

Beginnings . . .

The first axiom I learnt in Computer Science:

Computers might as well be made of green cheese

It is no longer safe to assume this!

Some Agendas for Quantum Computer Science

Some Agendas for Quantum Computer Science

- Information processing systems are **physically embodied**. The underlying physics is ultimately **quantum-mechanical**. Taking this seriously forces us to re-examine many of our basic assumptions about Computer Science.

Some Agendas for Quantum Computer Science

- Information processing systems are **physically embodied**. The underlying physics is ultimately **quantum-mechanical**. Taking this seriously forces us to re-examine many of our basic assumptions about Computer Science.
- It has already led to some exciting developments: remarkable new algorithms, cryptographic schemes, and basic questions in computational complexity.

Some Agendas for Quantum Computer Science

- Information processing systems are **physically embodied**. The underlying physics is ultimately **quantum-mechanical**. Taking this seriously forces us to re-examine many of our basic assumptions about Computer Science.
- It has already led to some exciting developments: remarkable new algorithms, cryptographic schemes, and basic questions in computational complexity.
- Beyond algorithms and complexity it offers new challenges and opportunities across the range of Computer Science: in programming languages and methods, logic and semantics.

Some Agendas for Quantum Computer Science

- Information processing systems are **physically embodied**. The underlying physics is ultimately **quantum-mechanical**. Taking this seriously forces us to re-examine many of our basic assumptions about Computer Science.
- It has already led to some exciting developments: remarkable new algorithms, cryptographic schemes, and basic questions in computational complexity.
- Beyond algorithms and complexity it offers new challenges and opportunities across the range of Computer Science: in programming languages and methods, logic and semantics.
- There is a fascinating two-way interplay developing between Computer Science and Physics, extending to the foundations of both, as well as to more practical matters. Quantum technology — "hacking matter" — will be a huge feature of 21st Century science and engineering, and a lot of it will be to do with information.

Some Agendas for Quantum Computer Science

- Information processing systems are **physically embodied**. The underlying physics is ultimately **quantum-mechanical**. Taking this seriously forces us to re-examine many of our basic assumptions about Computer Science.
- It has already led to some exciting developments: remarkable new algorithms, cryptographic schemes, and basic questions in computational complexity.
- Beyond algorithms and complexity it offers new challenges and opportunities across the range of Computer Science: in programming languages and methods, logic and semantics.
- There is a fascinating two-way interplay developing between Computer Science and Physics, extending to the foundations of both, as well as to more practical matters. Quantum technology — "hacking matter" — will be a huge feature of 21st Century science and engineering, and a lot of it will be to do with information.
- This is an exciting emerging area, attracting students with backgrounds in CS, Physics, Mathematics, Philosophy, . . .

Contextual Semantics

Contextual Semantics

- At the heart of quantum non-classicality are the phenomena of **non-locality**, **contextuality** and **entanglement**.

Contextual Semantics

- At the heart of quantum non-classicality are the phenomena of **non-locality**, **contextuality** and **entanglement**.
- These concepts play a central rôle in the rapidly developing field of quantum information, in delineating how quantum resources can transcend the bounds of classical information processing.

Contextual Semantics

- At the heart of quantum non-classicality are the phenomena of **non-locality**, **contextuality** and **entanglement**.
- These concepts play a central rôle in the rapidly developing field of quantum information, in delineating how quantum resources can transcend the bounds of classical information processing.
- They also have profound consequences for our understanding of the very nature of physical reality.

Contextual Semantics

- At the heart of quantum non-classicality are the phenomena of **non-locality**, **contextuality** and **entanglement**.
- These concepts play a central rôle in the rapidly developing field of quantum information, in delineating how quantum resources can transcend the bounds of classical information processing.
- They also have profound consequences for our understanding of the very nature of physical reality.
- We shall describe recent work in which tools from Computer Science are used to shed new light on these phenomena.

Contextual Semantics

- At the heart of quantum non-classicality are the phenomena of **non-locality**, **contextuality** and **entanglement**.
- These concepts play a central rôle in the rapidly developing field of quantum information, in delineating how quantum resources can transcend the bounds of classical information processing.
- They also have profound consequences for our understanding of the very nature of physical reality.
- We shall describe recent work in which tools from Computer Science are used to shed new light on these phenomena.
- There are also striking and unexpected connections with a number of topics in **classical** computer science, including relational databases and constraint satisfaction.

Contextual Semantics

- At the heart of quantum non-classicality are the phenomena of **non-locality**, **contextuality** and **entanglement**.
- These concepts play a central rôle in the rapidly developing field of quantum information, in delineating how quantum resources can transcend the bounds of classical information processing.
- They also have profound consequences for our understanding of the very nature of physical reality.
- We shall describe recent work in which tools from Computer Science are used to shed new light on these phenomena.
- There are also striking and unexpected connections with a number of topics in **classical** computer science, including relational databases and constraint satisfaction.
- So, let's go right to the heart of the quantum mystery . . .

