Math Does Not Represent

Erik Curiel

Munich Center For Mathematical Philosophy
LMU Munich

Black Hole Initiative
Harvard University

Quantum Information Structure of Spacetime Consortium

erik@strangebeautiful.com

Gefördert durch die Deutsche Forschungsgemeinschaft (DFG) – Projektnummer 312032894, CU 338/1-2. (Funded by the German Research Foundation (DFG) – project number 312032894, CU 338/1-2.)

The vague accuracy of events
dancing two by two with language,
which they forever surpass

William Carlos Williams
Paterson

Outline

Representation

Nine Problems

Theories, Knowledge, Concepts

Realism – An Envoi

Representation

Nine Problems

Theories, Knowledge, Concepts

Realism – An Envoi

By 'represent' here I mean in anything like a standard philosophical sense of designative, depictive or verisimilous representation of the world by the mathematics of our best physical theories, *e.g.*:

- a Tarskian relation of designation between elements of a mathematical space ("the space of states") and states of a physical system as in some versions of the semantic view (designative)
- or similarity along the lines of Giere (depictive)
- or the existence of a homomorphism of some kind between mathematical stuff and stuff in the world *à la* the structuralists (verisimilous)
- or a possible-worlds semantics for solutions to equations of motion or field equations (elements of all three)

meaning is fixed by ontology

Even empiricists and instrumentalist such as Carnap and van Fraassen subscribe to Tarskian-like semantics to give empirical content to the mathematical formalism of theories.

It is in that sense that I claim that they (and almost all other empiricists and instrumentalists I know of today) take the meaning of mathematical formalism to be determined by designative relations with respect to a fixed ontology, even if they are not realists about the ontology.

natural accompaniment:

> *empirical content accrues to the mathematical formalism of theory largely if not wholly by virtue of this kind of representation—any physical significance the mathematics has derives from it*

In today's world of philosophy of physics, standard practice is to require one of the following in order for one to feel that one has at least the grounds for understanding a theory:

1. a fixing of the fundamental ontology of the theory;
2. or a fixing of the semantics of the theory using some variant of the semantic view *à la* van Fraassen or (more popular these days) a possible-worlds semantics *à la* Lewis
3. or whatever is demanded by any of a number of more niche views about how to fix the empirical content of a theory, which nonetheless have ardent backers, like the use of category theory to characterize the models of a theory in such a way as (the sometimes implicit claim goes) to allow for a unique articulation of their physical significance.

Most often, proponents of such a thing think that that is *all* one needs in order to do all the work of understanding a theory.

We need the freedom to pick and choose what philosophical tools may be fruitful in any given investigative context, without the prior expectation that one approach or set of tools will work everywhere.

- sometimes, sketching a bare set of abstract interpretive principles does all the work one wants
- other times, one needs to get down and dirty with the fine details of theoretical calculations and of possible experimental design without any pretense that a possible-world semantics or abstract interpretive principles will do

1. complexity of the world

Steven French once asked me over beers (well, I was having beer and Steven was having a non-alcoholic beverage), what is the difference between mathematical space and physical space?

I replied, well, it's easier to find decent Chinese take-out in physical space.

It was a joke, but it also had a point: I know how to recognize decent Chinese take-out in physical space. I have no idea how even to *represent* "decent Chinese take-out", much less recognize it, in *any* mathematical structure we use to model physical systems in our best theories.[1]

1. "You think I open a restaurant in the middle of the 'hood and don't know what's going on? I fucking represent." – The Fugees, "Chinese Restaurant"

The world is a complex place, and our mathematical models of its parts are almost childishly, recklessly simple. How can a relation of "representation" hold between them?

(Steven and his ilk would be unimpressed by considerations of complexity, at least as presented so baldly and briefly as here. I return to ideas related to this theme and elaborate on them in problems 7–9 below)

2. levels of abstraction

Does—*can*—

$$\mathbf{F} = m\mathbf{a}$$

"represent" in the same way as

$$\mathbf{F} = G\frac{mM\hat{\mathbf{r}}}{r^2}$$

Indeed, can it "represent" at all?

- Can the latter do so in the same way as the equation modeling 2 perfect homogeneous spheres as a Keplerian binary system without a specified target?
- That modeling the Earth and the Sun as a concrete, individual gravitationally coupled system, with lunar and Jovian perturbations accounted for?
- Do different levels of abstraction "represent" in the same way?
- How does one decide when the mathematics is "concrete enough" to represent?

