LOCAL-VS-GLOBAL CONSISTENCY OF ANNOTATED RELATIONS

Albert Atserias
Universitat Politècnica de Catalunya
Barcelona, Catalonia, Spain

June 2025

Part I

Motivation

Escher's Local vs. Global

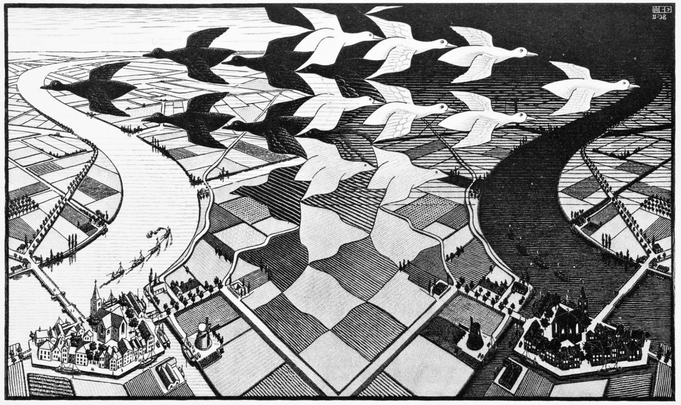

Day and Night, woodcut, Escher 1938.

[Low resolution image downloaded from Wikipedia]

Relational database consistency (on the triangle schema)

$R(X,Y)$		$S(Y,Z)$		$T(Z,X)$	
1	1	1	2	2	1
1	2	1	3	2	2
2	1	2	3	3	1
3	2			3	3

$W(X,Y,Z)$		
1	1	2
1	1	3
2	1	2
3	2	3
1	2	3

W[X,Y] = R   W[Y,Z] = S   W[Z,X] = T

R, S, T are consistent
W is a witness of their consistency.

Relational database consistency (arbitrary schema)

Let $X_1, \ldots, X_m$ be a schema.
Let $R_1(X_1), \ldots, R_m(X_m)$ be relations over that schema.

Definition [BFMY'83]
The relations $R_1(X_1), \ldots, R_m(X_m)$ are consistent if there exists a relation $W(X_1 \cdots X_m)$ that projects on X_i to R_i, for $i = 1, \ldots, m$; i.e.,

$$W[X_1] = R_1 \qquad W[X_2] = R_2 \qquad \cdots \qquad W[X_m] = R_m$$

We say that W is a witness of their consistency.

- pairwise consistent: any two are consistent,
- k-wise consistent: any k are consistent,
- globally consistent: all together are consistent.

Joins do the job ... right?

Basic fact about relations:

If $R(X)$ and $S(Y)$ are consistent relations,
then their join $R \bowtie S$ witnesses their consistency.

BUT not so for "*real-world* relations", i.e., bags.

The bag-join is *not* a witness of consistency for bags. [AK'21]

$R(X)$	$S(Y)$	$W(X,Y)$	$J(X,Y)$	$J[X]$	$J[Y]$
a_1	b_1	$a_1\ b_1$	$a_1\ b_1$	a_1	b_1
a_2	b_2	$a_2\ b_2$	$a_1\ b_2$	a_1	b_2
			$a_2\ b_1$	a_2	b_1
			$a_2\ b_2$	a_2	b_2

A more general problem

Let data come annotated with side information $(\alpha_i, \beta_j, \gamma_k, \ldots)$.

$R(X,Y)$	$S(Y,Z)$	$T(Z,X)$		$W(X,Y,Z)$
$a_1\ b_1 : \alpha_1$	$c_1\ d_1 : \beta_1$	$e_1\ f_1 : \gamma_1$		$g_1\ h_1\ i_1 :$ **?**
$a_2\ b_2 : \alpha_2$	$c_2\ d_2 : \beta_2$	$e_2\ f_2 : \gamma_2$		$g_2\ h_2\ i_2 :$ **?**
$\ldots$	$\ldots$	$\ldots$		$\ldots$
$a_m\ b_m : \alpha_m$	$c_n\ d_n : \beta_n$	$e_p\ f_p : \gamma_p$		$g_q\ h_q\ i_q :$ **?**

Questions:
How should data be annotated/measured/compared/aggregated?
Given data, does it come from a common source?
Can we reconstruct the source?