Alice and Bob look at bits

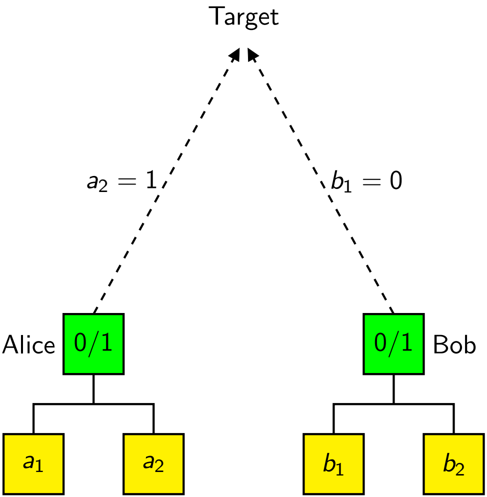

A Probabilistic Model Of An Experiment

A Probabilistic Model Of An Experiment

Example: The Bell Model

A	B	$(0,0)$	$(1,0)$	$(0,1)$	$(1,1)$
a_1	b_1	$1/2$	0	0	$1/2$
a_1	b_2	$3/8$	$1/8$	$1/8$	$3/8$
a_2	b_1	$3/8$	$1/8$	$1/8$	$3/8$
a_2	b_2	$1/8$	$3/8$	$3/8$	$1/8$

A Probabilistic Model Of An Experiment

Example: The Bell Model

A	B	$(0,0)$	$(1,0)$	$(0,1)$	$(1,1)$
a_1	b_1	$1/2$	0	0	$1/2$
a_1	b_2	$3/8$	$1/8$	$1/8$	$3/8$
a_2	b_1	$3/8$	$1/8$	$1/8$	$3/8$
a_2	b_2	$1/8$	$3/8$	$3/8$	$1/8$

The entry in row 2 column 3 says:

If Alice looks at a_1 and Bob looks at b_2, then $1/8$th of the time, Alice sees a 0 and Bob sees a 1.

A Probabilistic Model Of An Experiment

Example: The Bell Model

A	B	$(0,0)$	$(1,0)$	$(0,1)$	$(1,1)$
a_1	b_1	$1/2$	0	0	$1/2$
a_1	b_2	$3/8$	$1/8$	$1/8$	$3/8$
a_2	b_1	$3/8$	$1/8$	$1/8$	$3/8$
a_2	b_2	$1/8$	$3/8$	$3/8$	$1/8$

The entry in row 2 column 3 says:

If Alice looks at a_1 *and Bob looks at* b_2*, then* $1/8$*th of the time, Alice sees a* 0 *and Bob sees a* 1*.*

How can we explain this behaviour?

Classical Correlations: The Classical Source

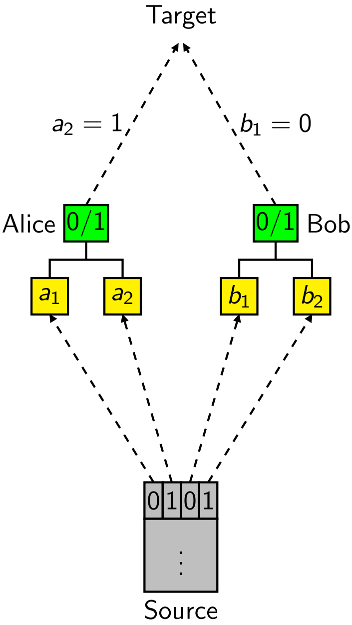

A Possibilistic Model Of An Experiment

A Possibilistic Model Of An Experiment

	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
(a_1, b_1)	1			
(a_1, b_2)	0			
(a_2, b_1)	0			
(a_2, b_2)				0

A Possibilistic Model Of An Experiment

	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
(a_1, b_1)	1			
(a_1, b_2)	0			
(a_2, b_1)	0			
(a_2, b_2)				0

The entry in row 1 column 1 says:

> *If Alice looks at a_1 and Bob looks at b_1, then* **sometimes** *Alice sees a* 0 *and Bob sees a* 0*.*

A Possibilistic Model Of An Experiment

	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
(a_1, b_1)	1			
(a_1, b_2)	0			
(a_2, b_1)	0			
(a_2, b_2)				0

The entry in row 1 column 1 says:

> *If Alice looks at a_1 and Bob looks at b_1, then* **sometimes** *Alice sees a 0 and Bob sees a 0.*

The entry in row 2 column 1 says:

> *If Alice looks at a_1 and Bob looks at b_2, then* **it never happens** *that Alice sees a 0 and Bob sees a 0.*

A Possibilistic Model Of An Experiment

	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
(a_1, b_1)	1			
(a_1, b_2)	0			
(a_2, b_1)	0			
(a_2, b_2)				0

The entry in row 1 column 1 says:

> *If Alice looks at a_1 and Bob looks at b_1, then* **sometimes** *Alice sees a 0 and Bob sees a 0.*

The entry in row 2 column 1 says:

> *If Alice looks at a_1 and Bob looks at b_2, then* **it never happens** *that Alice sees a 0 and Bob sees a 0.*

Can we explain this behaviour using a classical source?

What Do 'Observables' Observe?

What Do ‘Observables’ Observe?

Surely **objective properties** of a physical system, which are independent of our choice of which measurements to perform — of our **measurement context**.

