how to count structure

thomas barrett

introduction

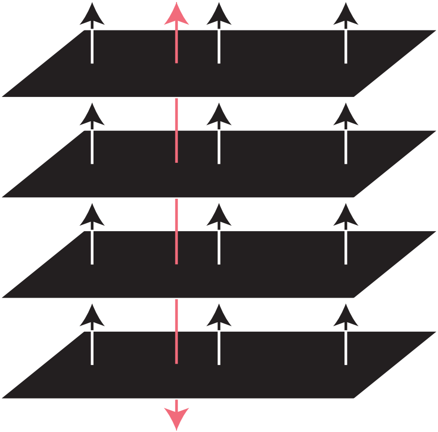

Aristotelian spacetime

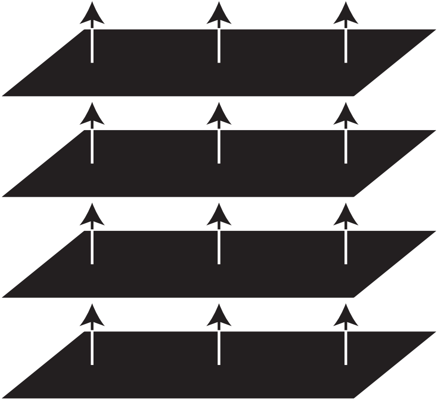

Newtonian spacetime

Galilean spacetime

Conclusion 1. Each of these classical spacetime theories *posits less structure*, or *ascribes less structure to the world*, than its predecessors. Galilean spacetime, for example, is obtained by 'taking something away' — the concept of absolute rest — from Newtonian spacetime.

Two reasons I care about this.

Conclusion 2. Each of these classical spacetime theories is *not equivalent* to its predecessors. Aristotelian and Newtonian spacetime, for example, disagree about whether or not there is a center of the universe.

Structural parsimony. All other things equal, we should prefer theories that posit less structure.

North: “This is a principle informed by Ockham’s razor; though it is not just that, other things being equal, it is best to go with the ontologically minimal theory. It is not that, other things being equal, we should go with the fewest entities, but that we should go with the least structure.”

Sider: “*structurally simpler* theories are more likely to be true”.

Geroch (!): “Although the evidence on this is perhaps a bit scanty, it seems to be the case that physics, at least in its fundamental aspects, always moves in this one direction. It may not be a bad rule of thumb to judge a new set of ideas in physics by the criterion of how many of the notions and relations that one feels to be necessary one is forced to give up.”

Structural parsimony. All other things equal, we should prefer theories that posit less structure.

The Onion Thesis: A physical theory is like an onion. It has layers of structure; some play important roles in the theory, and others are redundant or superfluous or surplus to the theory. One of the distinctive aims of philosophy of physics is to 'peel away' layers of the latter kind.

Chen: excising structure from a theory "eliminates the need for a large class of arbitrary conventions [...and] in the absence of these arbitrary conventions, we can look directly into the real structure of the [...] objects without worrying that we are looking at some merely representational artifact".


New Brunswick Voted No. 6 Most Exciting Place in NJ!
- as reported by Movoto.com
(click image for article)

outline

- the automorphism approach
- the category approach
- onions

the automorphism approach

Some (very recent) history.

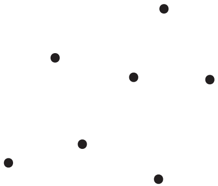

X

$<$

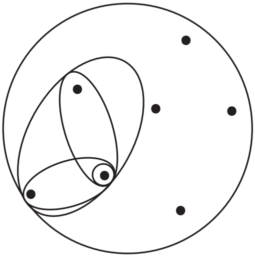

(X, τ)

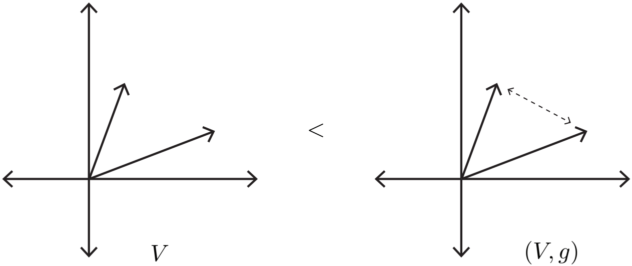
V
$<$
(V, g)

And there are many other examples. . .