Example 1: Image reconstruction and tomography

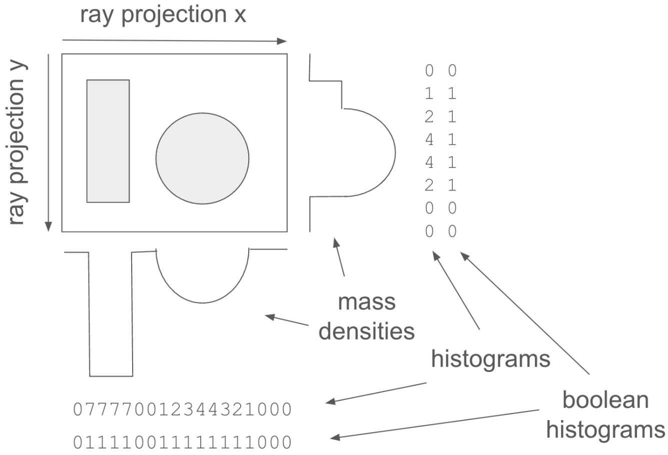

Cormack-Hounsfield 1979 Nobel in Physiology or Medicine

Example 2: Quantum theory, EPR, and Bell inequalities

Do measurements reflect "elements of physical reality"?
[EPR'35, NYT'35, B'64]


The New York Times

EINSTEIN ATTACKS QUANTUM THEORY; Scientist and Two Colleagues Find It Is Not 'Complete' Even Though 'Correct.'

Share full article

Copyright 1935 by Science Service.
May 4, 1935


EINSTEIN ATTACKS QUANTUM THEORY

BOOKS OF THE TIMES

https://www.nytimes.com/1935/05/04/archives/einstein-attacks-quantum-theory-scientist-and-two-colleagues-find.html?smid=url-share

Measuring two classical Head/Tail coins

In classical mechanics, the uncertainty of an experiment can be modelled by hidden variable theories. E.g.,

Coin 1	Coin 2	Wit : λ	Wit : λ	etc
$H : 1/2$	$H : 1/2$	$HH : 1/4$	$HH : 1/2$	
$T : 1/2$	$T : 1/2$	$HT : 1/4$	$HT : 0$	
		$TH : 1/4$	$TH : 0$	
		$TT : 1/4$	$TT : 1/2$	

The hidden variable theories model elements of physical reality.
Facts are "already there", just unknown before measurement.
The hidden variable theory need not be unique.

Measuring two quantum entangled particles

In quantum mechanics, states are unit vectors and measurements are orthogonal projection operators that collapse the state. E.g.,

$$B_1 = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \qquad C_1 = \begin{bmatrix} 1/2 & -1/2 \\ -1/2 & 1/2 \end{bmatrix}$$

$$B_2 = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \qquad C_2 = \begin{bmatrix} 1/2 & 1/2 \\ 1/2 & 1/2 \end{bmatrix}$$

A consequence of Bell's analysis in [B'64] is that the combined system does not admit compatible measurements.

Particle 1	Particle 2	Wit?
$b_1 : B_1$	$c_1 : C_1$	$b_1 c_1 : X_{11}$?
$b_2 : B_2$	$c_2 : C_2$	$b_1 c_2 : X_{12}$?
		$b_2 c_1 : X_{21}$?
		$b_2 c_2 : X_{22}$?

Ergo: no local hidden variable theory can model entanglement.

Part II

Annotated Relations

Annotations from an algebraic structure

- Set semantics use True (1) and False (0) as annotations.
- Bag semantics uses natural numbers as annotations.
- Semiring semantics uses annotations from a semiring:

A semiring is an algebraic structure $\mathbb{K} = (K, +, \times, 0, 1)$ where $(K, +, 0)$ and $(K, \times, 1)$ are commutative monoids (associative and commutative with neutral element), and $\times$ distributes over $+$.

Examples:
- Two-element Boolean algebra $\mathbb{B} = \{0, 1\}$ with $\vee$ and $\wedge$,
- Natural numbers $\mathbb{N}$ with $+$ and $\times$,
- Extended real numbers $\mathbb{R} \cup \{\pm\infty\}$ with min and max.
- Extended real numbers $\mathbb{R} \cup \{\pm\infty\}$ with min and $+$ (tropical).
- ...