What Do 'Observables' Observe?

Surely **objective properties** of a physical system, which are independent of our choice of which measurements to perform — of our **measurement context**.

More precisely, this would say that for each possible state of the system, there is a function λ which for each measurement m specifies an outcome $\lambda(m)$, **independently of which other measurements may be performed**.

What Do 'Observables' Observe?

Surely **objective properties** of a physical system, which are independent of our choice of which measurements to perform — of our **measurement context**.

More precisely, this would say that for each possible state of the system, there is a function λ which for each measurement m specifies an outcome $\lambda(m)$, **independently of which other measurements may be performed**.

This point of view is called **non-contextuality**. It is equivalent to the assumption of a classical source.

What Do 'Observables' Observe?

Surely **objective properties** of a physical system, which are independent of our choice of which measurements to perform — of our **measurement context**.

More precisely, this would say that for each possible state of the system, there is a function λ which for each measurement m specifies an outcome $\lambda(m)$, **independently of which other measurements may be performed**.

This point of view is called **non-contextuality**. It is equivalent to the assumption of a classical source.

However, this view is **impossible to sustain** in the light of our **actual observations of (micro)-physical reality**.

Hidden Variables: The Mermin instruction set picture

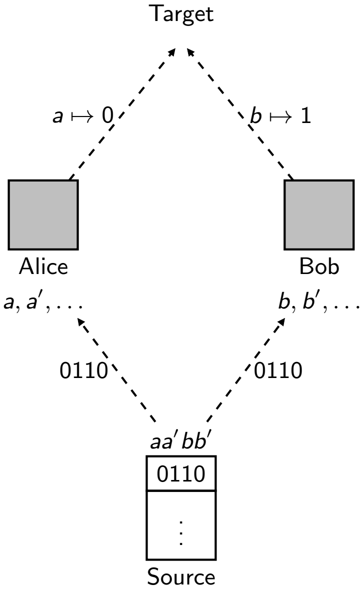

The ‘Hardy Paradox’

The 'Hardy Paradox'

Hardy models: those whose support satisfies

	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
(a_1, b_1)	1			
(a_1, b_2)	0			
(a_2, b_1)	0			
(a_2, b_2)				0

Which 'instruction set' λ could the outcomes $(0,0)$ for measurements (a_1, b_1) could come? Clearly, we must have

$$\lambda : a_1 \mapsto 0, \quad b_1 \mapsto 0.$$

The 'Hardy Paradox'

Hardy models: those whose support satisfies

	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
(a_1, b_1)	1			
(a_1, b_2)	0	?		
(a_2, b_1)	0			
(a_2, b_2)				0

Which 'instruction set' λ could the outcomes $(0,0)$ for measurements (a_1, b_1) could come? Clearly, we must have

$$\lambda : a_1 \mapsto 0, \quad b_1 \mapsto 0.$$

The 'Hardy Paradox'

Hardy models: those whose support satisfies

	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
(a_1, b_1)	1			
(a_1, b_2)	0	1		
(a_2, b_1)	0			
(a_2, b_2)				0

Which 'instruction set' λ could the outcomes $(0,0)$ for measurements (a_1, b_1) could come? Clearly, we must have

$$\lambda : a_1 \mapsto 0, \quad b_1 \mapsto 0.$$

The 'Hardy Paradox'

Hardy models: those whose support satisfies

	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
(a_1, b_1)	1			
(a_1, b_2)	0	1		
(a_2, b_1)	0			
(a_2, b_2)		?		0

Which 'instruction set' λ could the outcomes $(0,0)$ for measurements (a_1, b_1) could come? Clearly, we must have

$$\lambda : a_1 \mapsto 0, \quad b_1 \mapsto 0.$$

The 'Hardy Paradox'

Hardy models: those whose support satisfies

	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
(a_1, b_1)	1			
(a_1, b_2)	0	1		
(a_2, b_1)	0			
(a_2, b_2)		1		0

Which 'instruction set' λ could the outcomes $(0,0)$ for measurements (a_1, b_1) could come? Clearly, we must have

$$\lambda : a_1 \mapsto 0, \quad b_1 \mapsto 0.$$

The 'Hardy Paradox'

Hardy models: those whose support satisfies

	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
(a_1, b_1)	1			
(a_1, b_2)	0	1		
(a_2, b_1)	0			
(a_2, b_2)		1		0

So there is a unique 'instruction set' λ that outcomes $(0,0)$ for measurements (a_1, b_1) could come from:

$$\lambda : a_1 \mapsto 0, \quad a_2 \mapsto 0, \quad b_1 \mapsto 0, \quad b_2 \mapsto 1.$$

The 'Hardy Paradox'

Hardy models: those whose support satisfies

	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
(a_1, b_1)	1			
(a_1, b_2)	0	1		
(a_2, b_1)	0			
(a_2, b_2)		1		0

So there is a unique 'instruction set' λ that outcomes $(0,0)$ for measurements (a_1, b_1) could come from:

$$\lambda : a_1 \mapsto 0, \quad a_2 \mapsto 0, \quad b_1 \mapsto 0, \quad b_2 \mapsto 1.$$

However, this would require the outcome $(0,0)$ for measurements (a_2, b_1) to be possible, and this is **precluded**.