SYM: A mathematical object X has more structure than a mathematical object Y if and only if the automorphism group $\mathrm{Aut}(X)$ is "smaller than" the automorphism group $\mathrm{Aut}(Y)$.

SYM*: A mathematical object X has more structure than a mathematical object Y if and only if $\mathrm{Aut}(X) \subsetneq \mathrm{Aut}(Y)$.

Three arguments for SYM*:

- the argument from examples
- the argument from size
- the argument from definability

The argument from examples: SYM* makes intuitive verdicts in many easy cases of structural comparison.

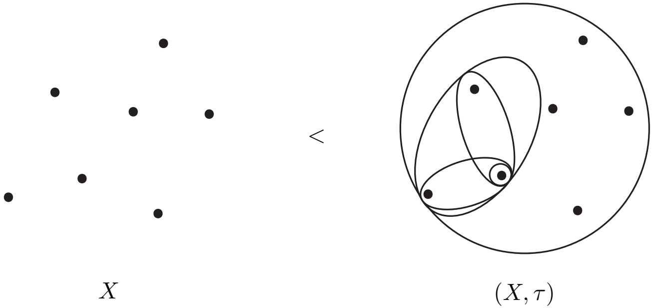
X
$<$
(X, τ)

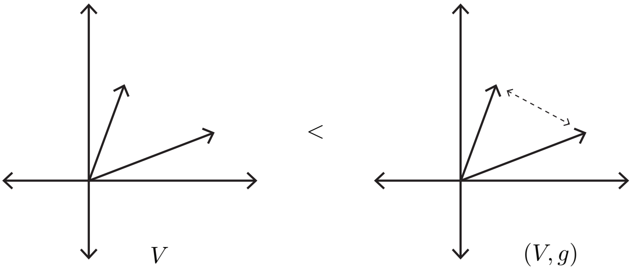
V
<
(V, g)

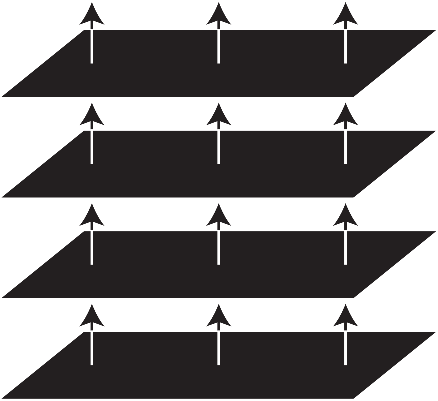

Newtonian spacetime

Galilean spacetime

The fact that SYM* captures some simple examples speaks in favor of the criterion, but it is not entirely convincing. . .

The argument from size: Since automorphisms are structure-preserving maps from an object to itself, if $\mathrm{Aut}(X) \subsetneq \mathrm{Aut}(Y)$, this means that X has *fewer* automorphisms than Y, which suggests that X has *more* structure that these automorphisms are required to preserve.

The argument from definability.

Desideratum. A mathematical object X has more structure than a mathematical object Y if and only if X has all of the structures that Y has and X has some structure that Y lacks.

A topological space (X, τ) vs a metric space (X, d).

Desideratum. A mathematical object X has more structure than a mathematical object Y if and only if X can define all of the structures that Y has, but X has some piece of structure that Y does not define.

Question. Is it the case that X has more structure than Y according to SYM* if and only if X can define all of the structures that Y has, but X has some piece of structure that Y does not define?

The basic set-up:

- Let Σ_1 and Σ_2 be signatures. We will think of the elements of Σ_1 and Σ_2 as the encoding the 'basic structures' on the two objects that we will consider.
- Let A be a Σ_1-structure and B a Σ_2-structure with the same underlying set. We will think of A and B as the two objects whose structures will we be comparing.
- Let $p \in \Sigma_2$ be one of the basic structures on B.

