Inferentialist semantics and scientific modeling

Kris Brown - JMM 2024

TOPOS INSTITUTE

(press s for speaker notes)

Welcome

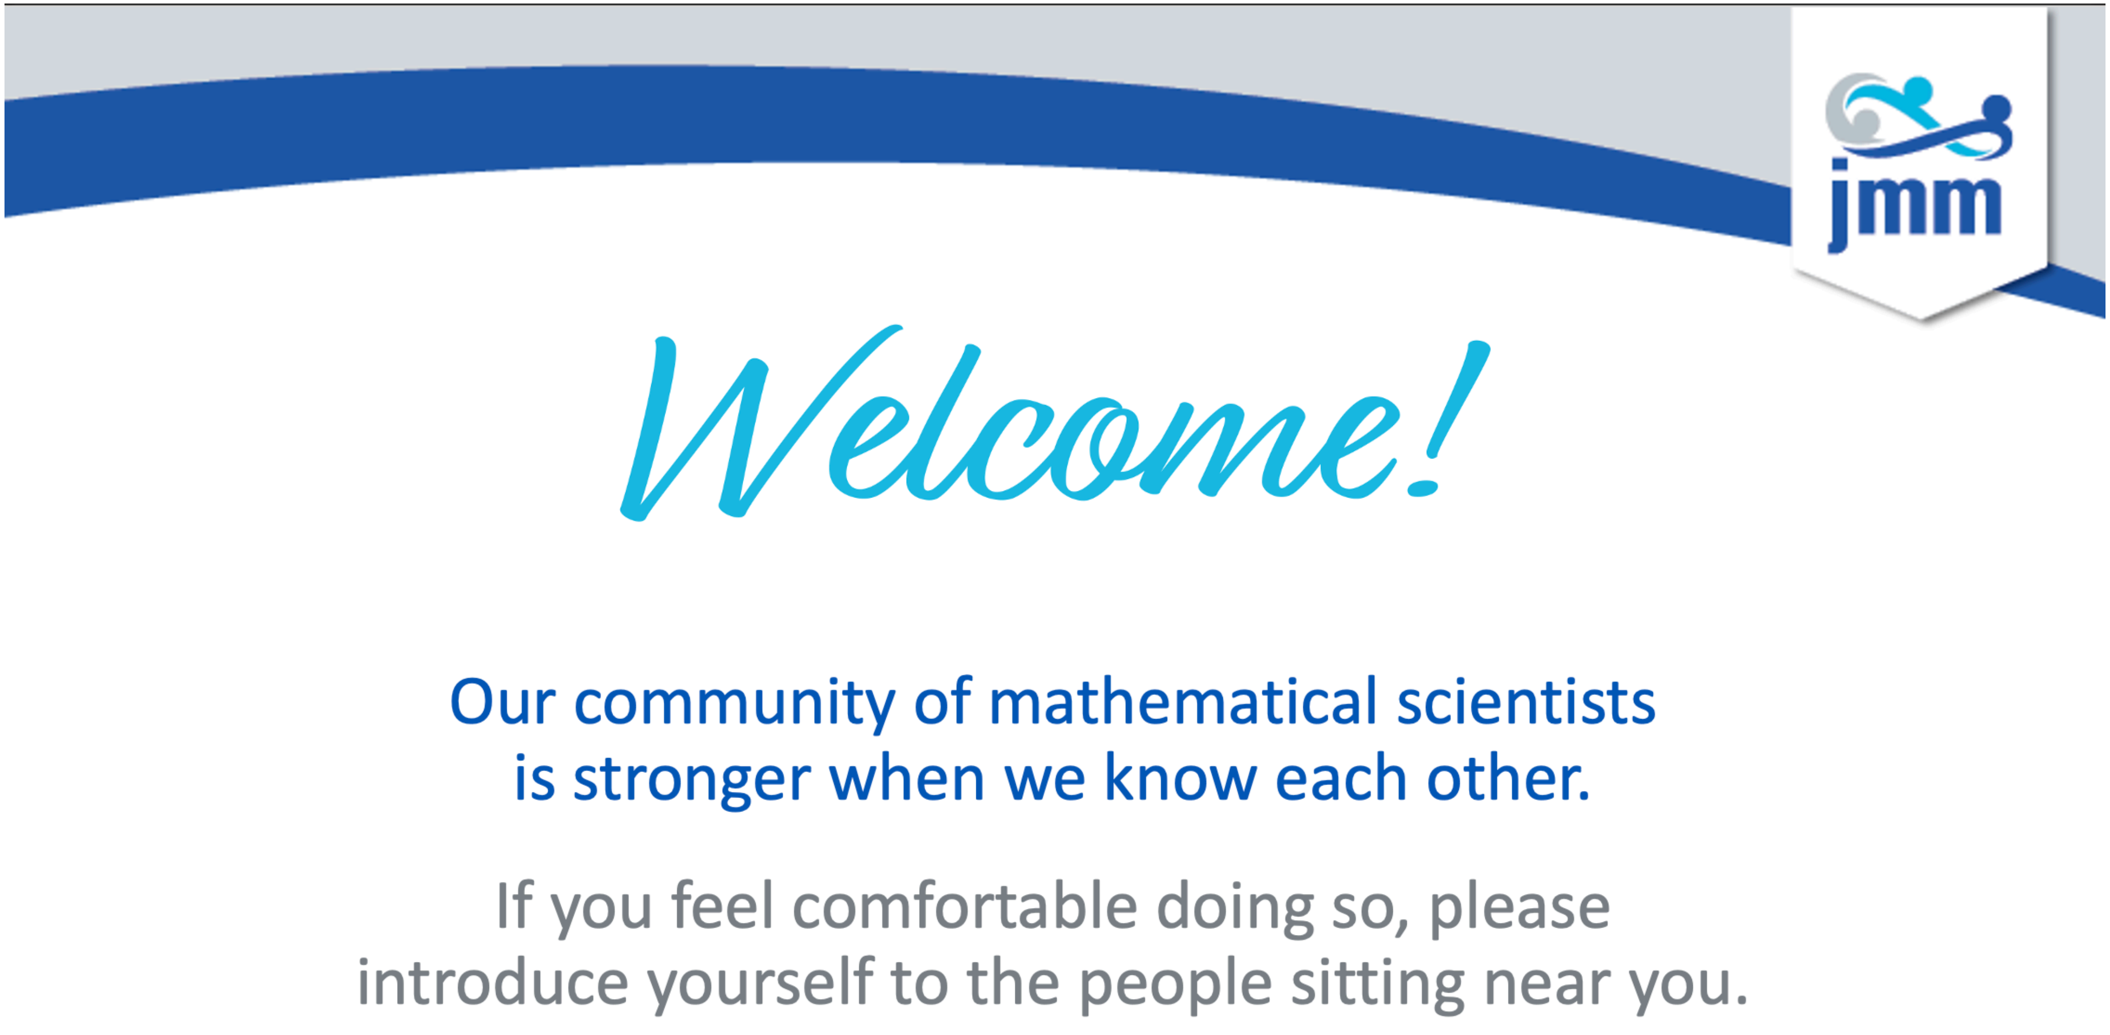
jmm
Welcome!
Our community of mathematical scientists
is stronger when we know each other.
If you feel comfortable doing so, please
introduce yourself to the people sitting near you.

Overview

- There are two competing paradigms for understanding linguistic meaning
  - Representationalism
  - Inferentialism
- These correspond to two paradigms of categorical semantics
  - Functorial semantics
  - Internal logic
- Both are valuable, but traditional thinking about natural language and scientific models is largely one-sided.
  - Software developed at the Topos Institute applies this other perspective.
  - This permits collaboration while demanding less agreement in background assumptions.