The 'Hardy Paradox'

Hardy models: those whose support satisfies

	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
(a_1, b_1)	1			
(a_1, b_2)	0	1		
(a_2, b_1)	0			
(a_2, b_2)		1		0

So there is a unique 'instruction set' λ that outcomes $(0,0)$ for measurements (a_1, b_1) could come from:

$$\lambda : a_1 \mapsto 0, \quad a_2 \mapsto 0, \quad b_1 \mapsto 0, \quad b_2 \mapsto 1.$$

However, this would require the outcome $(0,0)$ for measurements (a_2, b_1) to be possible, and this is **precluded**.

Thus Hardy models are **contextual**. They cannot be explained by a classical source.

Quantum Mechanics changes the game

Quantum Mechanics changes the game

It seems then that the kind of behaviour exhibited in these tables is not realisable.

Quantum Mechanics changes the game

It seems then that the kind of behaviour exhibited in these tables is not realisable.

However, if we use **quantum** rather than classical resources, it **is** realisable!

Quantum Mechanics changes the game

It seems then that the kind of behaviour exhibited in these tables is not realisable.

However, if we use **quantum** rather than classical resources, it **is** realisable!

More specifically, if we use an **entangled qubit** as a shared resource between Alice and Bob, who may be spacelike separated, then behaviour of exactly the kind we have considered **can** be achieved.

Quantum Mechanics changes the game

It seems then that the kind of behaviour exhibited in these tables is not realisable.

However, if we use **quantum** rather than classical resources, it **is** realisable!

More specifically, if we use an **entangled qubit** as a shared resource between Alice and Bob, who may be spacelike separated, then behaviour of exactly the kind we have considered **can** be achieved.

Alice and Bob's choices are now of **measurement setting** (e.g. which direction to measure spin) rather than "which register to load".

The Quantum Case: Spin Measurements

The Quantum Case: Spin Measurements

States of the system can be described by complex unit vectors in $\mathbb{C}^2$. These can be visualized as points on the unit 2-sphere:

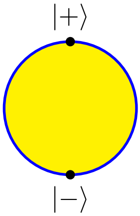

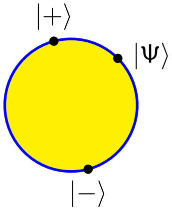

The Quantum Case: Spin Measurements

States of the system can be described by complex unit vectors in $\mathbb{C}^2$. These can be visualized as points on the unit 2-sphere:

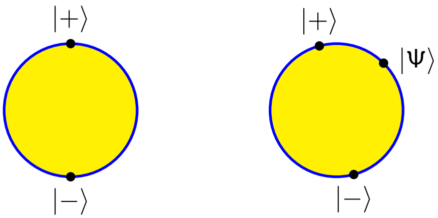

Spin can be measured in any direction; so there are a continuum of possible measurements. There are **two possible outcomes** for each such measurement; spin in the specified direction, or in the opposite direction. These two directions are represented by a pair of orthogonal vectors. They are represented on the sphere as a pair of **antipodal points**.

The Quantum Case: Spin Measurements

States of the system can be described by complex unit vectors in $\mathbb{C}^2$. These can be visualized as points on the unit 2-sphere:

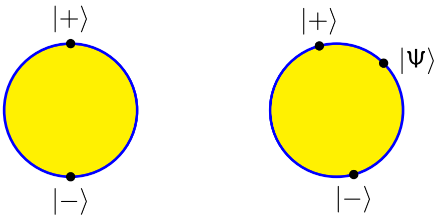

Spin can be measured in any direction; so there are a continuum of possible measurements. There are **two possible outcomes** for each such measurement; spin in the specified direction, or in the opposite direction. These two directions are represented by a pair of orthogonal vectors. They are represented on the sphere as a pair of **antipodal points**.

Note the appearance of **quantization** here: there are not a continuum of possible outcomes for each measurement, but only two!

The Stern-Gerlach Experiment

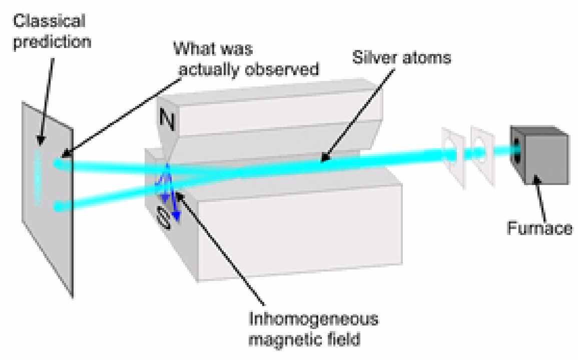

Quantum Entanglement

Quantum Entanglement

Bell state:

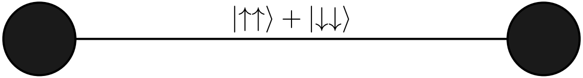

EPR state:

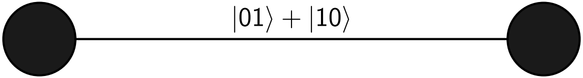

Quantum Entanglement

Bell state:

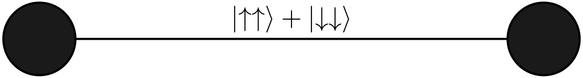

EPR state:

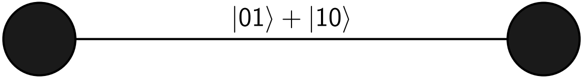

Compound systems are represented by **tensor product**: $\mathcal{H}_1 \otimes \mathcal{H}_2$. Typical element:

$$\sum_i \lambda_i \cdot \phi_i \otimes \psi_i$$

Superposition encodes **correlation**.