We say that the Σ_1-structure A **explicitly defines** p^B if there is a Σ_1-formula ϕ such that $\phi^A = p^B$.

We say that the Σ_1-structure A **implicitly defines** p^B if $h[p^B] = p^B$ for every automorphism $h : A \to A$ of A.

If A explicitly defines p^B, then A implicitly defines p^B. The converse does not hold.

Proposition 1. The following are equivalent:

1. For every symbol $p \in \Sigma_2$, A implicitly defines p^B, but there is a $q \in \Sigma_1$ such that B does not implicitly define q^A.
2. $\mathrm{Aut}(A) \subsetneq \mathrm{Aut}(B)$

Question. Is it the case that X has more structure than Y according to SYM* if and only if X can define all of the structures that Y has, but X has some piece of structure that Y does not define?

Answer. Yes, if we mean implicit definability.

Two problems with SYM*:

- sensitivity
- triviality

Sensitivity: Let (X, τ) be a topological space and Y a set that is not equal to X.

Triviality: When an object X has a trivial automorphism group, there is no object that has more structure than X.

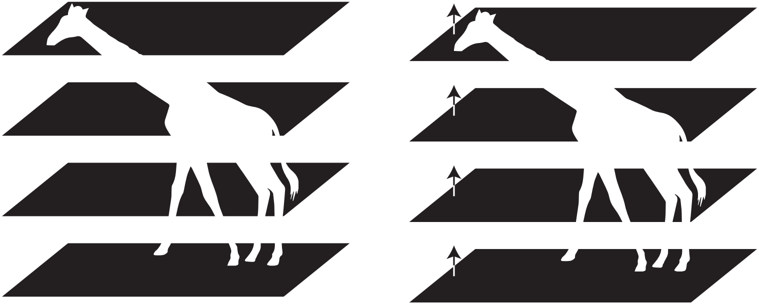

Malament's giraffe

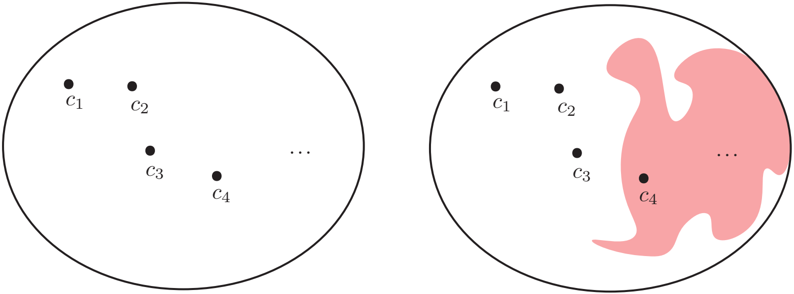

a $\{c_1, c_2, \ldots\}$-structure A vs a $\{p, c_1, c_2, \ldots\}$-structure B

Question. Is it the case that X has more structure than Y according to SYM* if and only if X can define all of the structures that Y has, but X has some piece of structure that Y does not define?

Answer. No, if we mean explicit definability.

the category approach

More (very recent) history.

SYM* answers the question *when does a mathematical object X have more structure than a mathematical object Y?*

The category approach answers the question *when does one type of mathematical object have more structure than another type of mathematical object?* or *when does one theory posit more structure than another theory?*

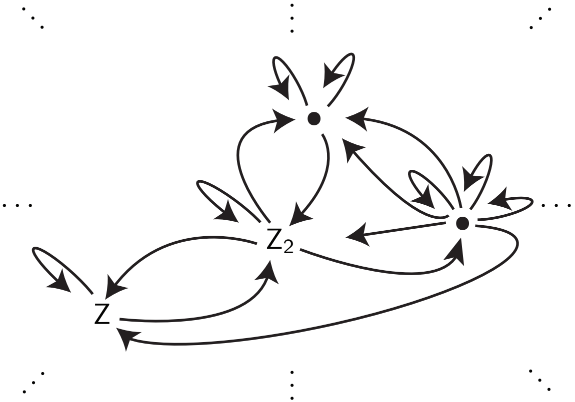

a category

A functor is a structure-preserving map between categories.