Representationalism: language is descriptive

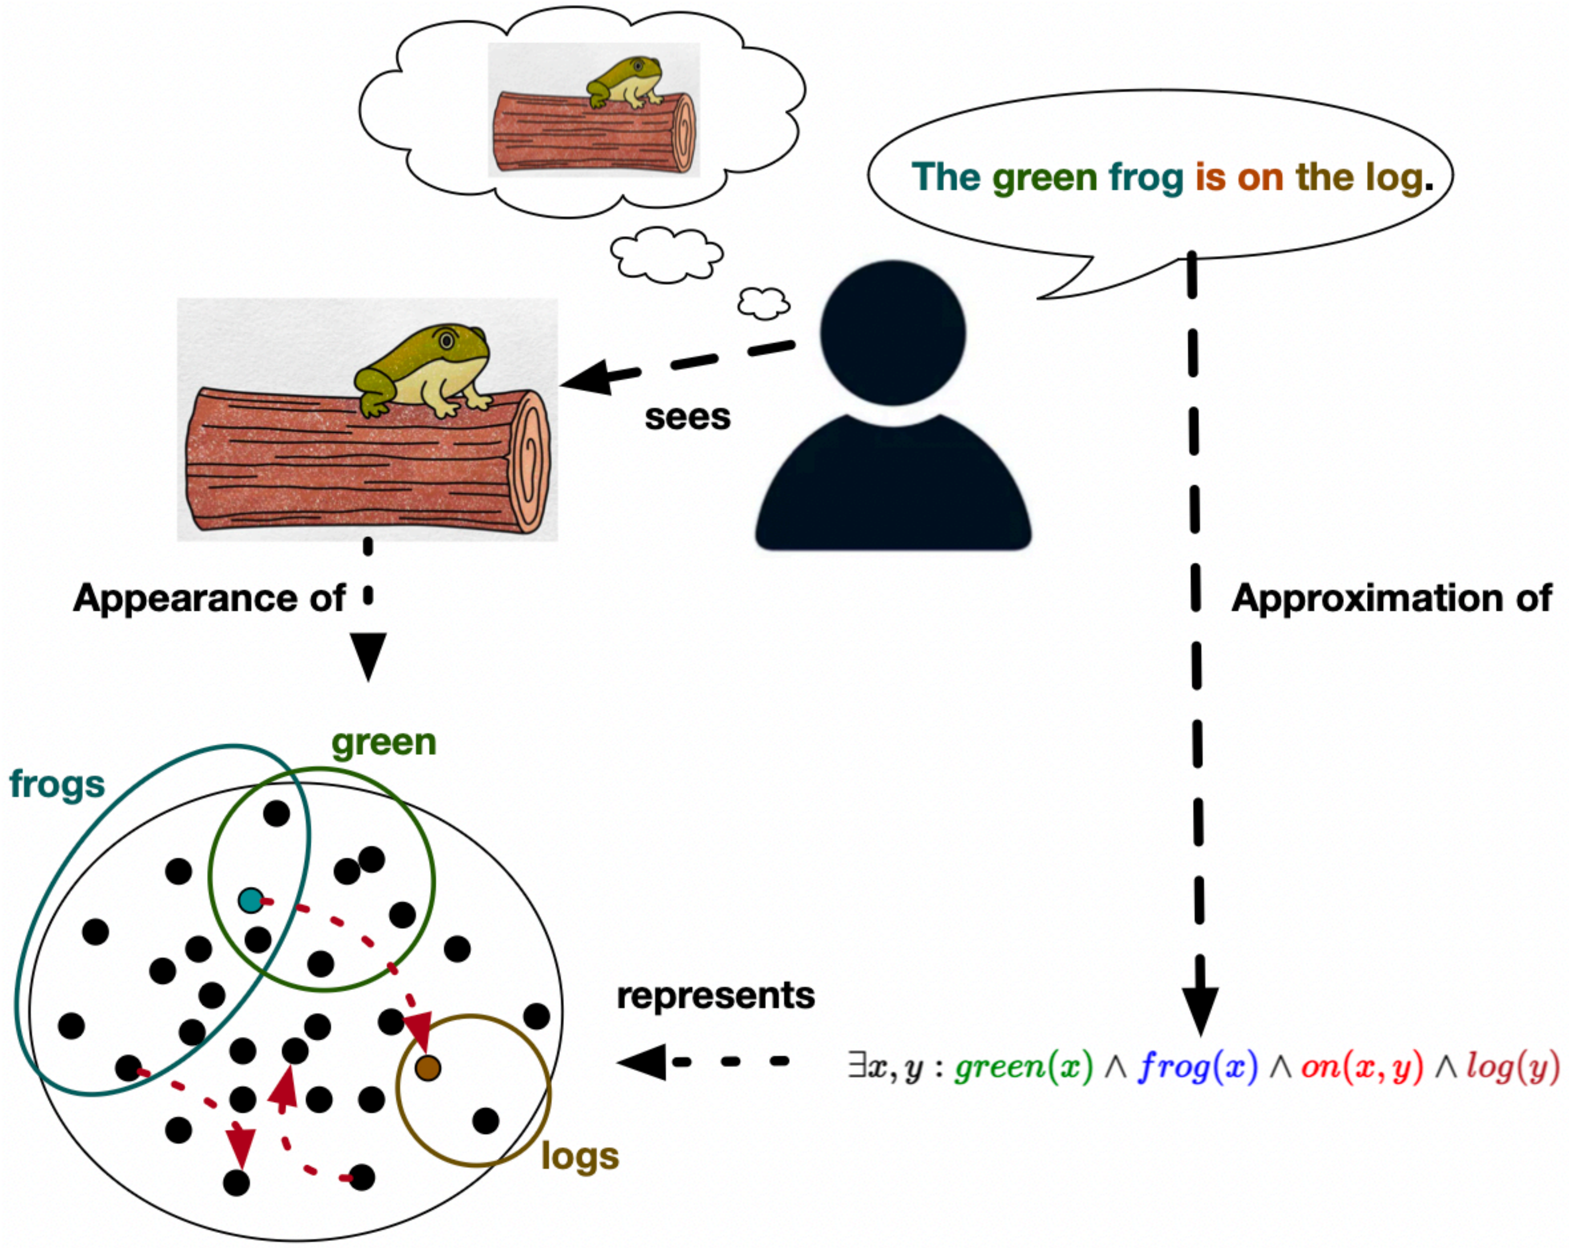

Representationalism: order of explanation

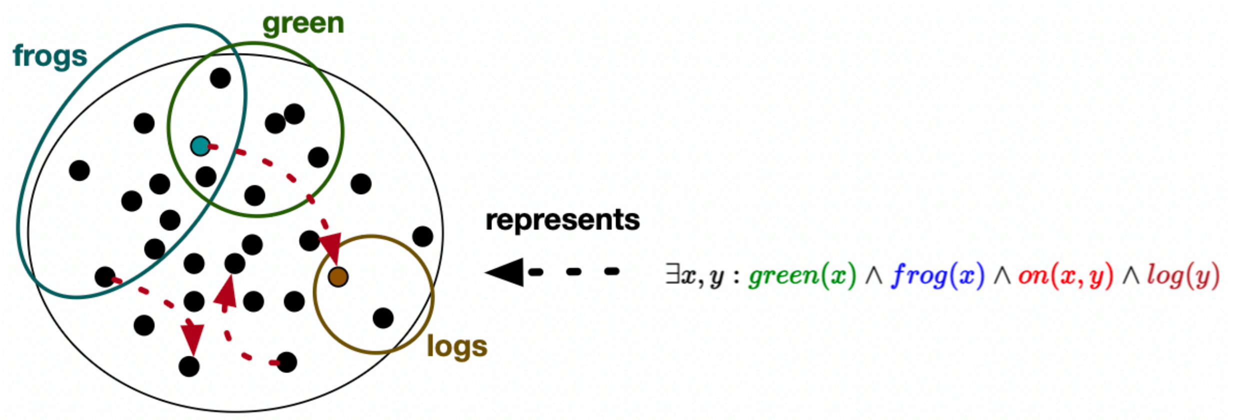

Bottom-up order of explanation

1. Taking the world to (just) be a certain way
2. … explains what our words mean,
3. … which explains what our sentences mean,
4. … which explains what inferences are good.

This leads to **atomistic** semantics:

- Each sentence is made true independent of the other sentences.
- Each meaning can be specified independent of each other meaning.

Representationalism: facts and values

This picture of meaning can serve as a criterion to identify when certain sentences shouldn't be thought of as actually contentful:

- "Twas brillig, and the slithy toves …" / "Colorless green ideas sleep furiously"
- "It's immoral to do that."

Understanding these statements as (covertly) meaningless explains why they aren't truth-apt, why we can't seem to methodically resolve debates or find definitions (e.g. ethics). We're *licensed* to stop trying to resolve these debates.

However, many important concepts in science and engineering are non-representational:

- Scientific methodologies of resolving debate via looking to 'the data' require common ontologies / practices of interpreting data
  - Interesting science occurs in the regions where this isn't true
- Choosing between theories depends on aesthetic norms (e.g. *beauty*, *simplicity*)
- Engineering decisions are frequently entangled with ethical concepts
- All of our concepts "have attitudes".

Inferentialism

Bottom-up (atomistic) order of explanation

1. Taking the world to (just) be a certain way.
2. … explains what our words mean.
3. … which explains what our sentences mean.
4. … which explains what inferences are good.

Top-down (holistic) order of explanation

1. Taking some inferences to (just) be good.
2. … explains what our sentences mean.
3. … which explains what our words mean.
4. … which explains what we're talking *about*.

Inferential takes the *use* of language to be conceptually prior to the *meaning*.

Sequent rules can express these usage patterns:

$$\wedge\text{-}I\ \frac{A \qquad B}{A \wedge B} \qquad\qquad \wedge\text{-}E\ \frac{A \wedge B}{A} \quad \frac{A \wedge B}{B}$$

To understand what *and* means doesn't require locating *and* in the world somewhere.

Mastery of the proper inferential moves of *and* suffice to grasp the concept.

Two kinds of category-theoretic semantics

CT since Lawvere has a tradition understanding the meaning of something via representation via various notions of **functorial semantics**: $\mathrm{Mod} : \mathsf{Syn} \to \mathsf{Sem}$

However, CT *also* comes with opinions about what things mean without saying what they represent. These are purely based on how these things are related to each other.

Universal properties characterize what it means to be a 'side-by-side placement' (coproduct $\mathrm{A} + \mathrm{B}$), or an 'element' $(1 \to \mathrm{X})$ or a 'subobject' $(\mathrm{X} \rightarrowtail \mathrm{Y})$, etc.

> **Duality of the approaches**
>
> Functorial semantics is characterized by maps *out* of a category
>
> Universal properties are characterized by maps *into* a category.

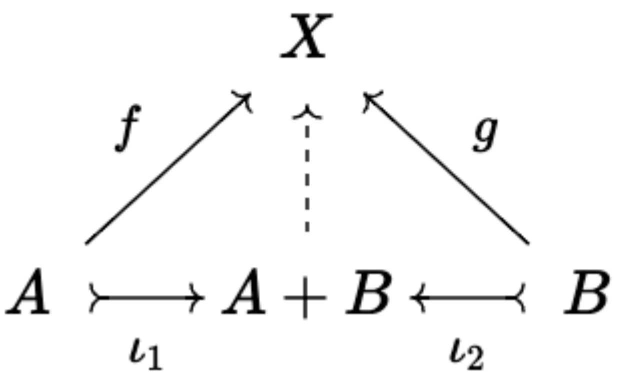

Examples

Example 1: chemical kinetics modeling

Imagine a very basic form of chemistry that develops around trying to make sense of changing quantities of substances.