Quantum Entanglement

Bell state:

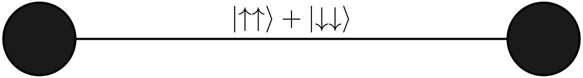

EPR state:

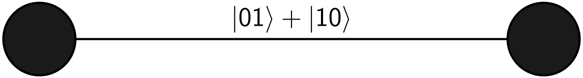

Compound systems are represented by **tensor product**: $\mathcal{H}_1 \otimes \mathcal{H}_2$. Typical element:

$$\sum_i \lambda_i \cdot \phi_i \otimes \psi_i$$

Superposition encodes **correlation**.

Einstein's 'spooky action at a distance'. Even if the particles are spatially separated, measuring one has an effect on the state of the other.

Quantum Entanglement

Bell state:

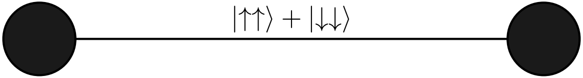

EPR state:

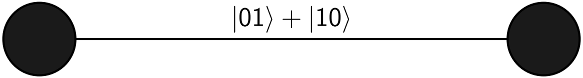

Compound systems are represented by **tensor product**: $\mathcal{H}_1 \otimes \mathcal{H}_2$. Typical element:

$$\sum_i \lambda_i \cdot \phi_i \otimes \psi_i$$

Superposition encodes **correlation**.

Einstein's 'spooky action at a distance'. Even if the particles are spatially separated, measuring one has an effect on the state of the other.

Bell's theorem: QM is **essentially non-local**.

A Possibilistic Model Of An Experiment

A Possibilistic Model Of An Experiment

	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
(a_1, b_1)	1			
(a_1, b_2)	0			
(a_2, b_1)	0			
(a_2, b_2)				0

A Possibilistic Model Of An Experiment

	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
(a_1, b_1)	1			
(a_1, b_2)	0			
(a_2, b_1)	0			
(a_2, b_2)				0

This model can be **physically realised** in quantum mechanics.

A Possibilistic Model Of An Experiment

	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
(a_1, b_1)	1			
(a_1, b_2)	0			
(a_2, b_1)	0			
(a_2, b_2)				0

This model can be **physically realised** in quantum mechanics.

There is an entangled state of two qubits, and directions for spin measurements a_1, a_2 for Alice and b_1, b_2 for Bob, which generate this table according to the predictions of quantum mechanics.

A Possibilistic Model Of An Experiment

	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
(a_1, b_1)	1			
(a_1, b_2)	0			
(a_2, b_1)	0			
(a_2, b_2)				0

This model can be **physically realised** in quantum mechanics.

There is an entangled state of two qubits, and directions for spin measurements a_1, a_2 for Alice and b_1, b_2 for Bob, which generate this table according to the predictions of quantum mechanics.

Moreover, behaviour of this kind has been extensively experimentally confirmed.

A Possibilistic Model Of An Experiment

	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
(a_1, b_1)	1			
(a_1, b_2)	0			
(a_2, b_1)	0			
(a_2, b_2)				0

This model can be **physically realised** in quantum mechanics.

There is an entangled state of two qubits, and directions for spin measurements a_1, a_2 for Alice and b_1, b_2 for Bob, which generate this table according to the predictions of quantum mechanics.

Moreover, behaviour of this kind has been extensively experimentally confirmed.

This is really how the world is!

A Possibilistic Model Of An Experiment

	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
(a_1, b_1)	1			
(a_1, b_2)	0			
(a_2, b_1)	0			
(a_2, b_2)				0

This model can be **physically realised** in quantum mechanics.

There is an entangled state of two qubits, and directions for spin measurements a_1, a_2 for Alice and b_1, b_2 for Bob, which generate this table according to the predictions of quantum mechanics.

Moreover, behaviour of this kind has been extensively experimentally confirmed.

This is really how the world is!

This proves a **strong version of Bell's theorem**.