Semiring semantics in database theory

Extensively studied in last two decades:
- Provenance [GKT'07, KG'12, DGNT'21].
- Query containment [G'11, KRS'14].
- Datalog and recursion [KNPSW'24, DGNT'21].

Key idea 1:
- Addition $+$ is used for "alternative information" (or/projection).
- Multiplication $\times$ is used for "joint information" (and/join).

Key idea 2:
- To define consistency, only the additive structure is relevant.
- As noted earlier for bags, the standard join ($\times$) may not witness.
- To model consistency, positivity (defined next) is natural.

Positive commutative monoids

A positive commutative monoid is an algebraic structure $\mathbb{K} = (K, +, 0)$ where $+$ is an associative and commutative operation on K, with neutral element 0, which satisfies positivity:

$$x + y = 0 \quad \text{implies} \quad x = 0 \text{ and } y = 0.$$

Examples

- Boolean monoid: $\mathbb{B} = \{\texttt{true}, \texttt{false}\}$ with $\vee$ and $\texttt{false}$.
- Powerset monoid: $\mathcal{P}(S)$ with $\cup$ and $\emptyset$ for some set S.
- Bag monoid: $\mathbb{N}$ with $+$ and 0.
- Real-valued measures: $\mathbb{R}_{\geq 0}$ with $+$ and 0.
- POVM of dimension d: $\mathbb{R}^{d \times d}_{\succeq 0}$ with $+$ and 0.
- Tropical/Cost monoid: $\mathbb{R} \cup \{\infty\}$ with min and ∞.
- Access control [GT'17]: $\texttt{P} < \texttt{S} < \texttt{T} < \texttt{I}$ with min and $\texttt{I}$.

$\mathbb{K}$-relations and their projections

Let $\mathbb{K} = (K, +, 0)$ be a positive commutative monoid.
Let X be a set of attributes with domain D.

A $\mathbb{K}$-relation $R(X)$ of schema X is a map $R : D^X \to K$ with finite support

$$\mathrm{Supp}(R) := \{t \in D^X \ : \ R(t) \neq 0\}.$$

For $Y \subseteq X$, the Y-projection denoted $R[Y]$ is the $\mathbb{K}$-relation of schema Y defined on every Y-tuple r by

$$R[Y](r) := \sum_{\substack{t \in D^X : \\ t[Y] = r}} R(t)$$

Fact. Positivity of $\mathbb{K}$ ensures projection commutes with support:

$$\mathrm{Supp}(R[Y]) = \mathrm{Supp}(R)[Y].$$

Consistency of relations over monoids

Let $\mathbb{K}$ be a positive commutative monoid.
Let $X_1, \ldots, X_m$ be a schema.
Let $R_1(X_1), \ldots, R_m(X_m)$ be $\mathbb{K}$-relations over that schema.

Definition [AK'24]
The $\mathbb{K}$-relations $R_1(X_1), \ldots, R_m(X_m)$ are consistent if there exists a $\mathbb{K}$-relation $W(X_1 \cdots X_m)$ that projects on X_i to R_i, for $i = 1, \ldots, m$; i.e.,

$$W[X_1] = R_1 \qquad W[X_2] = R_2 \qquad \cdots \qquad W[X_m] = R_m$$

We say that W is a witness of their consistency.

- pairwise consistent: any two are consistent
- k-wise consistent: any k are consistent
- globally consistent: all together are consistent

Part III

Inner Consistency

A weaker form of consistency

Definition [AK'24]
Two $\mathbb{K}$-relations $R(X)$ and $S(Y)$ are called inner consistent if

$$R[X \cap Y] = S[X \cap Y].$$

Fact. For every positive commutative monoid:

$$\text{consistency} \implies \text{inner consistency}$$

Indeed:
Let $R(X)$ and $S(Y)$ be $\mathbb{K}$-relations.
Let $Z = X \cap Y$ be the common attributes.
Let $W(X, Y)$ witness consistency.
Then:

$$R[Z] = W[X][Z] = W[Z] = W[Y][Z] = S[Z].$$

Monoids and inner consistency

Question. For which positive commutative monoids

$$\text{inner consistency} \overset{?}{\Longrightarrow} \text{consistency}$$

Coming up:

- Boolean monoid:	YES	the standard join
- Powerset monoid:	YES	intersections of annotations
- Tropical/cost monoid:	YES	maxima(!) of annotations
- Real-valued measures:	YES	normalized volume
- Bag monoid:	YES!	flow theory
- POVM:	NO.	

plus

a characterization and its consequences

From inner consistency to consistency by solving equations

$R(X, Y)$
$a_1\ c_1 : \alpha_1$
$a_2\ c_1 : \alpha_2$
$a_3\ c_1 : \alpha_3$

$S(Y, Z)$
$c_1\ d_1 : \beta_1$
$c_1\ d_2 : \beta_2$

$W(X, Y, Z)$
$a_1\ c_1\ d_1 :\ x_{11}?$
$a_1\ c_1\ d_2 :\ x_{12}?$
$a_2\ c_1\ d_1 :\ x_{21}?$
$a_2\ c_1\ d_2 :\ x_{22}?$
$a_3\ c_1\ d_1 :\ x_{31}?$
$a_3\ c_1\ d_2 :\ x_{32}?$

The inner consistency assumption is, in this case,

$$\alpha_1 + \alpha_2 + \alpha_3 = \beta_1 + \beta_2$$

The consistency witness is, in this case, any solution to the system

$$\begin{aligned} x_{11} + x_{12} &= \alpha_1 \\ x_{21} + x_{22} &= \alpha_2 \\ x_{31} + x_{32} &= \alpha_3 \end{aligned} \qquad \begin{aligned} x_{11} + x_{21} + x_{31} &= \beta_1 \\ x_{12} + x_{22} + x_{32} &= \beta_2 \end{aligned}$$

Instances of the Transportation Problem **TP(m,n)**

Given $\alpha_1, \ldots, \alpha_m$ **and** $\beta_1, \ldots, \beta_n$ **such that**

$$\alpha_1 + \cdots + \alpha_m = \beta_1 + \cdots + \beta_n$$

find x_{ij} **such that**

$$\begin{array}{ccccccccc}
x_{11} & + & x_{12} & + & \cdots & + & x_{1n} & = & \alpha_1 \\
+ & & + & & & & + & & \\
x_{21} & + & x_{22} & + & \cdots & + & x_{2n} & = & \alpha_2 \\
+ & & + & & & & + & & \\
\vdots & & \vdots & & \ddots & & \vdots & & \\
+ & & + & & & & + & & \\
x_{m1} & + & x_{m2} & + & \cdots & + & x_{mn} & = & \alpha_m \\
\shortparallel & & \shortparallel & & & & \shortparallel & & \\
\beta_1 & & \beta_2 & & & & \beta_n & &
\end{array}$$

Characterization

Characterization Theorem [AK'24].
For every positive commutative monoid $\mathbb{K}$, the following statements are equivalent:

(1) Every two $\mathbb{K}$-relations that are inner consistent are consistent.
(2) Every instance of TP over $\mathbb{K}$ is feasible.
(3) Every instance of TP(2,2) over $\mathbb{K}$ is feasible.

We say that $\mathbb{K}$ has the inner consistency property
We say that $\mathbb{K}$ has the transportation property

Indeed:
(1) $\Longleftrightarrow$ (2) : done; see two slides back from this one.
(2) $\Longleftrightarrow$ (3) : known from the theory of weighted automata [S'07]; see also two slides forward from this one.

Transportation property failling: an example

$\mathbb{N}_q$ = bag monoid with addition truncated to q, for $q \geq 2$.

It's a positive commutative monoid.
The precondition

$$1 + (q-1) = 1 + q \qquad (\text{in } \mathbb{N}_q)$$

holds, BUT the following system is infeasible:

$$\begin{array}{ccccc} x_{11} & + & x_{12} & = & 1 \\ + & & + & & \\ x_{21} & + & x_{22} & = & q-1 \\ \shortparallel & & \shortparallel & & \\ 1 & & q & & \end{array}$$

It suffers from the "short blanket dilemma" at x_{22}; i.e.,
By Row 1 & Col 2 we need $x_{22} \geq q - 1$.
By Col 1 & Row 2 we need $x_{22} \leq q - 2$.