The changes in quantities are attributed to "reactions", and an experiment is explained by constructing a model such as "$\mathrm{A} \to \mathrm{B}, \mathrm{B} \to \mathrm{C}, \mathrm{C} \to \mathrm{A}$".

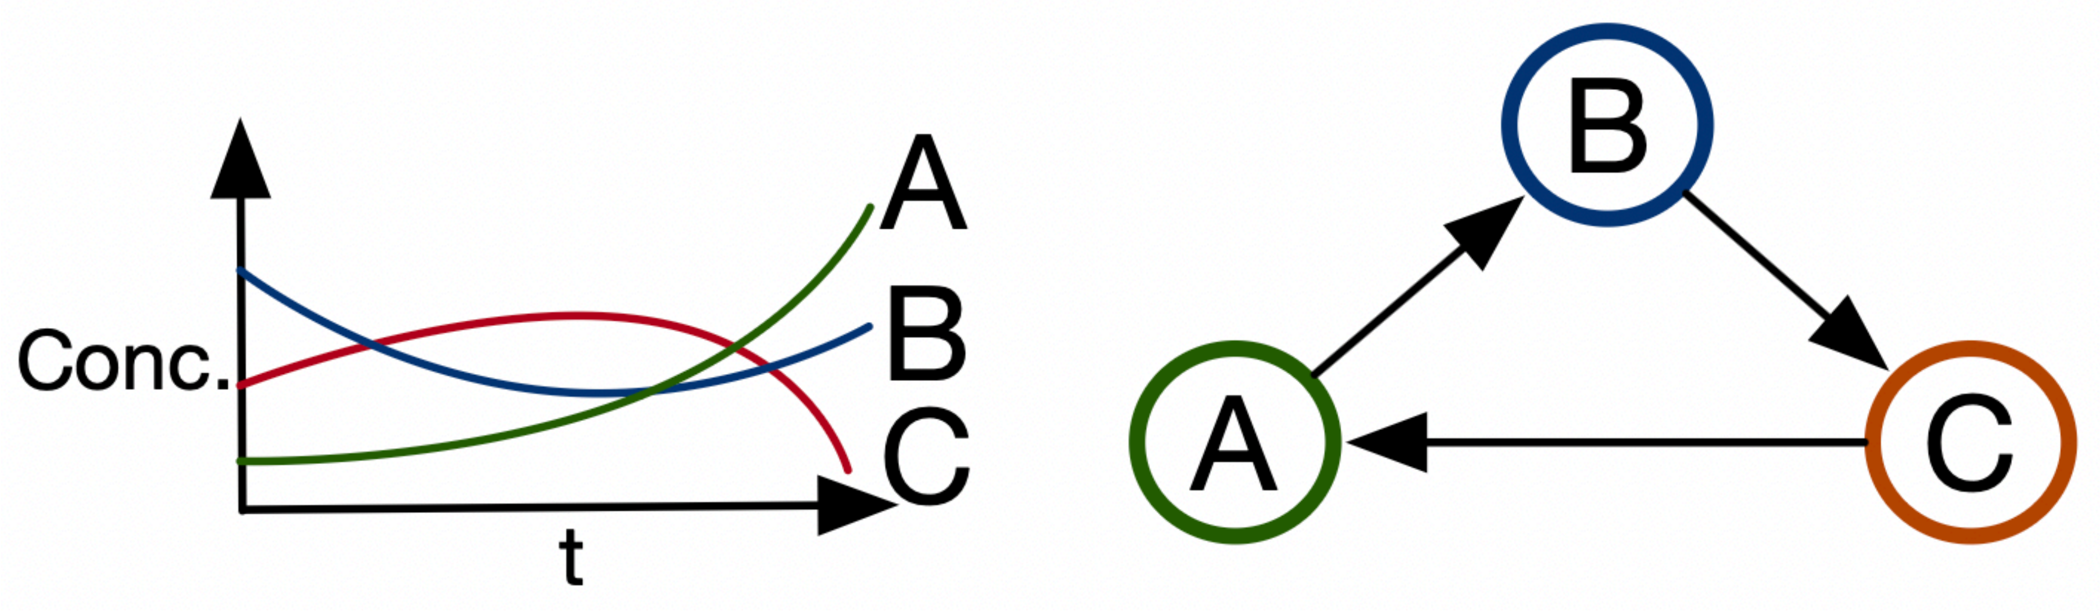

C-Sets: Graphs as representational modeling language

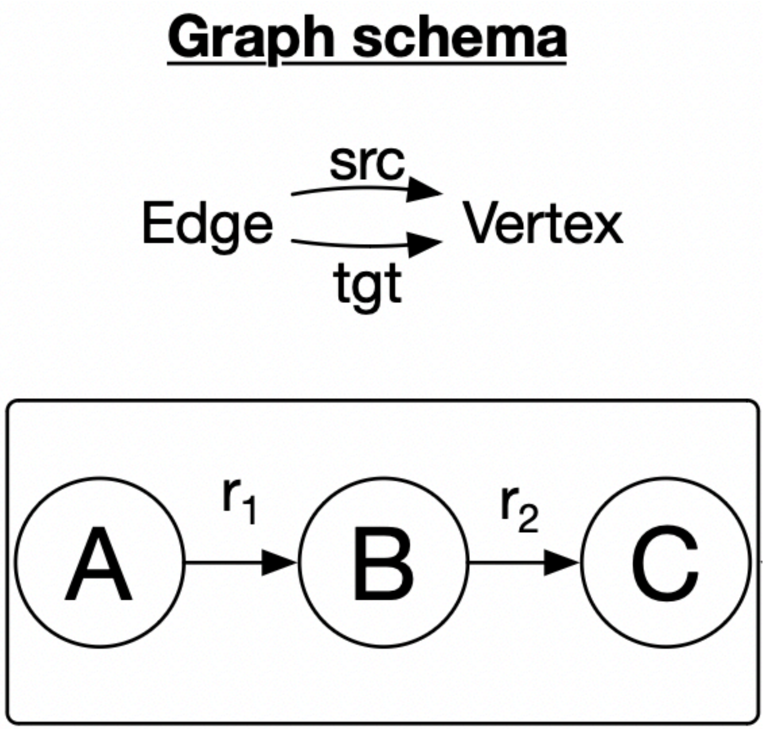

E	src	tgt
r_1	A	B
r_2	B	C

- We can think of directed multi-graphs as functors into Set from a particular ‘schema’ category.
- We can depict these C-Sets in a tabular format, where the outgoing arrows correspond to functions (which are represented as columns).
- The natural choice of morphism between functors, i.e. natural transformations, lines up perfectly with the ordinary notion of graph homomorphism.
- Being a functor into Set, this model can be thought of representationally: as describing the world as really having three species and two reactions which form a path.
- But different graphs purporting to represent the same phenomenon are making mutually incompatible claims.