Mathematical Structure of Possibility Tables

Mathematical Structure of Possibility Tables

	$(0,0)$	$(1,0)$	$(0,1)$	$(1,1)$
(a, b)	1	1	1	1
(a', b)	0	1	1	1
(a, b')	0	1	1	1
(a', b')	1	1	1	0

Mathematical Structure of Possibility Tables

	$(0,0)$	$(1,0)$	$(0,1)$	$(1,1)$
(a,b)	1	1	1	1
(a',b)	0	1	1	1
(a,b')	0	1	1	1
(a',b')	1	1	1	0

The **measurement contexts** are

$$\{a,b\}, \quad \{a',b\}, \quad \{a,b'\}, \quad \{a',b'\}.$$

Mathematical Structure of Possibility Tables

	$(0,0)$	$(1,0)$	$(0,1)$	$(1,1)$
(a,b)	1	1	1	1
(a',b)	0	1	1	1
(a,b')	0	1	1	1
(a',b')	1	1	1	0

The **measurement contexts** are

$$\{a,b\}, \quad \{a',b\}, \quad \{a,b'\}, \quad \{a',b'\}.$$

Each measurement has possible outcomes 0 or 1. The matrix entry at row (a',b) and column $(0,1)$ indicates the **event**

$$\{a' \mapsto 0, \; b \mapsto 1\}.$$

Mathematical Structure of Possibility Tables

	$(0,0)$	$(1,0)$	$(0,1)$	$(1,1)$
(a,b)	1	1	1	1
(a',b)	0	1	1	1
(a,b')	0	1	1	1
(a',b')	1	1	1	0

The **measurement contexts** are

$$\{a,b\}, \quad \{a',b\}, \quad \{a,b'\}, \quad \{a',b'\}.$$

Each measurement has possible outcomes 0 or 1. The matrix entry at row (a',b) and column $(0,1)$ indicates the **event**

$$\{a' \mapsto 0, \; b \mapsto 1\}.$$

Each row of the table specifies a **Boolean distribution** on events O^C for a given choice of measurement context C.

The Geometry of Contextuality

The Geometry of Contextuality

We have the following formal structure:

The Geometry of Contextuality

We have the following formal structure:

- a set of measurements X (the 'space');

The Geometry of Contextuality

We have the following formal structure:

- a set of measurements X (the 'space');
- a family of subsets of X, the **measurement contexts** (a 'cover');

The Geometry of Contextuality

We have the following formal structure:

- a set of measurements X (the 'space');
- a family of subsets of X, the **measurement contexts** (a 'cover');
- to each such set C a boolean distribution (finite non-empy subset) on **local sections** $s : C \to O$, where O is the set of **outcomes**.

The Geometry of Contextuality

We have the following formal structure:

- a set of measurements X (the 'space');
- a family of subsets of X, the **measurement contexts** (a 'cover');
- to each such set C a boolean distribution (finite non-empy subset) on **local sections** $s : C \to O$, where O is the set of **outcomes**.
- a distribution on C restricts to $C' \subseteq C$ by pointwise restriction of the local sections.

The Geometry of Contextuality

We have the following formal structure:

- a set of measurements X (the 'space');
- a family of subsets of X, the **measurement contexts** (a 'cover');
- to each such set C a boolean distribution (finite non-empy subset) on **local sections** $s : C \to O$, where O is the set of **outcomes**.
- a distribution on C restricts to $C' \subseteq C$ by pointwise restriction of the local sections.

These local sections correspond to the directly observable **joint outcomes** of **compatible measurements**, which can actually be performed jointly on the system.

The Geometry of Contextuality

We have the following formal structure:

- a set of measurements X (the 'space');
- a family of subsets of X, the **measurement contexts** (a 'cover');
- to each such set C a boolean distribution (finite non-empy subset) on **local sections** $s : C \to O$, where O is the set of **outcomes**.
- a distribution on C restricts to $C' \subseteq C$ by pointwise restriction of the local sections.

These local sections correspond to the directly observable **joint outcomes** of **compatible measurements**, which can actually be performed jointly on the system.

The different sets of compatible measurements correspond to the different contexts of measurement and observation of the physical system.

The Geometry of Contextuality

We have the following formal structure:

- a set of measurements X (the 'space');
- a family of subsets of X, the **measurement contexts** (a 'cover');
- to each such set C a boolean distribution (finite non-empty subset) on **local sections** $s : C \to O$, where O is the set of **outcomes**.
- a distribution on C restricts to $C' \subseteq C$ by pointwise restriction of the local sections.

These local sections correspond to the directly observable **joint outcomes** of **compatible measurements**, which can actually be performed jointly on the system.

The different sets of compatible measurements correspond to the different contexts of measurement and observation of the physical system.

The fact that the behaviour of these observable outcomes cannot be accounted for by some context-independent global description of reality corresponds to the geometric fact that these local sections cannot be glued together into a **global section**.