A functor $F : C \to D$ is **full** if for all objects c_1, c_2 in C and arrows $g : Fc_1 \to Fc_2$ in D there exists an arrow $f : c_1 \to c_2$ in C with $Ff = g$.

The Baez et al. method. We say that a functor $F : C \to D$ **forgets structure** if it is not full. This captures a sense in which (relative to the comparison generated by F) objects in C have more structure than objects in D.

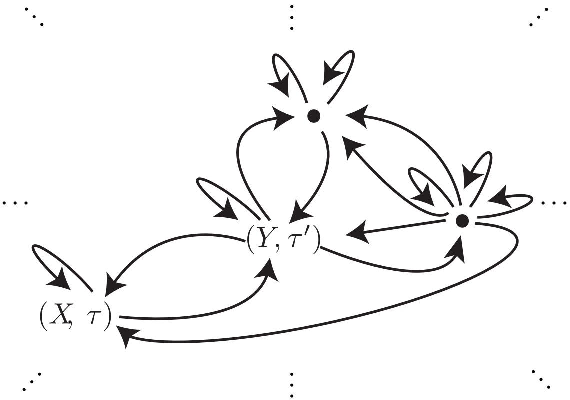

The category of topological spaces

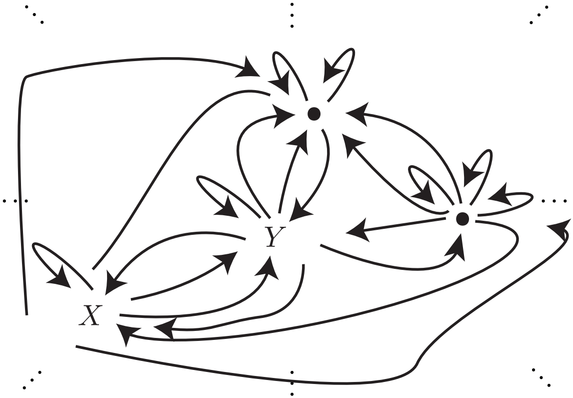

The category of sets

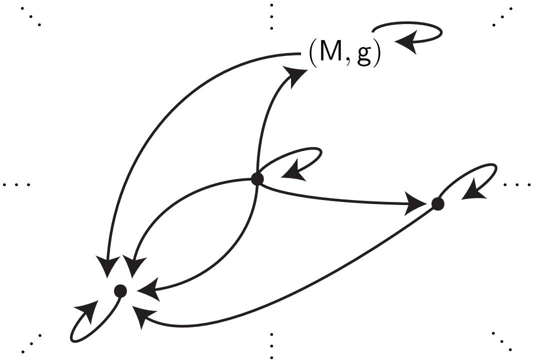

(1-3) general relativity

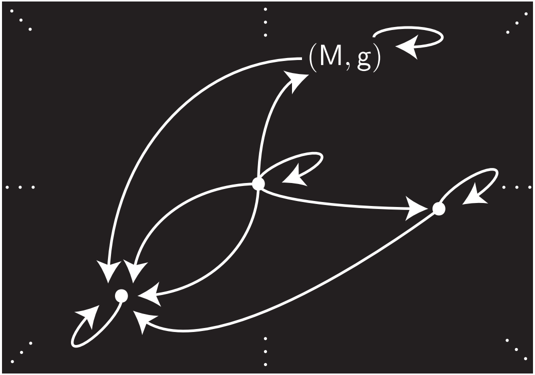

(3-1) general relativity

Three arguments for the Baez method:

- the argument from examples
- the argument from size
- the argument from definability

The argument from examples: The Baez method makes intuitive verdicts in many easy cases of structural comparison.