3. “reading off” the meaning of mathematics

Not every mathematical operation one can meaningfully perform in the structures used to formulate a theory has physical significance.

Just because one *can* apply a mathematical result or operation does not *ipso facto* entail that it will have physical significance.

And even if it does seem to have physical significance, one must always take some care and think about what that physical significance is, not assume one knows it by “reading it off” the mathematics in a way naively suggested by depictive representation.

the Hole Argument is a perfect example:

1. assume spacetime points are substantival in the sense that it makes sense to imagine the world as it is, and the world such that each body occupies different spacetime points
2. the relationship between those two physical situations, the Hole Argument assumes, is represented by a diffeomorphism, because that is the mathematical operation representing "pushing stuff (tensor fields) around on top of fixed spacetime (manifold) points"
3. but I have never seen an argument that dragging tensors around without also dragging the points around, as a mathematical operation, has physical meaning at all, much less *that* meaning, *in the investigative context in which the Hole Argument is posed*
4. that diffeomorphisms are properly interpreted in that way in such a context is exactly what I want an argument for: the very fact that trying to interpret diffeomorphisms in that way leads to a quagmire such as the Hole Argument is a *reductio* of such interpretation
5. the reply will inevitably come, "But that's just what diffeomorphisms mean, it's obvious from the mathematics." No. Mathematics does not interpret itself, no matter how "obvious" it may seem. There are always different possible ways of interpreting the formalism, and one must have principled grounds for choosing one.

4. same operation, different meanings

Sometimes, one and the same mathematical operation has different physical significance depending on the investigative context in which it is performed.

It makes no physical sense to “add two spatial positions” as a possible representation of a physical operation, even though they are encoded in the structure of a 3-dimensional vector space.

It does make sense to add them when computing a factitious quantity such as the center of mass of a system of particles.

The formalism alone cannot distinguish such cases, as would be the case if the math represented the world in a depictive sense.

Only pragmatics can distinguish them, a judgment of how the math is to be applied and interpreted in any given investigative context not itself determined by any formal or structural features of the math.

5. different formulations of the same theory

A naive reading of the mathematics misleads. General relativity:

1. manifold and metric tensor
2. manifold, metric tensor, and stress-energy tensor
3. manifold and tetrad field
4. manifold, tetrad field, and stress-energy tensor
5. section of $SO(3, 1)$ principal fiber bundle
6. section of $SO(3, 1)$ principal fiber bundle over a manifold and section of the tensor bundle of 2-index symmetric covariant tensors over the same manifold
7. manifold, choice of spin structure, and cross-section of the spinor bundle over the double covering space of the manifold
8. manifold, choice of spin structure, cross-section of the spinor bundle over the double covering space of the manifold, and spin coefficients of the stress-energy tensor
9. manifold, a diffeomorphism-invariant gauge theory of the Lorentz group, with Lagrangian of the type $f(F \wedge F)$, where F is the curvature 2-form of the spin connection, with 6 extra primary and secondary constraints

cont.

10. manifold and Synge's biscalar "world function" (geodesic distance)
11. manifold, Synge's biscalar "world function" (geodesic distance), and stress-energy tensor
12. manifold and fixed set of values for scalars in parametrized post-Newtonian formalism
13. manifold and fixed set of values for scalars in parametrized post-Newtonian formalism, and stress-energy tensor
14. Einstein algebra
15. Einstein-Hilbert action
16. Palatini action (and variants)
17. Eddington-Schrödinger affine formulation
18. Cartan's tetrad formulation (and variants)
19. Plebański's chiral formulation
20. chiral pure-connection formulation
21. all the many, many 3+1 formulations
22. all the many, many 2+2 formulations
23. ...

Do some “really represent” and others not?

Do some “really represent” only in so far as one can transform their results into the language of The One Formulation, canonical and privileged above all others?

That seems silly to me. In so far as they all work equally well, they all represent equally well. Or not.

There is no principled way to say one is “canonical” or “privileged”.

6. meaning without representation

There can be meaning, physical significance, without representation: the physical significance of the symplectic structure in Hamiltonian mechanics is, in part, energy conservation.

Not even the hardest of the hardcore realists would be tempted to think that the symplectic structure "represents".

Indeed, one can not even make sense of the idea that it designates something in, *e.g.*, the semantic view of theories or possible world semantics, because it is a relation on the entire class of models.