Gluing functional sections

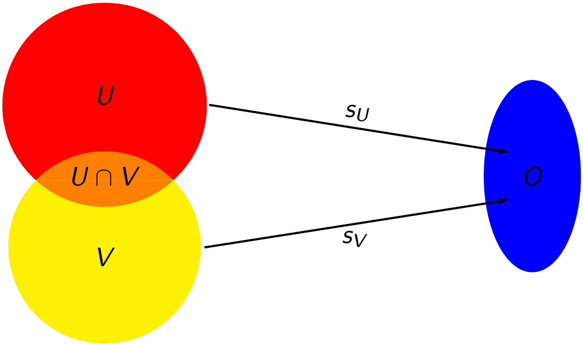

Gluing functional sections

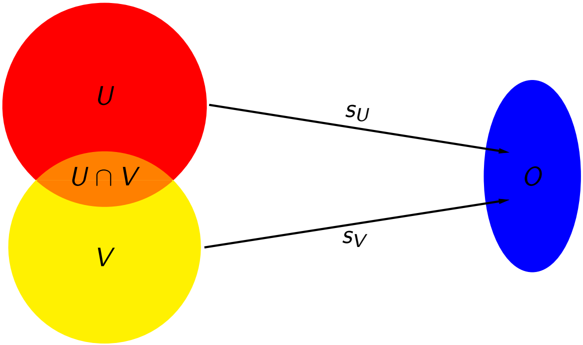

If $s_U|_{U \cap V} = s_V|_{U \cap V}$, they can be glued to form

$$s : U \cup V \longrightarrow O$$

such that $s|_U = s_U$ and $s|_V = s_V$.

Obstructions to gluing distributions

Obstructions to gluing distributions

In geometric language, the Hardy paradox corresponds to the fact that there is a **local section** which cannot be extended to a **global section** which is compatible with the family of boolean distributions.

Obstructions to gluing distributions

In geometric language, the Hardy paradox corresponds to the fact that there is a **local section** which cannot be extended to a **global section** which is compatible with the family of boolean distributions.

In other words, the space of **local possibilities** is sufficiently logically ‘twisted’ to **obstruct** such an extension.

Obstructions to gluing distributions

In geometric language, the Hardy paradox corresponds to the fact that there is a **local section** which cannot be extended to a **global section** which is compatible with the family of boolean distributions.

In other words, the space of **local possibilities** is sufficiently logically 'twisted' to **obstruct** such an extension.

The quantum phenomena of **non-locality** and **contextuality** correspond exactly to the existence of obstructions to global sections in this sense.

Obstructions to gluing distributions

In geometric language, the Hardy paradox corresponds to the fact that there is a **local section** which cannot be extended to a **global section** which is compatible with the family of boolean distributions.

In other words, the space of **local possibilities** is sufficiently logically 'twisted' to **obstruct** such an extension.

The quantum phenomena of **non-locality** and **contextuality** correspond exactly to the existence of obstructions to global sections in this sense.

This geometric picture and the associated methods can be applied to a wide range of situations in classical computer science.

Obstructions to gluing distributions

In geometric language, the Hardy paradox corresponds to the fact that there is a **local section** which cannot be extended to a **global section** which is compatible with the family of boolean distributions.

In other words, the space of **local possibilities** is sufficiently logically 'twisted' to **obstruct** such an extension.

The quantum phenomena of **non-locality** and **contextuality** correspond exactly to the existence of obstructions to global sections in this sense.

This geometric picture and the associated methods can be applied to a wide range of situations in classical computer science.

In particular, as we shall now see, there is an isomorphism between the formal description we have given for the quantum notions of non-locality and contextuality, and basic definitions and concepts in relational database theory.

Relational databases

Relational databases

Samson Abramsky, 'Relational databases and Bell's theorem', In *In Search of Elegance in the Theory and Practice of Computation: Essays Dedicated to Peter Buneman*, Springer 2013.

Relational databases

Samson Abramsky, ‘Relational databases and Bell’s theorem’, In *In Search of Elegance in the Theory and Practice of Computation: Essays Dedicated to Peter Buneman*, Springer 2013.

branch-name	**account-no**	**customer-name**	**balance**
Cambridge	10991-06284	Newton	£2,567.53
Hanover	10992-35671	Leibniz	€11,245.75
...	...	...	...

From possibility models to databases

From possibility models to databases

Consider again the Hardy model:

	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
(a_1, b_1)	1	1	1	1
(a_1, b_2)	0	1	1	1
(a_2, b_1)	0	1	1	1
(a_2, b_2)	1	1	1	0

From possibility models to databases

Consider again the Hardy model:

	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
(a_1, b_1)	1	1	1	1
(a_1, b_2)	0	1	1	1
(a_2, b_1)	0	1	1	1
(a_2, b_2)	1	1	1	0

Change of perspective:

$a_1,\ a_2,\ b_1,\ b_2$	attributes
0, 1	data values
joint outcomes of measurements	tuples

The Hardy model as a relational database

The four rows of the model turn into four **relation tables**:

a_1	b_1
0	0
0	1
1	0
1	1

a_1	b_2
0	1
1	0
1	1

a_2	b_1
0	1
1	0
1	1

a_2	b_2
0	0
1	0
0	1

The Hardy model as a relational database

The four rows of the model turn into four **relation tables**:

a_1	b_1
0	0
0	1
1	0
1	1

a_1	b_2
0	1
1	0
1	1

a_2	b_1
0	1
1	0
1	1

a_2	b_2
0	0
1	0
0	1

What is the DB property corresponding to the presence of non-locality/contextuality in the Hardy table?

The Hardy model as a relational database

The four rows of the model turn into four **relation tables**:

a_1	b_1
0	0
0	1
1	0
1	1

a_1	b_2
0	1
1	0
1	1

a_2	b_1
0	1
1	0
1	1

a_2	b_2
0	0
1	0
0	1

What is the DB property corresponding to the presence of non-locality/contextuality in the Hardy table?