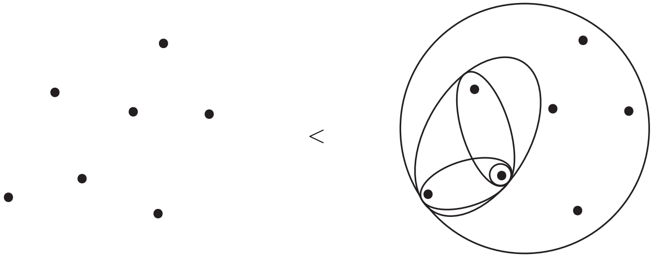

$X \qquad < \qquad (X, \tau)$

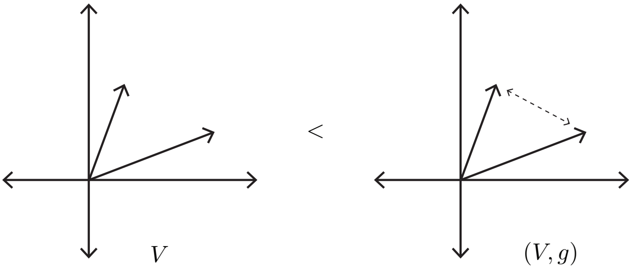
V
$<$
(V, g)

The argument from size: Since arrows in a category are structure-preserving maps between objects, if $F : C \to D$ is not full, this means that objects of C have *fewer* arrows between them than objects of D, which suggests that the former have *more* structure that these arrows are required to preserve.

The argument from definability.

A Σ-**theory** T is a set of Σ-sentences.

A Σ-structure is a **model** of the Σ-theory T if $M \models \phi$ for all $\phi \in T$.

A Σ-theory T has a **category of models** $\text{Mod}(T)$.

Question. Suppose that we have two theories T_1 and T_2 and a functor $F : \mathrm{Mod}(T_1) \to \mathrm{Mod}(T_2)$ between their categories of models. Is it the case that F forgets structure — thus capturing a sense in which T_1 posits more structure than T_2 — if and only if T_2 fails to define all of the structures that T_1 posits?

Consider the two signatures $\Sigma_1 = \{p_0, p_1, p_2, \ldots\}$ and $\Sigma_2 = \{q_0, q_1, q_2, \ldots\}$, each of which have a countable infinity of unary predicate symbols. We define the Σ_1-theory T_1 and the Σ_2-theory T_2 as follows.

$$T_1 = \{\exists_{=1} x(x = x)\}$$
$$T_2 = \{\exists_{=1} y(y = y), \forall y(q_0(y) \rightarrow q_1(y)), \forall y(q_0(y) \rightarrow q_2(y)), \ldots\}$$

Question. Suppose that we have two theories T_1 and T_2 and a functor $F : \mathrm{Mod}(T_1) \to \mathrm{Mod}(T_2)$ between their categories of models. Is it the case that F forgets structure — thus capturing a sense in which T_1 posits more structure than T_2 — if and only if T_2 fails to define all of the structures that T_1 posits?

Answer. In general, no.

Question. Suppose that we have two theories T_1 and T_2 and a *nicely behaved* functor $F : \text{Mod}(T_1) \to \text{Mod}(T_2)$ between their categories of models. Is it the case that F forgets structure — thus capturing a sense in which T_1 posits more structure than T_2 — if and only if T_2 fails to define all of the structures that T_1 posits?

Proposition 2. Let T^+ be a Σ^+-theory that is an extension of the Σ-theory T. The functor $\Pi : \mathrm{Mod}(T^+) \to \mathrm{Mod}(T)$ forgets structure if and only if there is a symbol $r \in \Sigma^+$ such that there is no Σ-formula ϕ that satisfies $T^+ \vDash \forall x(r(x) \leftrightarrow \phi(x))$.

Question. Suppose that we have two theories T_1 and T_2 and a nicely behaved functor $F : \text{Mod}(T_1) \rightarrow \text{Mod}(T_2)$ between their categories of models. Is it the case that F forgets structure — thus capturing a sense in which T_1 posits more structure than T_2 — if and only if T_2 fails to define all of the structures that T_1 posits?

Answer. Yes, when F is a projection functor.

But one would like some more generality. Not every functor is a projection functor. . .