C-Sets: Comparing models

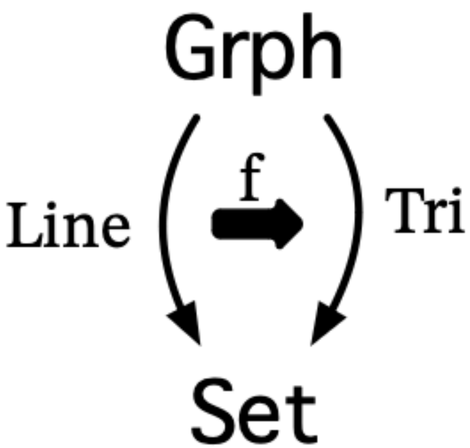

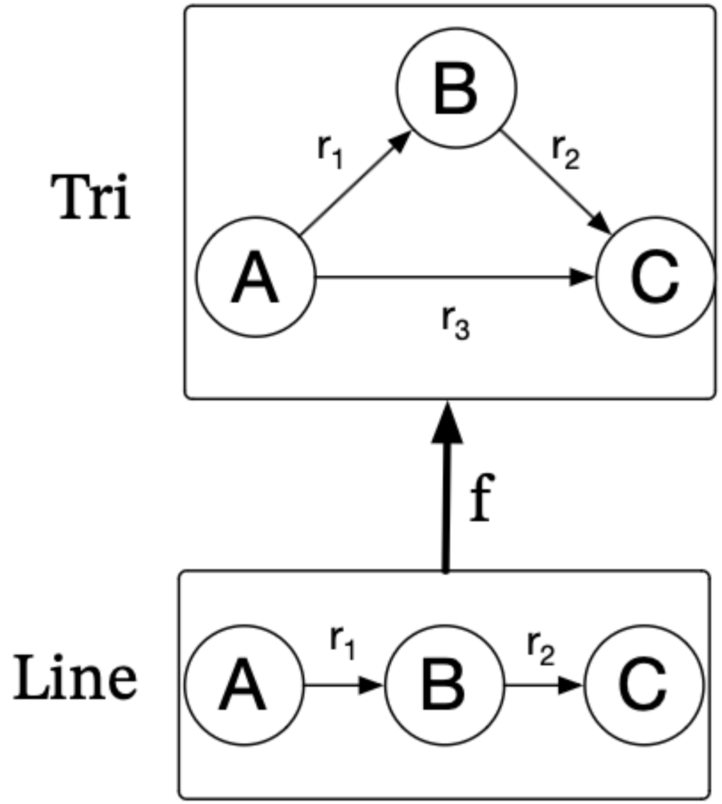

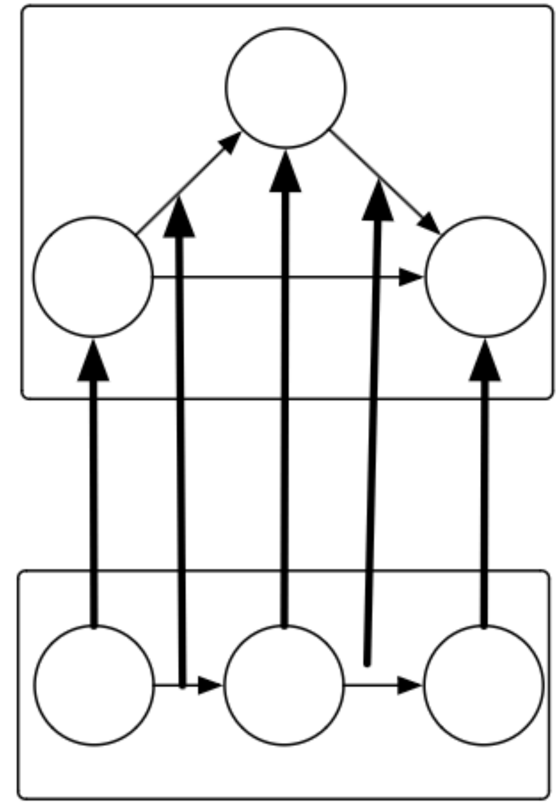

We can gain understanding of a model by recognizing it as, e.g., a subobject or (co)product of other models we feel we have a better grip on.

C-Sets: Comparing models at a *lower* level of abstraction

Level shift down

We view each model as a “schema” of sorts with models being assignments of parameters[1] to obtain dynamical systems.

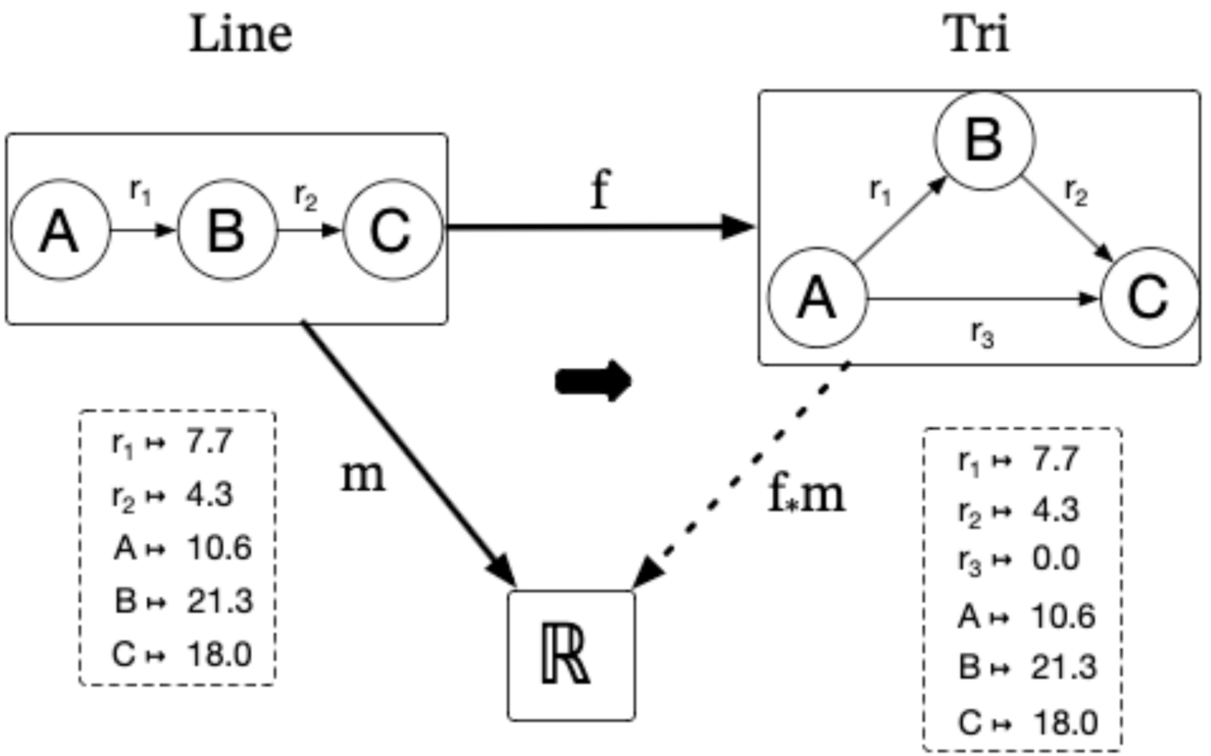

f is a recipe for taking Line models and producing Tri models.

$\mathrm{f}_* \mathrm{m}$ is characterized by a universal property. It is the most conservative way to do this interpretation in a precise sense.

C-Sets: Comparing models at a *higher* level of abstraction

Level shift up

We view each schema as a model of a theory of categories (Cat is the category of models).

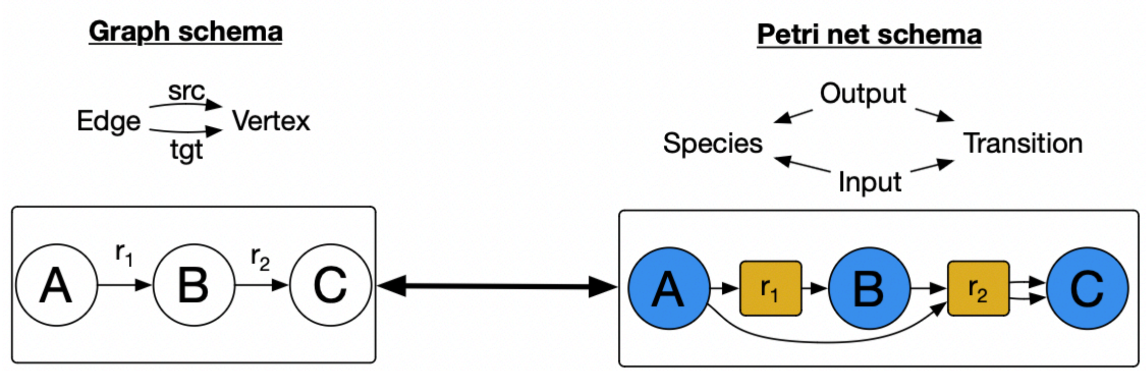

E	src	tgt
r_1	A	B
r_2	B	C

Input	S	T
i_1	A	r_1
i_2	B	r_2
i_3	A	r_2

Output	S	T
o_1	B	r_1
o_2	C	r_2
o_3	C	r_2

- Challenge assumptions: why can't reactions have more than one src or tgt?
- These different perspectives (different schemas) come with different trade-offs.

These people *ought* be able to productively communicate with each other, rather than:

- "Edges by definition aren't the sort of thing that have two sources. You're saying nonsense!"
- "Your model claims that 'inputs' and 'outputs' *really exist*. But all there really is in the chemical world are species and reactions. Show me where the 'inputs' are in the plot!"