Reduction from $(m \times n)$ to $(m \times 2)$, then to (2×2)

1. Set $\beta = \beta_1 + \cdots + \beta_{n-1}$.
2. Split into two systems (variables $y_1, \ldots, y_m$ are new):

$$
\begin{array}{ccccccc}
y_1 & + & x_{1n} & = & \alpha_1 \\
+ & & + & & \\
\vdots & & \vdots & & \\
+ & & + & & \\
y_m & + & x_{mn} & = & \alpha_m \\
\shortparallel & & \shortparallel & & \\
\beta & & \beta_n & &
\end{array}
\qquad
\begin{array}{ccccccc}
x_{11} & + & \cdots & + & x_{1(n-1)} & = & y_1 \\
+ & & & & + & & \\
\vdots & & \ddots & & \vdots & & \\
+ & & & & + & & \\
x_{m1} & + & \cdots & + & x_{m(n-1)} & = & y_m \\
\shortparallel & & & & \shortparallel & & \\
\beta_1 & & & & \beta_{n-1} & &
\end{array}
$$

3. Recurse.

Solving 2×2 instances in special cases

$$\alpha_1 + \alpha_2 = \gamma = \beta_1 + \beta_2$$

$$\begin{array}{ccccc} x_{11} & + & x_{12} & = & \alpha_1 \\ + & & + & & \\ x_{21} & + & x_{22} & = & \alpha_2 \\ \shortparallel & & \shortparallel & & \\ \beta_1 & & \beta_2 & & \end{array}$$

Boolean monoid/powerset/tropical/... distributive lattices:

$$(\alpha_i \wedge \beta_1) \vee (\alpha_i \wedge \beta_2) = \alpha_i \wedge (\beta_1 \vee \beta_2) = \alpha_i \wedge (\alpha_1 \vee \alpha_2) = \alpha_i$$

Real-valued measures/tropical semiring/... semifields:

$$\alpha_i\beta_1/\gamma + \alpha_i\beta_2/\gamma = \alpha_i(\beta_1 + \beta_2)/\gamma = \alpha_i\gamma/\gamma = \alpha_i$$

Solving 2×2 transportation for bag monoid

Bag monoid 2×2 instance:

$$
\begin{array}{ccccccccc}
x_{11} & + & x_{12} & = & \alpha_1 & & & & \\
+ & & + & & & & & & \\
x_{21} & + & x_{22} & = & \alpha_2 & & & & \\
\shortparallel & & \shortparallel & & & & & & \\
\beta_1 & & \beta_2 & & \gamma & = & \alpha_1 + \alpha_2 = \beta_1 + \beta_2 & &
\end{array}
$$

Set:

$$
\begin{array}{rcll}
x_{11} & = & \min(\alpha_1, \beta_1) & \text{totally ordered} \\
x_{12} & = & \alpha_1 - x_{11} & \text{non-negative!} \\
x_{21} & = & \beta_1 - x_{11} & \text{non-negative!} \\
x_{22} & = & \gamma - \max(\alpha_1, \beta_1) & \text{non-negative!}
\end{array}
$$

Unrolling the induction gives the Northwest Corner Method from the theory of linear programming [AK'24].

Part IV

Local vs Global Consistency

Vorobe'v and BFMY Theorems

Vorobe'v Theorem. [V'68]
For all collections $X_1, \ldots, X_m$ of sets of random variables, TFAE:

(1) $X_1, \ldots, X_m$ forms a regular simplicial complex.
(2) Every collection of probability measures on $X_1, \ldots, X_m$ that is pairwise consistent is consistent.

Beeri-Fagin-Maier-Yannakakis Theorem. [BFMY'83]
For all collections $X_1, \ldots, X_m$ of sets of attributes, TFAE:

(1) $X_1, \ldots, X_m$ is the set of edges of an acyclic hypergraph.
(2) Every collection of relations over $X_1, \ldots, X_m$ that is pairwise consistent is consistent.

Hypergraph acyclicity: a database theory classic

1. Acyclicity was introduced in the BFMY paper

"On the Desirability of Acyclic Schemes"

with *many different* equivalent characterizations.