A **reconstrual** F of Σ_1 into Σ_2 is a map from elements of the signature Σ_1 to Σ_2-formulas.

A reconstrual $F : \Sigma_1 \to \Sigma_2$ is a **translation** of a Σ_1-theory T_1 into a Σ_2-theory T_2 if $T_1 \vDash \phi$ implies that $T_2 \vDash F\phi$ for all Σ_1-sentences ϕ.

(The existence of a translation $F : T_1 \to T_2$ captures a sense in which T_2 can define the structures of T_1.)

A translation $F : T_1 \to T_2$ naturally induces a functor $F^* : Mod(T_2) \to Mod(T_1)$.

A translation $F : T_1 \to T_2$ is **essentially surjective** if for every Σ_2-formula ψ there is a Σ_1-formula ϕ such that $T_2 \vDash \forall x_1 \dots \forall x_n(\psi(x_1, \dots, x_n) \leftrightarrow F\phi(x_1, \dots, x_n))$.

(If a translation $F : T_1 \to T_2$ is essentially surjective, that captures a sense in which T_1 also can define the structures of T_2.)

Proposition 3. Let T_1 be a Σ_1-theory and T_2 a Σ_2-theory with $F : T_2 \to T_1$ a translation. The following are equivalent:

1. F is essentially surjective.
2. $F^* : \mathrm{Mod}(T_1) \to \mathrm{Mod}(T_2)$ is full, i.e. does not forget structure.

Question. Suppose that we have two theories T_1 and T_2 and a nicely behaved functor $F : \text{Mod}(T_1) \to \text{Mod}(T_2)$ between their categories of models. Is it the case that F forgets structure — thus capturing a sense in which T_1 posits more structure than T_2 — if and only if T_2 fails to define all of the structures that T_1 posits?

Answer. Yes, when F is a functor induced by a translation.

Problems with the Baez method:

- triviality
- relativization to the functor

onions

Summing up.

Structural parsimony. All other things equal, we should prefer theories that posit less structure.

Chen writes, for example, that excising structure from a theory "eliminates the need for a large class of arbitrary conventions [. . . and] in the absence of these arbitrary conventions, we can look directly into the real structure of the [. . .] objects without worrying that we are looking at some merely representational artifact".

The Onion Thesis: A physical theory is like an onion. It has layers of structure; some play important roles in the theory, and others are redundant or superfluous or surplus to the theory. One of the distinctive aims of philosophy of physics is to ‘peel away’ layers of the latter kind.

Remark 1: Theories are not like normal onions.

Two mistakes one can make when trying to excise structure:

- excising not enough
- excising too much

Let $\Sigma = \{p, r\}$ where p is a binary predicate symbol and r is unary predicate symbol. Consider the Σ-theory

$$T = \{\forall x(r(x) \leftrightarrow p(x, x))\}$$

How do we excise the structure r from this theory?

Let $\Sigma = \{p, r\}$ where p is a binary predicate symbol and r is unary predicate symbol. Consider the Σ-theory

$$T = \{\forall x(r(x) \leftrightarrow p(x, x))\}$$

How do we excise the structure r from this theory?

- excising not enough

Let $\Sigma = \{p, r\}$ where p is a binary predicate symbol and r is unary predicate symbol. Consider the Σ-theory

$$T = \{\forall x(r(x) \leftrightarrow p(x, x))\}$$

How do we excise the structure r from this theory?

- excising not enough
- excising too much

Suppose we wanted to excise "straightness structure" — i.e. the derivative operator ∇ — from general relativity. How do we excise this?

Suppose we wanted to excise “straightness structure” — i.e. the derivative operator ∇ — from general relativity. How do we excise this?

- excising not enough

Suppose we wanted to excise "straightness structure" — i.e. the derivative operator ∇ — from general relativity. How do we excise this?

- excising not enough
- excising too much

Layers of the onion can be connected in interesting ways. And 'peeling away' a layer is not as simple as just reformulating the theory in such a way that the piece of structure is not explicitly referred to.

Thanks!