C-Sets: Comparing models at a *higher* level of abstraction

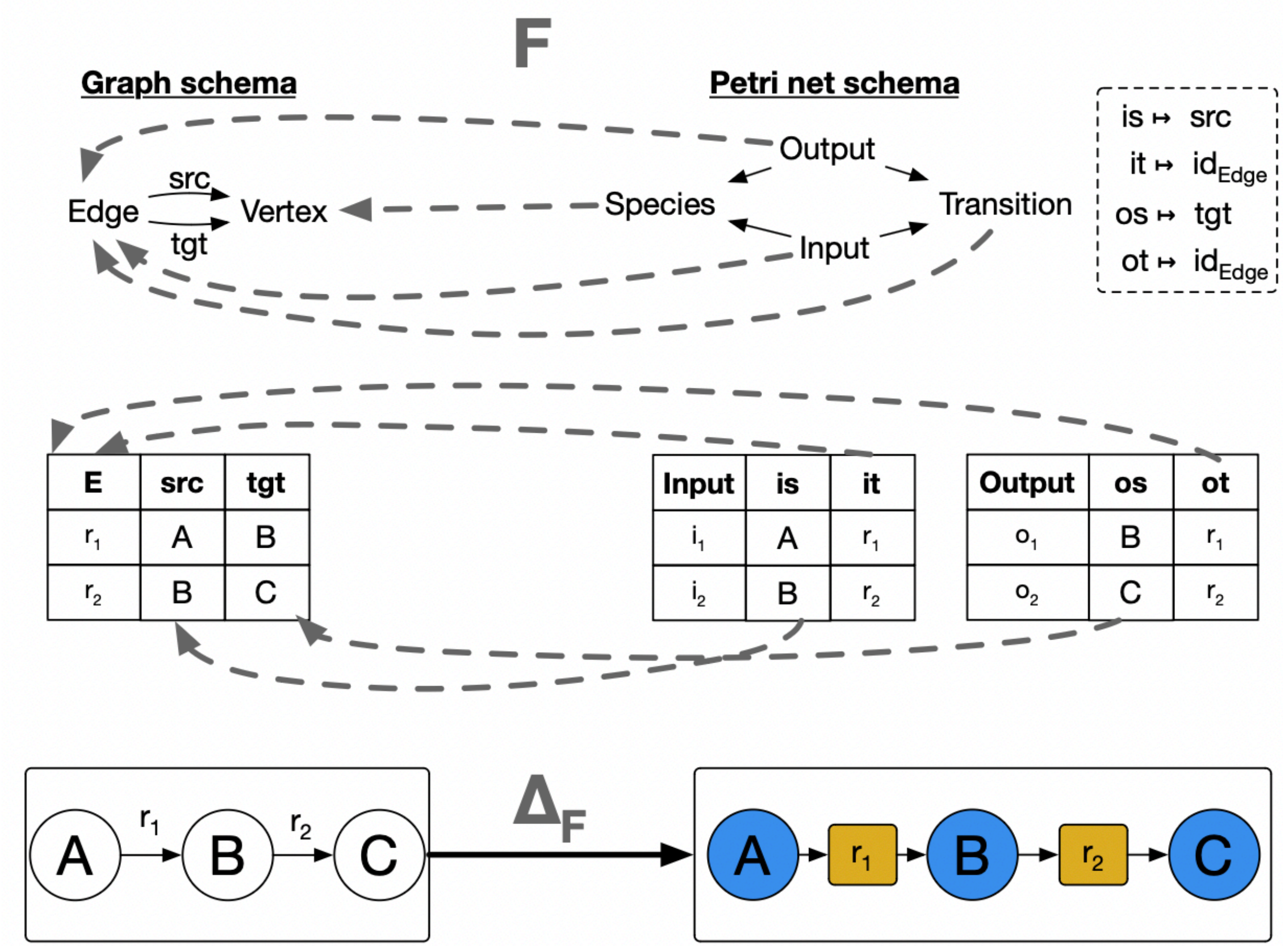

F is a recipe for taking Grph models and producing Petri models.

Δ_{F} is characterized by a universal property: we understand what the model is in relation to F and all other models of the relevant kind.

C-Sets: What have we accomplished?

1. Showed that, from Petri's perspective, that Grph isn't nonsense (and vice-versa)
    - Made explicit how it's possible for Petri and Grph to be about the same thing (via F)
2. Shown Petri under which conditions (in Petri's vocabulary) they aren't in conflict at all
    - Namely, all transitions have exactly one input and output.

A pure representationalist attitude encourages us to view conflict as irreconcilable:

- Tri is literally referring to things which don't exist
- Petri is literally saying that reactions have more than one source/target.

Here, our common sense lies squarely opposed to this:

Example 2: software design

Computational science involves lots of messy, unmaintainable, untrustable code.

Interfaces play the role of schemas in the previous example. They help give code structure.

The interface of a *monoid* is defined in GATlab below:

```
@theory ThMonoid begin
  # Carrier type constructor
  X::TYPE
  # Unit element term constructor
  e()::X
  # Multiplication term constructor
  (x ⋅ y)::X ⊣ [x::X, y::X]
  # Unitality axiom
  e ⋅ x == x == x ⋅ e ⊣ [x::X, y::X]
  # Associativity axiom
  x⋅(y⋅z) == (x⋅y)⋅z ⊣ [x::X, y::X, z::X]
end
```

A **generalized algebraic theory** consists in judgments (**type constructors**, **term constructors**, and **axioms**).

These judgments are built out of **types**, **terms**, and **contexts**.

Defining a monoid interface is helpful; we can write the following code *once*:

```
function aggregate_monoid_list(model::Monoid{M}, elems::Vector{M})::M where M
  @withmodel model (e, ⋅) begin
    result = e()
    for elem in elems
      result = result ⋅ elem
    end
    return result
  end
end
```

Verifying interface implementations

Implementations of an interface (i.e. models of a theory) send the type constructors to datatypes in one's language and term constructors to functions. Valid implementations respect the axioms.

```
@instance ThMonoid{Int} [model::NatPlus] begin
  e() = 0
  ·(x::Int,y::Int) = x + y
end
```

```
@instance ThMonoid{Int} [model::NatTimes] begin
  e() = 1
  ·(x::Int,y::Int) = x * y
end
```

How do we know if the implementation is valid given that the axioms are quantified over all possible inputs? There are two plausible approaches:

1. **Static verification**

- This will not help us verifying models in isolation because arbitrary, general-purpose code is allowed in the declaration of implementations.

2. **Property-based testing** (systematically check a finite number of inputs)

- This will not help us verifying models in isolation because there is no canonical way to enumerate values of arbitrary types in a general-purpose language.

Verifying implementations compositionally

What we care about is verifying a *particular* interface. But this is intractable in isolation.

A GAT morphism $\mathrm{F} : \mathrm{A} \to \mathrm{B}$ maps type constructors in A to types in B and term constructors in A to terms in B such that axioms are preserved.

This finite data *can* be verified (via a semidecidable e-graph algorithm).

Given such a morphism F, we have canonical ways of translating models between the theories, analogously to what we saw for C-Sets.[1]

```
# ThNatPlus defines ℕ, Z(ero), S(uccessor), and +
@theory ThVect <: ThNatPlus begin
  X::TYPE
  Vect(len::ℕ)::TYPE
  empty()::Vect(Z)
  concat(x::Vect(n), y::Vect(m))::Vect(n+m) ⊣ [(n,m)::ℕ]
  push(x::X, v::Vect(n))::Vect(S(n)) ⊣ [n::ℕ]

  concat(v, push(x,w)) == push(x, concat(v,w))
    ⊣ [(n,m)::ℕ,v::Vect(n),w::Vect(m), x::X]
end
```

```
@theorymap F(ThMonoid, ThVect) begin
  default() => Vect
  e() => empty()
  (a⋅b) ⊣ [(a,b)::default] => concat(a,b)
end

@assert is_natural(F) # run e-graph test

@withmodel Δ(F, MyVectImpl) (e, ⋅) begin
  @assert [1,2,3]⋅e()⋅[4,5] == [1,2,3,4,5]
end
```

We haven't verified `Δ(F, MyVectImpl)` is a valid `ThMonoid` model, but it is *if* `MyVectImpl` is a valid `ThVect` model. The problem is reduced if have a better understanding of the `ThVect` code or choose to trust that code.

Verifying implementations compositionally

If our *datatypes* are models of GATs, then we *do* have a systematic way of enumerating the possible values for property-based testing:

```
@theory ThReflGraph begin
  V::TYPE;
  E(src::V, tgt::V)::TYPE
  refl(v::V)::E(v,v) # a reflexive edge
end
```

```
# assuming coproduct(ReflGraph, ReflGraph) is defined
@instance ThMonoid{ReflGraph} [model::Coprod] begin
  e() = ReflGraph() # empty graph
  ⋅(x::ReflGraph,y::ReflGraph) = apex(coproduct(x, y))
end
```

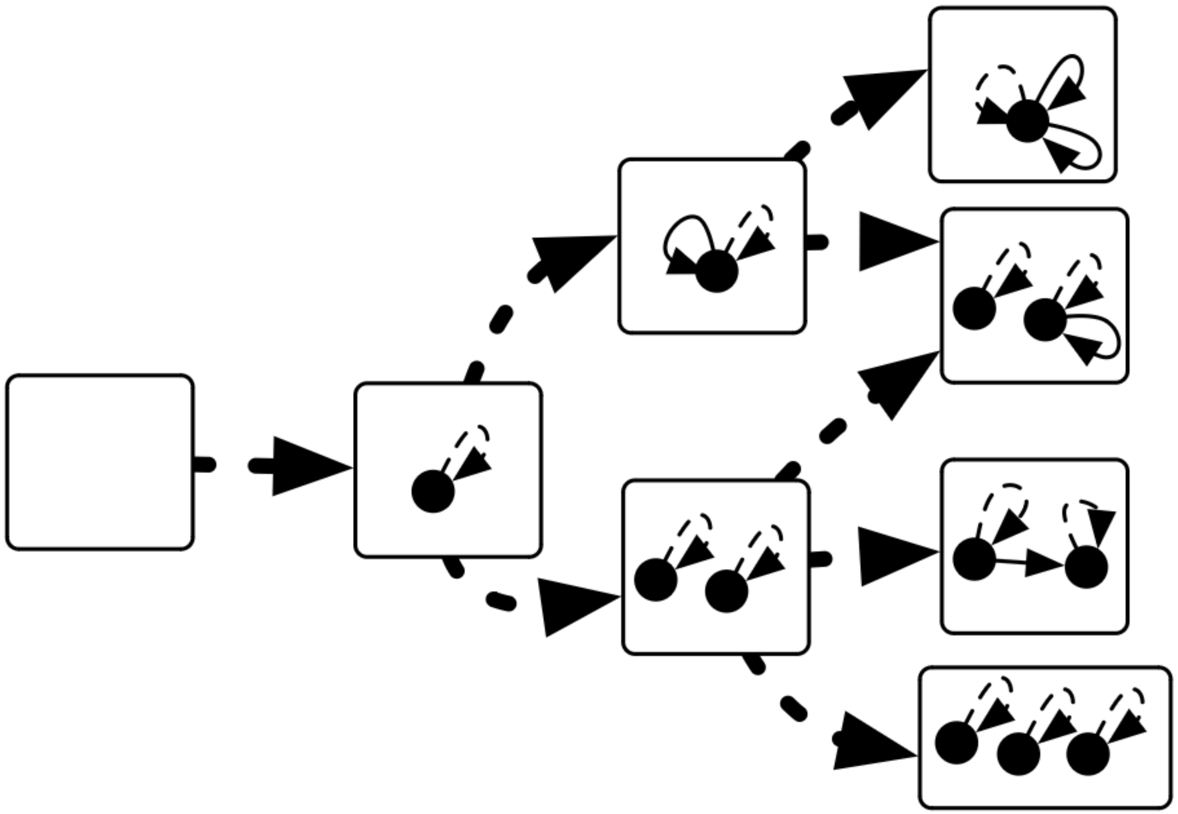

Conclusions

Conclusions: problems with representationalism?