There is no **universal relation**: no table

a_1	a_2	b_1	b_2
$\vdots$	$\vdots$	$\vdots$	$\vdots$

whose projections onto $\{a_i, b_i\}$, $i = 1, 2$, yield the above four tables.

A dictionary

A dictionary

Relational databases	measurement scenarios
attribute	measurement
set of attributes defining a relation table	compatible set of measurements
database schema	measurement cover
tuple	local section (joint outcome)
relation/set of tuples	boolean distribution on joint outcomes
universal relation instance	global section/hidden variable model
acyclicity	Vorob'ev condition

A dictionary

Relational databases	measurement scenarios
attribute	measurement
set of attributes defining a relation table	compatible set of measurements
database schema	measurement cover
tuple	local section (joint outcome)
relation/set of tuples	boolean distribution on joint outcomes
universal relation instance	global section/hidden variable model
acyclicity	Vorob'ev condition

We can also consider probabilistic databases and other generalisations; cf. provenance semirings.

Contextual Semantics

Contextual Semantics

Why do such similar structures arise in such apparently different settings?

Contextual Semantics

Why do such similar structures arise in such apparently different settings?

The phenomenon of contextuality is pervasive. Once we start looking for it, we can find it everywhere!

Physics, computation, logic, natural language, . . . biology, economics, . . .

Contextual Semantics

Why do such similar structures arise in such apparently different settings?

The phenomenon of contextuality is pervasive. Once we start looking for it, we can find it everywhere!
Physics, computation, logic, natural language, . . . biology, economics, . . .

The **Contextual semantics hypothesis**: we can find common mathematical structure in all these diverse manifestations, and develop a widely applicable theory.

Contextual Semantics

Why do such similar structures arise in such apparently different settings?

The phenomenon of contextuality is pervasive. Once we start looking for it, we can find it everywhere!
Physics, computation, logic, natural language, . . . biology, economics, . . .

The **Contextual semantics hypothesis**: we can find common mathematical structure in all these diverse manifestations, and develop a widely applicable theory.

More than a hypothesis! Already extensive results in

Contextual Semantics

Why do such similar structures arise in such apparently different settings?

The phenomenon of contextuality is pervasive. Once we start looking for it, we can find it everywhere!
Physics, computation, logic, natural language, . . . biology, economics, . . .

The **Contextual semantics hypothesis**: we can find common mathematical structure in all these diverse manifestations, and develop a widely applicable theory.

More than a hypothesis! Already extensive results in

- Quantum information and foundations: hierarchy of contextuality, logical characterisation of Bell inequalities, classification of multipartite entangled states, cohomological characterisation of contextuality, structural explanation of macroscopic locality, . . .

Contextual Semantics

Why do such similar structures arise in such apparently different settings?

The phenomenon of contextuality is pervasive. Once we start looking for it, we can find it everywhere!
Physics, computation, logic, natural language, ... biology, economics, ...

The **Contextual semantics hypothesis**: we can find common mathematical structure in all these diverse manifestations, and develop a widely applicable theory.

More than a hypothesis! Already extensive results in

- Quantum information and foundations: hierarchy of contextuality, logical characterisation of Bell inequalities, classification of multipartite entangled states, cohomological characterisation of contextuality, structural explanation of macroscopic locality, ...
- And beyond: connections with databases, robust refinement of the constraint satisfaction paradigm, application of contextual semantics to natural language semantics, connections with team semantics in Dependence logics, ...

Contextual Semantics

Why do such similar structures arise in such apparently different settings?

The phenomenon of contextuality is pervasive. Once we start looking for it, we can find it everywhere!
Physics, computation, logic, natural language, ... biology, economics, ...

The **Contextual semantics hypothesis**: we can find common mathematical structure in all these diverse manifestations, and develop a widely applicable theory.

More than a hypothesis! Already extensive results in

- Quantum information and foundations: hierarchy of contextuality, logical characterisation of Bell inequalities, classification of multipartite entangled states, cohomological characterisation of contextuality, structural explanation of macroscopic locality, ...
- And beyond: connections with databases, robust refinement of the constraint satisfaction paradigm, application of contextual semantics to natural language semantics, connections with team semantics in Dependence logics, ...

For an accessible overview of Contextual Semantics, see the article in the *Logic in Computer Science* Column, Bulletin of EATCS No. 113, June 2014 (and `arXiv`).

People

Comrades in Arms in Contextual Semantics:

People

Comrades in Arms in Contextual Semantics:










People

Comrades in Arms in Contextual Semantics:

Adam Brandenburger, Lucien Hardy, Shane Mansfield, Rui Soares Barbosa, Ray Lal, Mehrnoosh Sadrzadeh, Phokion Kolaitis, Georg Gottlob, Carmen Constantin, Kohei Kishida

Logical Dynamics from a Dynamic Logician

Logical Dynamics from a Dynamic Logician

We look forward to much more to come!

Logical Dynamics from a Dynamic Logician

We look forward to much more to come!

Thank You Johan