7. "Oh, Lord, ooh, you are so big, so absolutely huge...

...Gosh, we're all really impressed down here I can tell you."

The overwhelmingly, unimaginably vast majority of structures, solutions, *etc.*, associated with the mathematical formalism of every single physical theory I know is unknown to us, and, moreover, *unknowable* by us, and by any epistemic agents we can imagine as being remotely like us in any way.

GR as class of Lorentzian 4-geometries:

- even just at the level of topology, the problem is inconceivable
- the issue of individual metrics exponentiates it
- even trying to use "general principles" to narrow down and fix on "manageable, physically reasonable" classes of spacetimes doesn't help

I think we tend to forget in philosophy, we lose sight of, how complex real physical systems are and what a miracle it is that our almost naively, recklessly simple-minded theories, and the childishly sketched models we construct in those theories, can still capture them with astonishing accuracy, and do so in ways, moreover, that seem to give us real understanding of the nature of the world in a broader sense.

And correlatively, we forget how much distance there is between those simple models we do know in any given theory and the real physical systems they purport to represent, and so we forget how many of the theory's other models would also perform the depictive representational tasks better and *prima facie* do them better (by adding ever more finely grained detail, for example), if one of them does at all—how many of those models that we have no idea how to construct in any way graspable by the human mind, or indeed even how to identify if someone gave us one gift-wrapped.

Thus we are saying that, on depictive views, the overwhelming majority of the empirical content of a theory is something that we do not know, are in fact nowhere near knowing, and have good reason to think we will never know in anything resembling thoroughness and detail

epistemologically speaking, the idea that a theory even has non-trivial cognizable empirical content is an act of faith.

What possible use could such a thing have to a fruitful analysis of the content and structure of our knowledge and understanding in physics?

8. which parts of the math represent?

Consider the problem of knowing *which parts* of "realistic" or "physically significant" solutions and models do represent parts of physically possible systems. This is not the problem of "gauge" or "surplus structure".

Take Navier-Stokes theory (the classical theory of viscöelastic, thermoconductive fluids), a continuum theory—the solutions to its equations allow one to make, indeed they *necessarily encode*, physical predictions at arbitrarily small spatial and temporal scales (say, 10^{-100} cm); but we know that it is not appropriate for that task, and we know we can safely ignore those "predictions" of the theory in assessing its propriety and adequacy.

We know this, however, by nothing determined by the formalism, by the mathematics. It is wholly a pragmatic affair to determine the regime of applicability of a physical theory. ("It needs a lot of intuitive physical sense to know when to expect actual things to behave like the idealized models we make of them." –Synge, *Relativity: The Special Theory*, ch. i, §18, p. 32.)

9. the breakdown of theories

We use theories to model systems, and to model them well, *i.e.*, fruitfully, in a way that teaches us much and also displays our conceptual mastery of them and their behavior. . .

even when the systems are in states such that the models are in no way predictively accurate

when, that is, there can be no question of the math representing the system in any depictive sense.

Every theory, for every kind of physical system it treats, has characteristic breakdown scales—spatial, temporal, energetic, *et al.*—beyond which the theory is no longer appropriately and adequately applied.

The predictive accuracy of Navier-Stokes theory breaks down as a fluid approaches turbulence.

Nevertheless, even while the dynamical equations of the theory no longer yield accurate predictions, other parts of the theory, the kinematical relations among its quantities (*e.g.*, that the shear-stress tensor is symmetric, and that heat flux is always independent of the pressure gradient), still are meaningfully applied to model some aspects and features of the system in those states.

It is those kinematical relations that we use, among other purposes, to guide the design of instruments to probe the systems, which shows that the theory can capture something of deep physical significance about the systems even when it is not predictively accurate of them.

But, again, there can be no question of depictive representation without predictive accuracy.

The root of all these problems is that contemporary philosophy has a fundamentally mistaken picture of what a physical theory is.

It is not a formal system in conjunction with rules of interpretation, or principles of coordination, or formal semantics, or sketches of representational capacities, or some such thing to bring the formalism into empirical contact with phenomena by virtue of relations of representation.

A physical theory is a body of inter-related knowledge and understanding, both theoretical and experimental, about some more or less clearly delineated part of the physical world (classical viscöelastic, thermoconductive fluids, *e.g.*, or weakly ionized, inviscid plasmas).

The theory is a structured network of our concepts about such things in combination with the practical mastery of experimentation on and observation of them.