Key components of the representationalist worldview are undercut by *philosophical* arguments from Quine, Sellars, Wittgenstein, Putnam, and many others.[1]

Formal theories have **logical** and **extralogical** symbols. E.g. $\mathrm{F} = \mathrm{m} \cdot \mathrm{a}$

The practical usage of formal theories involves taking the extralogical symbols to refer to aspects of our ordinary (informal) understanding of the world.

It can be very helpful to express a scientific model as a map $\mathrm{M} : \mathrm{T} \to \mathrm{W}$. However, it's easy to trick ourselves into thinking we've formalized the reference-binding process itself. However, the W side is just as much of a formal object as the T side.

It is only in virtue of collaborators performing this informal activity in a compatible way that models can appropriately function. But we are interested in cases where people **don't** do this in the same way, or disagree on what aspects of the world exist to be represented.

Conclusions: negotiating frame vs content

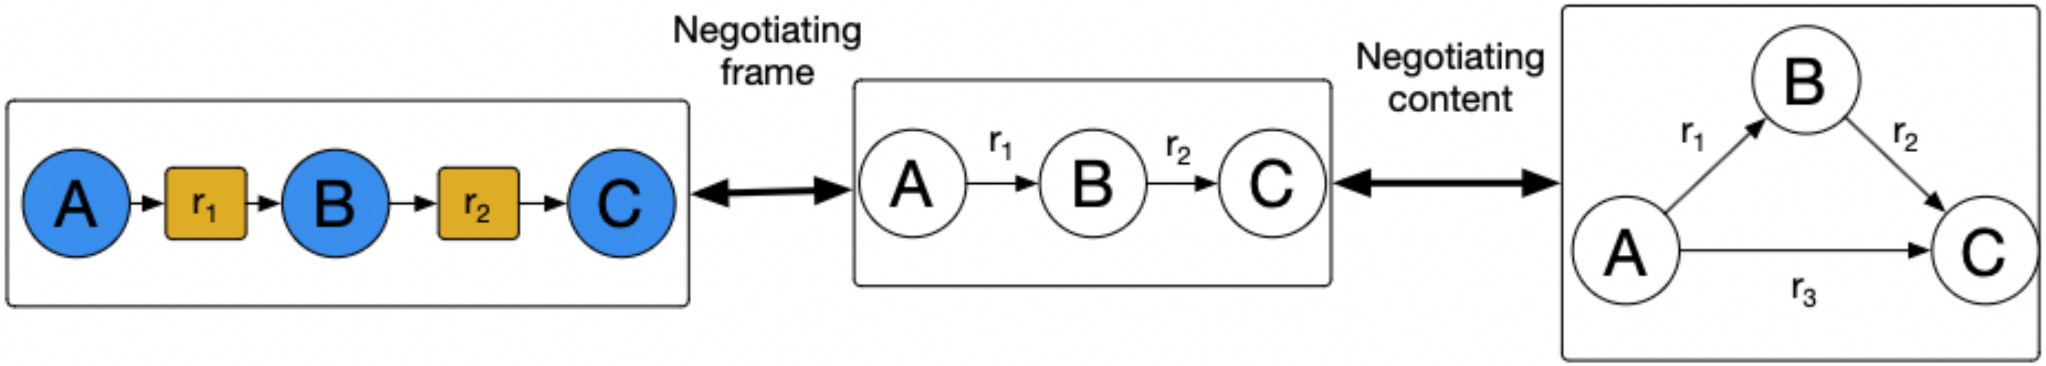

For a fixed frame, there might be a relatively straightforward process of negotiating content. But if we find this difficult, we can level shift upwards and seek a new frame in which our two frames are the content itself.

This is just codifying common sense of how reasonable people work out disagreements (offering arguments, finding common ground). We're at risk of losing our grip on this common sense when we start thinking that science and mathematics are an exception: that one can be speaking 'nonsense' by making apparent contradictions.

We can *take* someone to be saying nonsense (meaning: we give up on finding common ground) but *math* does not obligate us to do this; it's something we must occasionally do, but also it's something we should do deliberately and take responsibility for.

Consequences for scientific practice

- Operationalize some common sense:
  - apparent conflicts in our usage of terminology can be rationally resolved
- Pluralism of theories and methodologies and ontologies
  - Not a trivial kind of pluralism: finding common ground requires hard conceptual work
- An openness to confronting ethical / aesthetic aspects of scientific practice
  - Opening the possibility of these being meaningful (even if not located ‘in the world’)
  - Allowing there to be *reasons* for or against claims of this kind.
- Common templates for collective sensemaking
  - Categorification, (Co)limits, data migrations

Credits

Many have contributed to Catlab and GATlab, including:

Topos Institute

- Evan Patterson
- Owen Lynch
- Kevin Carlson

University of Florida

- James Fairbanks

Thank you!

Extra slides

Twenty minutes isn't a lot of time

What is good about representationalism?

- Language is **compositional**:[1] we can explain this (miraculous) fact by appealing to formal semantic models.

- Making a distinction between what we are saying and what we are talking about.
- A great deal of *inference* naturally accounted for by this bottom-up, atomistic story.
- In particular, *models* are essential to science. They allow one to make inferences, provided one assigns referents to the (formal) aspects of the theory.
- Atomistic semantics are practically convenient: the meaning of ‘green’ can’t possibly depend on the meaning of ‘red’ or whether or not it is a Tuesday. Meaningful sentences are made true independent of any context, and managing context is hard.

Conclusions: normativity of reference

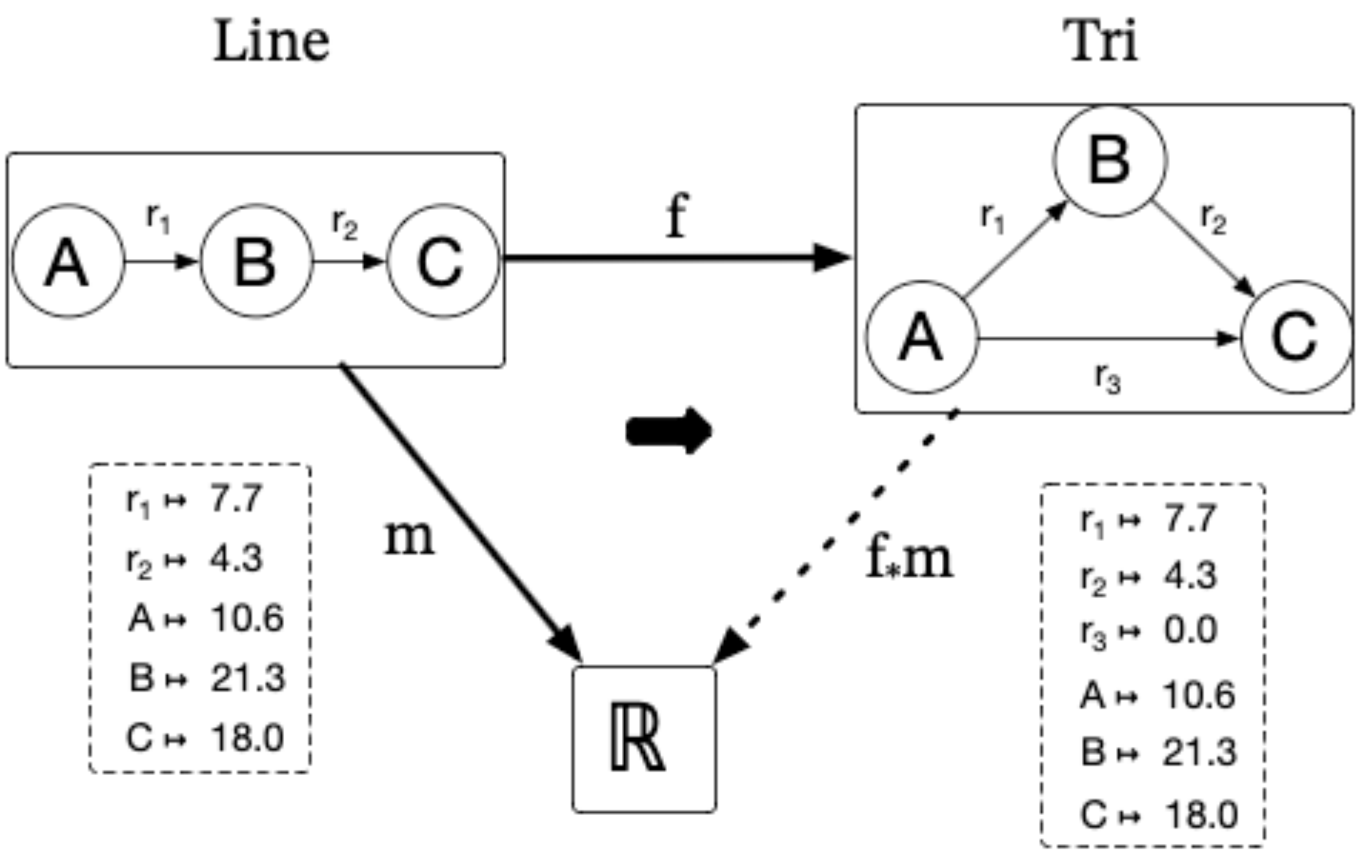

Here the categorical approach is making explicit the ι data: in order to relate these two models, *we* have to make a commitment (a choice). The co-reference does not come from the names anyone chose.