2. The equivalent concept of join-tree is contemporary to Roberston and Seymour's tree-width and tree-decompositions of Graph Minors I/II (early 80's).

3. It is a key component in Yannakakis' (1981) fundamental join-tree algorithm for conjunctive query evaluation.

4. Non-trivially generalizes graph acyclicity [F'83].

$$\text{Berge acyclic} < \gamma\text{-acyclic} < \beta\text{-acyclic} < \alpha\text{-acyclic}$$

Generalizing Vorobe'v and BFMY Theorems

Theorem. [AK'24]
Let $\mathbb{K}$ be a positive commutative monoid that has the transportation property. For all collections $X_1, \ldots, X_m$ of sets of attributes, TFAE:

(1) $X_1, \ldots, X_m$ is the set of edges of an acyclic hypergraph.
(2) $X_1, \ldots, X_m$ has the local-to-global (L2G) property on $\mathbb{K}$; i.e., every collection of $\mathbb{K}$-relations on $X_1, \ldots, X_m$ that is pairwise consistent is consistent.

Corollary. BFMY acyclicity and Vorobe'v regularity coincide.

[A direct proof of corollary is also doable... and more natural.]

Proof of necessity : L2G implies acyclicity (1/2)

Note: The proofs in BFMY and V do not generalize (at all).

Theorem [AK'24]
If H_0 is a d-regular & k-uniform hypergraph with $d \geq 2$ and $k \geq 2$, then H_0 fails L2G on $\mathbb{K}$.

Theorem (reformulated from [BFMY'83])
If H is non-acyclic, then $C_n \leq_{\mathsf{G}} H$ or $S_n \leq_{\mathsf{G}} H$ for some $n \geq 3$.

Lemma [AK'24]
If $H_0 \leq_{\mathsf{G}} H$ and H_0 fails L2G on $\mathbb{K}$, then H fails L2G on $\mathbb{K}$.

Fact
C_n is 2-regular and 2-uniform.
S_n is $(n-1)$-regular and $(n-1)$-uniform.

Proof of necessity : L2G implies acyclicity (2/2)

Assume $H_0 = \{X_1, \ldots, X_m\}$ is d-regular and k-uniform.
Fix arbitrary $\ell : \{1, \ldots, m\} \to \mathbb{Z}/d\mathbb{Z}$ such that

$$\ell(1) + \cdots + \ell(m) \not\equiv 0 \bmod d.$$

Such a labelling ℓ exists if $d \geq 2$ and $m \geq 1$.

Fix arbitrary $\alpha^* \in K \setminus \{0\}$.
Define $R_i(X_i) : (\mathbb{Z}/d\mathbb{Z})^k \to K$ by

$$R_i(a_1, \ldots, a_k) := d^k \cdot \alpha^* \quad \text{iff} \quad a_1 + \cdots + a_k \equiv \ell(i) \mod d$$

where $d^k \cdot \alpha^* := \alpha^* + \cdots + \alpha^*$ (d^k times).

pairwise consistent: by $k \geq 2$ and uniformity of $\mathbb{Z}/d\mathbb{Z}$-subspaces.
globally inconsistent: by d-regularity and $\sum_{i=1}^{m} \ell(i) \not\equiv 0 \bmod d$.

Proof of sufficiency : acyclicity implies L2G (1/1)

Assume $\mathbb{K}$ has the transportation property.
Then:

Yannakakis join-tree algorithm
specialized to full conjunctive queries
works also for $\mathbb{K}$-relations

when

inner consistency $\Longrightarrow$ consistency

which, here, is the case by the Characterization Theorem.

Part V

Back to Bell

Transportation property failing for POVMs

Let's argue that the analysis of Bell Inequalities ([B'64]) leads to:

Fact.
For all $d \geq 2$, the positive commutative monoid $\mathbb{R}_{\succeq}^{d\times d}$ of POVM with component-wise $+$ does not have the transportation property.

Recall:
$\mathbb{R}_{\succeq}^{d\times d}$: the set of positive semi-definite (PSD) matrices $M \in \mathbb{R}^{d\times d}$.
PSD: symmetric and such that $z^{\mathrm{T}}Mz \geq 0$ holds for all $z \in \mathbb{R}^{d}$.