It is our concepts that (we attempt to use to) represent. They are the Peircean symbols.

Math does not represent, but rather serves as the Peircean icon or index mediating the relation of representation between our thought (the Peircean symbol) and the world.

Math is a tool we use to bring our concepts and the world into contact.

Indeed, math provides us a wealth of different tools to use in order to bring our concepts and the world into contact (and that itself in a number of different ways), nothing more, nothing less. Some of those tools function in ways that superficially resemble standard ideas of representation, even though they are nothing like it in any important sense; others do not.

Math does not represent the world in any interesting sense—it rather mediates the way that our concepts attempt to represent the world.

When I say "math does not represent," I mean in a very broad sense: philosophers of science, realist and instrumentalist alike, tend to think that all math in physical theory gets its meaning in exactly (or primarily) one way—representation, whether they are realists about the ontology or not

and so, correlatively, that it can play only one role (or primarily one role) in giving an account of how we come to know about the world and what it is that we know when we do.

But that's not how terms in a theory's formalism acquire physical significance. They do so by the use we make of them as tools in mediating the connection between our concepts and the world, and how each functions as a tool is something we must work out on a case-by-case basis; it is not something fixed once and for all.

a sampling from the mathematical toolbox of general relativity

- contributes to T_{ab} (*e.g.*, Maxwell field)
- required for initial-value formulation of manifestly physical fields (*e.g.*, Maxwell field, g_{ab})
- dynamically couples to manifestly physical entities (*e.g.*, Maxwell field, g_{ab})
- dynamically couples to manifestly physical quantities that more than one type of physical system can bear (*e.g.*, Einstein tensor)
- acts as a measure of an observable aspect of manifestly physical entities (*e.g.*, Riemann tensor)
- enters the field equation of a manifestly physical structure (*e.g.*, Einstein tensor)
- constrains the behavior of a manifestly physical entity (*e.g.*, Killing field, conformal structure)
- plays an ineliminable (albeit physically obscure) role in the mathematical structure required to formulate the theory (*e.g.*, Einstein tensor)

Still, it sometimes seems difficult to deny that math "directly represents" in the depictive, verisimilous sense. Think, *e.g.*, of the ellipses of Kepler that Newton put to such profound use.

I believe we feel so strongly that the math "directly latches on to the physical world" in such cases because the physical concepts embodying the theory enlivens in our understanding the mathematical concepts constituting the formal structure we use as a tool to help us bring our concepts to bear in attempting to reason and learn about the world.

Part of the common way of grasping the physical notion of “circle”, *e.g.*, just is “body moving in a particular way”. That physical idea then is commonly brought to bear, for the most part unconsciously, in our picture of the mathematical concept of “circle”, in how we grasp and make use of the concept.

And so we mistake the physical content enlivening the mathematical concept for a relation of representation, for a fantastical capacity of the mathematical concept to reach out and make direct contact with the physical world.

How deeply the physical penetrates, grounds and enlivens the mathematical, in a way familiarized by the pedestrian-perceptual, how much we use the physical to understand the math, measures how strongly we feel the math "directly represents" the world.

This is why it is easy to think that the mathematical formalism of quantum theory "does not represent"—because the way we understand an element of a complex function space with a sesquilinear form is not enlivened by any physical concepts familiarized from the pedestrian-perceptual, nor the way we understand a linear functional on an abstract normed algebra with involution, closed under the weak topology.

van Fraassen, *The Scientific Image* (p. 8, his italics):

> *Science aims to give us, in its theories, a literally true story of what the world is like; and acceptance of a scientific theory involves the belief that it is true.* This is the correct statement of scientific realism.

My picture allows for a more satisfactory explication of the idea of realism.

Realism has nothing to do with representation—that "a theory tells us what the world would be like if the theory were true". It is rather the idea that:

1. our theories are such that they can teach us more about the world than how to predict experimental outcomes;
2. and we are in an epistemic state such that we can learn from our theories to understand the world in ways that lie deeper and that go beyond the prediction of experimental outcomes.

I can believe that “there is a way the world is”—and do so with good reason—without being a realist in the sense that I believe our best scientific theories are, with respect to ontology, a veridical or even merely adequately accurate representation of the world in any traditional sense.

Math does not give us verisimilous descriptions of the world to believe in.

Rather, we use math to confabulate ways of thinking about the world that conduce to understanding.