This is intimately connected to inferentialist solutions to paradoxes of representationalism and reference: reference is *normative*.

Inferentialism

Imagine you are a baby. You hear people saying opaque things.

You naturally pick up the practical skill of saying things (first: repeating the noises) yourself.

Via feedback, you start to gain partial mastery of **norms**: what actions are permissible / obligatory, including making vocalizations among these actions.

You can't help but pick up on certain correlations: you hear "You are a baby" and learn the move to saying yourself "I am a baby". "You are being loud" to "I am being loud". You start to act in accordance with the rule "You are φ" to "I am φ". One aspect of what "You are φ" **means** to you is that you are entitled to assert "I am φ".

Eventually you gain enough proficiency that you taken to have an understanding of what you say. You can be held responsible (it becomes *possible* for you to lie, to say *that* such-and-such is the case, to think).

Sometimes we can very concisely express what role something plays in inference:

$$\wedge\text{-}I\ \frac{A \qquad B}{A \wedge B} \qquad\qquad \wedge\text{-}E\ \frac{A \wedge B}{A} \quad \frac{A \wedge B}{B}$$

Representationalist platitudes: the fact-value dichotomy

One who says “Is that supposed to be a *fact* or a *value judgment*?” often presupposes:

- If it’s a value judgment, it can’t be a fact.
- Value judgments are subjective and cannot be verified.
- Moral/aesthetic arguments involve only persuasion, not reasoning or logical argument.
- What makes science rational is that its subject matter is matters of fact. Science consists of deductive/inductive methods which rationally settle disagreements
- The **analytic/synthetic dichotomy**:
  - if p is true, then p is true either by definition (convention) xor made true by the world.

Putnam: These beliefs are “culturally institutionalized”[1]

- They tend to persist even though the philosophers who accept them acknowledge that the *arguments* for them are terrible.

Arguments against representationalism: Quine

Two dogmas of empiricism: we can't specify the meanings of words first and then look to the world to see what facts are true. What something means is contingent on what facts we take to be true, so the dichotomy of statements between *analytic* (which are true purely in virtue of the meanings alone, such as "all bachelors are men" and "1+1=2") and *synthetic* ("the Earth is larger than the Moon") is bankrupt.

Arguments against representationalism: Sellars

Myth of the Given: experience is conceptually-articulated - the dichotomy of mental events between raw sense impressions and conceptual inferences is bankrupt. ‘Raw sense impressions’ makes sense in a scientific *causal* picture, applying equally to thermostats and animals. But our perceptions are the kind of thing that can serve as evidence in justifying claims, meaning they are *conceptually*-articulated.

Arguments against representationalism: Wittgenstein

Private language argument: intentionality is a social product, not a private relation between thinker and world. One cannot follow a rule (e.g. the meaning of a word) in isolation because any behavior can be understood as in accord with the rule (e.g. the multiplicity of rules one could have been following to generate any particular sequence of numbers). The social practices which litigate whether one is following a rule are themselves constitutive of the rule.

Arguments against representationalism: Putnam

Arguments against moral facts cut against ‘epistemic facts’ as well. Humean noncognitivism about ethics stems from that a fact is something we can perceive, and we don’t have sense impressions of ‘goodness’.

But neither do we have them for ‘simplicity’ or ‘coherence’, which are essential epistemic virtues for science.

Without the hope of agreement, argument would be pointless.

- It doesn’t follow that, without reaching an agreement, the argument was pointless.

Although there is something *characteristic* about objective (descriptive) talk and subjective (normative) talk, these are not *dichotomies*.

Representationalism is wrapped up with problematic dichotomies such as facts and values, meaning and use, perceptions and beliefs.

Implication frames

The philosophical problems with representationalism came from first trying to specify *meaning* and then specify *use*.

In an implication frame,

The models hitherto discussed were calculational devices for managing commitments and entitlements. What would a language look like that represents these directly?

There is much to say about the mathematics underlying logical expressivism, and there is a lot of interesting future work to do. A future blog post will methodically go over this, but this section will just give a preview.

Implication frames

An **implication frame** (or: **vocabulary**) is the data of a $\mathrel{|\!\sim}$ relation, i.e. a **lexicon** L (a set of claimables: things that can be said) and a set of **incoherences**, $\mathbb{I} \subseteq \mathrm{P}(\mathrm{L} + \mathrm{L})$, i.e. the good implications where it is incoherent to deny the conclusions while accepting the premises.

Given any base vocabulary $\mathrm{X} = (\mathrm{L}_{\mathrm{X}}, \mathbb{I}_{\mathrm{X}})$, we can introduce a **logically-elaborated vocabulary** whose lexicon includes L_{X} but is also closed under $\neg$, $\rightarrow$, $\land$, $\lor$. The $\mathbb{I}$ of the logically-elaborated relation is precisely $\mathbb{I}_{\mathrm{X}}$ when restricted to the nonlogical vocabulary (i.e. logical vocabulary must be *harmonious*), and the following sequent rules indicate how the goodness of implications including logical vocabulary is determined by implications without logical vocabulary.

$$\frac{\Gamma \mathrel{|\!\sim} \Delta, A}{\Gamma, \neg A \mathrel{|\!\sim} \Delta} \qquad \frac{\Gamma, A, B \mathrel{|\!\sim} \Delta}{\Gamma, A \land B \mathrel{|\!\sim} \Delta} \qquad \frac{\Gamma, A \mathrel{|\!\sim} \Delta \quad \Gamma, B \mathrel{|\!\sim} \Delta \quad \Gamma, A, B \mathrel{|\!\sim} \Delta}{\Gamma, A \lor B \mathrel{|\!\sim} \Delta} \qquad \frac{\Gamma, B \mathrel{|\!\sim} \Delta \quad \Gamma \mathrel{|\!\sim} A, \Delta \quad \Gamma, B \mathrel{|\!\sim} A, \Delta}{\Gamma, A \rightarrow B \mathrel{|\!\sim} \Delta}$$

$$\frac{\Gamma, A \mathrel{|\!\sim} \Delta}{\Gamma \mathrel{|\!\sim} \Delta, \neg A} \qquad \frac{\Gamma \mathrel{|\!\sim} \Delta, A \quad \Gamma \mathrel{|\!\sim} \Delta, B \quad \Gamma \mathrel{|\!\sim} \Delta, A, B}{\Gamma \mathrel{|\!\sim} \Delta, A \land B} \qquad \frac{\Gamma \mathrel{|\!\sim} \Delta, A, B}{\Gamma \mathrel{|\!\sim} \Delta, A \lor B} \qquad \frac{\Gamma, A \mathrel{|\!\sim} \Delta, B}{\Gamma \mathrel{|\!\sim} \Delta, A \rightarrow B}$$

Implication frames

The double bars are *bidirectional* meta-inferences: thus they provide both introduction and elimination rules for each connective being used as a premise or a conclusion. They are quantified over all possible sets Γ and Δ. The top of each rule makes no reference to logical vocabulary, so the logical expressions can be seen as making explicit the implications of the non-logical vocabulary.

Vocabularies can be given a semantics based on **implicational roles**, where the role of an inference $\mathrm{a} \mid\!\sim \mathrm{b}$ is the set of implications $\Gamma \mid\!\sim \Delta$ in which $\mathrm{a} \mid\!\sim \mathrm{b}$ is a good inference:

$$(\mathrm{a} \mid\!\sim \mathrm{b})^* := \{(\Gamma, \Delta) \mid \Gamma, \mathrm{a} \mid\!\sim \mathrm{b}, \Delta \in \mathbb{I}\}$$

The role of an implication can also be called its **range of subjunctive robustness**.

Implication frames

To see an example, first let's remind ourselves of our q ("The cat has four legs") and r ("The cat lost a leg") example, a vocabulary which we'll call $C = (L_C = \{q, r\}, \mathbb{I}_C)$:

$\mathbb{I}_C$	**0**	q^-	r^-	$q^- r^-$
0	$\checkmark$	$\boxed{\checkmark}$	$\boxtimes$	$\boxed{\checkmark}$
q^+	$\boxtimes$	$\checkmark$	$\boxtimes$	$\checkmark$
r^+	$\boxtimes$	$\boxtimes$	$\checkmark$	$\checkmark$
$q^+ r^+$	$\boxed{\checkmark}$	$\checkmark$	$\checkmark$	$\checkmark$

The role of q^- (i.e. $\mid\!\sim q$) in vocabulary C is the set of all 16 possible implications *except* for $r \mid\!\sim$ and $r \mid\!\sim q$.