The counterexample

The matrices

$$B_1 = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \qquad C_1 = \begin{bmatrix} 1/2 & -1/2 \\ -1/2 & 1/2 \end{bmatrix}$$

$$B_2 = \begin{bmatrix} 0 & 0 \\ 0 & 1 \end{bmatrix} \qquad C_2 = \begin{bmatrix} 1/2 & 1/2 \\ 1/2 & 1/2 \end{bmatrix}$$

are PSD, they satisfy

$$B_1 + B_2 = C_1 + C_2$$

BUT the following system is infeasible in PSD matrices:

$$\begin{array}{ccccc} X_{11} & + & X_{12} & = & B_1 \\ + & & + & & \\ X_{21} & + & X_{22} & = & B_2 \\ \shortparallel & & \shortparallel & & \\ C_1 & & C_2 & & \end{array}$$

The certificate of infeasibility is a Bell inequality

Set

$$A_1 = \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \qquad A_2 = \begin{bmatrix} -1 & 1 \\ 1 & 1 \end{bmatrix}$$

and

$$M := A_1 X_{11} + A_2 X_{12} - A_1 X_{21} - A_2 X_{22}.$$

Using the assumption that the X_{ij} satisfy the system we have:

$$\begin{aligned} (1) \quad & \mathrm{tr}(M) = \mathrm{tr}(2\mathrm{I}) = 4 \\ (2) \quad & \mathrm{tr}(M) \leq (\textstyle\sum_{ij} \mathrm{tr}(X_{ij}))(\max_i \|A_i\|) \leq \mathrm{tr}(\mathrm{I})\sqrt{2} = 2\sqrt{2} \end{aligned}$$

I.e., A Bell inequality fails;
a "quantum short blanket dilemma" of sorts.

Part VI

CONCLUSIONS

Key ideas and findings

Key ideas:
- Following recent trend (e.g., provenance), data comes annotated.
- In the study of consistency, positive monoids are enough.

Key finding 1:
- Inner consistency of data is necessary for its coherent existence.
- But it is not always sufficient: see $\mathbb{N}_q$, and Bell scenarios.

Key finding 2:
- Sufficiency of locality is ensured by the transportation property.
- The sufficiency extends to all acyclic scenarios.
- It does not extend to any non-acyclic scenario whatsoever.

An open-ended question

Conventional wisdom has been that data ought to be globally consistent – reminiscent of “an element of physical reality”.

BUT IT IS RARELY ENFORCED!

Are there computational or information-theoretic advantages in explicitly giving up on global consistency of data?

[If witnessed, quantum advantage may be an answer, but maybe not the only answer]

Maybe not the only answer...

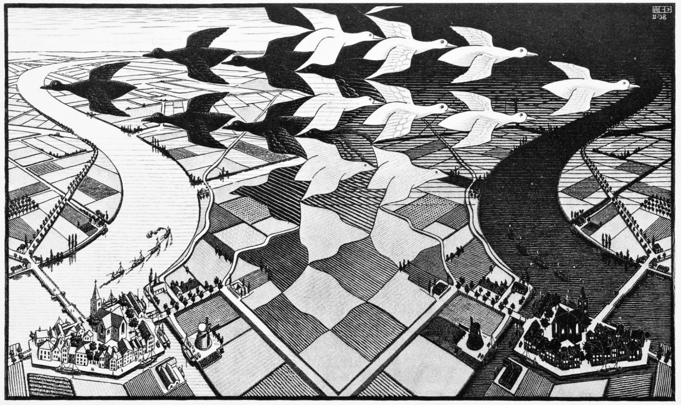

Day and Night, woodcut, Escher 1938.

[Low resolution image downloaded from Wikipedia]

Three references

[AK'21] A. Atserias and Ph. G. Kolaitis. *Structure and Complexity of Bag Consistency*. PODS'21.

[AK'24/25] A. Atserias and Ph. G. Kolaitis. *Consistency of Relations over Monoids*. Journal of the ACM, Volume 72, Issue 3 Article No.: 18, Pages 1 - 47, 2025. Preliminary version in PODS'24.

[AK'25] A. Atserias and Ph. G. Kolaitis. *Consistency Witnesses for Annotated Relations*. To appear in SIGMOD Record 2025.