The role of a set of implications is defined as the intersection of the roles of each element:

$$\Gamma^* := \bigcap_{\gamma \in \Gamma} \gamma^*$$

Implication frames

The power set of implications for a given lexicon have a quantale structure with the $\otimes$ operation:

$$\Gamma \otimes \Delta := \{\gamma \cup \delta \mid (\gamma, \delta) \in \Gamma \times \Delta\}$$

Roles are naturally combined via a dual operation, $\sqcup$:

$$r_1 \sqcup r_2 := (r_1^* \otimes r_2^*)^*$$

A pair of roles (a premisory role and a conclusory role) is called a **conceptual content**: to see why *pairs* of roles are important, consider how the sequent rules for logical connectives are quite different for introducing a logically complex term on the left vs the right of the turnstile; in general, the inferential role of a sentence is different depending on whether it is being used as a premise or a conclusion. Any sentence $a \in L$ has its premisory and conclusory roles as a canonical conceptual content, written in typewriter font:

$$\mathtt{a} := ((a \mathrel{\sim\!\!\!\mid})^*, (\mathrel{\mid\!\!\sim} a)^*)$$

Implication frames

Below are recursive semantic formulas for logical connectives: given arbitrary conceptual contents $\mathrm{A} = (\mathrm{a}^+, \mathrm{a}^-)$ and $\mathrm{B} = (\mathrm{b}^+, \mathrm{b}^-)$, we define the premisory and conclusory roles of logical combinations of A and B. Because $\sqcup$ is an operation that depends on all of ▯, this is both a compositional *and* a holistic semantics.

Connective	Premisory role	Conclusory role
$\neg\mathrm{A}$	a^-	a^+
$\mathrm{A} \wedge \mathrm{B}$	$\mathrm{a}^+ \sqcup \mathrm{b}^+$	$\mathrm{a}^- \cap \mathrm{b}^- \cap (\mathrm{a}^- \sqcup \mathrm{b}^-)$
$\mathrm{A} \vee \mathrm{B}$	$\mathrm{a}^+ \cap \mathrm{b}^+ \cap (\mathrm{a}^+ \sqcup \mathrm{b}^+)$	$\mathrm{a}^- \sqcup \mathrm{b}^-$
$\mathrm{A} \rightarrow \mathrm{B}$	$\mathrm{a}^- \cap \mathrm{b}^+ \cap (\mathrm{a}^- \sqcup \mathrm{b}^+)$	$\mathrm{a}^+ \sqcup \mathrm{b}^-$

Note: each cell in this table corresponds directly to a sequent rule further above, where combination of sentences *within* a sequent corresponds to $\sqcup$, and multiple sequents are combined via $\cap$.

Implication frames

There are other operators we can define on conceptual contents, such as $\mathtt{A}^{+} := (a^{+}, a^{+})$ and $\mathtt{A} \sqcup \mathtt{B} := (a^{+} \sqcup b^{+}, a^{-} \sqcup b^{-})$.

Given two sets $\mathtt{G} = \{g_1, \ldots, g_m\}, \mathtt{D} = \{d_1 \ldots, d_n\}$ of conceptual contents, we can define **content entailment**:

$$\mathtt{G} \Vdash\!\sim \mathtt{D} := \mathbb{I}^* \subseteq g_1^+ \sqcup \ldots \sqcup g_m^+ \sqcup d_1^- \sqcup \ldots \sqcup d_n^-$$

Implication frames

A preliminary computational implementation (in Julia, available on Github) supports declaring implication frames, computing conceptual roles / contents, and computing the logical combinations and entailments of these contents. This can be used to demonstrate that this is a supraclassical logic:[1] this semantics validates all the tautologies of classical logic while also giving one the ability to reason about the entailment of nonlogical contents (or combinations of both logical and nonlogical contents).

```
C = ImpFrame([[]=>[:q], []=>[:q,:r], [:q,:r]=>[]], [:q,:r]; containment=true)
q, r = contents(C)            # canonical contents for the bearers q and r
∅ = nothing                   # empty set of contents
@test ∅ ⊩ (((q → r) → q) → q) # Pierce's law
@test ∅ ⊮ ((q → r) → q)       # not Pierce's law
```

Interpretations

We can *interpret* a lexicon in another vocabulary. An **interpretation function** $[\![-]\!] : \mathrm{A} \to \mathrm{B}$ between vocabularies assigns conceptual contents in B to sentences of A. We often want the interpretation function to be compatible with the structure of the domain and codomain: it is **sound** if for any candidate implication in A, we have $\Gamma \mid\!\sim_{\mathrm{C}} \Delta$ iff $[\![\Gamma]\!] \mid\!\mid\!\sim_{\mathrm{B}} [\![\Delta]\!]$.

Interpretations

To see an example of interpretations, let's first define a new vocabulary S with $L_S = \{x, y, z\}$.

- x: "It started in state s"
- y: "It's presently in state s"
- z: "There has been a net change in state".

$[\![\,]\!]_S$	**0**	x^-	y^-	z^-	$x^- y^-$	$x^- z^-$	$y^- z^-$	$x^- y^- z^-$
0	✓	☒	☒	☒	☒	☒	☒	☒
x^+	☒	✓	☑	☒	✓	✓	☑	✓
y^+	☒	☒	✓	☒	✓	☒	✓	✓
z^+	☒	☒	☒	✓	☒	✓	✓	✓
$x^+ y^+$	☒	✓	✓	☒	✓	✓	✓	✓
$x^+ z^+$	☒	✓	☒	✓	✓	✓	✓	✓
$y^+ z^+$	☒	☒	✓	✓	✓	✓	✓	✓
$x^+ y^+ z^+$	☑	✓	✓	✓	✓	✓	✓	✓

Interpretations

$\mathbb{S}$ claims it is part of our concept of 'state' that something stays in a given state, unless its state has changed (hence there is a similar non-monotonicity to the one in $\mathbb{C}$, but now with $x \mathrel{|\!\sim} y$ and $x, z \not\mathrel{|\!\sim} y$). We can understand what someone is saying by r or q in terms of interpreting these claimables in $\mathbb{S}$. The interpretation function $q \mapsto x^+ \sqcup y$ and $r \mapsto x^+ \sqcup z$ is sound. We can offer a full account for what we meant by our talk about cats and legs in terms of the concepts of states and change.

Interpretations

We could also start with a lexicon $L_D = \{x, y, z\}$ and interpretation function $q \mapsto x \to y$ and $r \mapsto x \to z$; we can compute what structure $\mathbb{I}_D$ must have in order for us to see $\mathbb{I}_C$ as generated by the interpretation of q, r in D. Below [?] means that it doesn't matter whether that implication is in $\mathbb{I}_D$:

$\mathbb{I}_D$	0	x^-	y^-	z^-	$x^- y^-$	$x^- z^-$	$y^- z^-$	$x^- y^- z^-$
0	✓	[✓]	[?]	[?]	[?]	[?]	[?]	[?]
x^+	[×]	✓	[✓]	[×]	✓	✓	[✓]	✓
y^+	[×]	[✓]	✓	[?]	✓	[?]	✓	✓
z^+	[?]	[✓]	[?]	✓	[?]	✓	✓	✓
$x^+ y^+$	[?]	✓	✓	[×]	✓	✓	✓	✓
$x^+ z^+$	[?]	✓	[×]	✓	✓	✓	✓	✓
$y^+ z^+$	[✓]	[✓]	✓	✓	✓	✓	✓	✓
$x^+ y^+ z^+$	[?]	✓	✓	✓	✓	✓	✓	✓

Interpretations

We can do the same with $q \mapsto x \wedge y$ and $r \mapsto x \wedge z$.

$\llbracket\,\rrbracket_D$	**0**	x^-	y^-	z^-	$x^- y^-$	$x^- z^-$	$y^- z^-$	$x^- y^- z^-$
0	$\checkmark$	$\boxed{\checkmark}$	$\boxed{\checkmark}$	$\boxed{\times}$	$\boxed{\checkmark}$	$\boxed{\checkmark}$	$\boxed{\checkmark}$	$\boxed{\checkmark}$
x^+	$\boxed{?}$	$\checkmark$	$\boxed{?}$	$\boxed{?}$	$\checkmark$	$\checkmark$	$\boxed{?}$	$\checkmark$
y^+	$\boxed{?}$	$\boxed{?}$	$\checkmark$	$\boxed{?}$	$\checkmark$	$\boxed{?}$	$\checkmark$	$\checkmark$
z^+	$\boxed{?}$	$\boxed{?}$	$\boxed{?}$	$\checkmark$	$\boxed{?}$	$\checkmark$	$\checkmark$	$\checkmark$
$x^+ y^+$	$\boxed{\times}$	$\checkmark$	$\checkmark$	$\boxed{\times}$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$
$x^+ z^+$	$\boxed{\times}$	$\checkmark$	$\boxed{\times}$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$
$y^+ z^+$	$\boxed{\times}$	$\boxed{\times}$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$
$x^+ y^+ z^+$	$\boxed{\checkmark}$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$

Interpretations

Interpretation functions can be used to generate vocabularies in via a `sound_dom` function in our software implementation, which constructs the domain ▯ of an interpretation function from the assumption that it is sound. The following code witnesses how we recover our earlier vocabulary C via interpretation functions into the vocabularies S and D above.

```
S = ImpFrame([[:x]=>[:y], [:x]=>[:y,:z], [:x,:y,:z]=>[]]; containment=true)
x, y, z = contents(S)
x⁺ = Content(prem(x), prem(x))
f = Interp(q = x⁺ ⊔ y, r = x⁺ ⊔ z)
@test sound_dom(f) == C

D = ImpFrame([[]=>[:x], []=>[:y], []=>[:x,:y], []=>[:x,:z],
              []=>[:y,:z],[]=>[:x,:y,:z],[:x,:y,:z]=>[]]; containment=true)
x, y, z = contents(D)
f = Interp(q = x ∧ y, r = x ∧ z)
@test sound_dom(f) == C
```