AVANT, Vol. X, No. 3/2019
ISSN: 2082-6710 avant.edu.pl/en
DOI: 10.26913/avant.2019.02.05

The Allegory of Isomorphism

Alistair M. C. Isaac
School of Philosophy, Psychology and Language Sciences
University of Edinburgh
a.m.c.isaac@ed.ac.uk

Received 29 November 2018; accepted 30 December 2018; published 12 December 2019

Abstract

Isomorphism has become a key concept for the analysis of representation in many contexts: perceptual experience, mental imagery, scientific theories, and visual artwork may all be described as standing in isomorphisms to their targets. Yet isomorphism is a technical term from mathematics—how are we to evaluate its use in fields such as philosophy, psychology, neuroscience, or physics? I suggest that we should understand appeals to isomorphism as *allegorical*; the upshot of this suggestion is that isomorphism claims always operate on two distinct levels of significance, with different standards of precision and evaluation. Recognizing these levels as distinct changes the landscape of debate for isomorphism-based accounts of representation: it both dissolves the well-known triviality objection to these accounts and undermines strong forms of structural realism.

Keywords: representation; structure; homomorphism; Newman's problem; structural realism.

1. Introduction

Analytic philosophy strives for precise terminology, in the expectation that clarity of conceptual (*a fortiori* metaphysical, epistemological) analysis requires an exactness of language. Nevertheless, this inclination should be tempered in interdisciplinary projects, where words must be flexible enough to apply in a variety of contexts. *Isomorphism*, for instance, is a term that has come to play a major role in accounts of representation across philosophy and the sciences—percepts, mental images, scientific theories, diagrams, and photographs may all be described as standing in isomorphisms to their targets. Nevertheless, as a term with quite varied standards of application in mathematics, logic, physics, and psychology, isomorphism claims cannot be evaluated with respect to a mono-dimensional semantic criterion—doing so would introduce misunderstanding and confusion rather than clarity.

Most philosophers first encounter isomorphism as a technical term in the set-theoretic foundations of logic—we are shown that the natural numbers and the rationals are the same "size" of infinity by establishing an isomorphism between them. But this is by no means the general notion of isomorphism, and misapplication of this narrow, set-theoretic understanding to richer contexts has resulted in the famous, yet misguided, worry that isomorphisms preserve only the cardinality of their targets—a result inconsistent with isomorphism claims in physics and psychology. If isomorphism is to serve a central role in the interdisciplinary analysis of representation across multiple domains and varieties of inquiry, we need a broad, unified conception applicable to all pertinent contexts.

I propose that isomorphism-talk in analyses of representation be understood *allegorically*. In literary analysis, *allegory* is the more sophisticated cousin to metaphor, simile, and analogy; where these techniques establish a simple correspondence between static items (*my love is* [*like*] *a rose*), allegory establishes a complex correspondence between rich, structured narratives (children pass through a wardrobe and face a series of trials against the backdrop of a battle between a lion and a witch ⟷ the convert confronts new ethical conflicts and faces a series of trials of faith against the backdrop of the struggle between Christ and the devil). Three characteristic features of allegory should inform our understanding of isomorphism. First, allegories operate on two levels of significance: the adventures of the children through the wardrobe may be assessed *qua* fantastical adventure of magic and violence **or** *qua* spiritual journey of discovery and conversion. Second, for an allegory to be successful, both levels require internal coherence: the story of the children must make narrative sense, as must the arc of Christian conversion it symbolizes have an internal logic. Finally, care must be made in identifying which aspects of an allegory are doing allegorical work, and which are not: the Turkish delights have allegorical significance *qua* temptation, but not *qua* confectionary.

Likewise, isomorphism claims function on two levels: a formal level, requiring assessment in accordance with the standards of mathematics, and a theoretical level, requiring assessment by standards appropriate for concept analysis and theory construction; and on both these levels, there are many moving parts and a rich "narrative" structure that must exhibit internal coherence. Nevertheless, we must be careful when moving between levels, and not illegitimately import (say) a feature of the mathematics of isomorphism into the conceptual claim, if it plays no allegorical role.

After establishing the basic features of isomorphism and the allegorical understanding of isomorphism-based theories of representation, I'll examine some applications of this view. I argue that well-known triviality objections to isomorphic accounts of representation dissolve, or are transformed, once we recognize that isomorphism claims operate on two levels. I conclude by critically assessing the use of isomorphism-based reasoning in philosophy of science, for instance to motivate *ontic structural realism*, arguing that this positive project also problematically confuses the two levels of the allegory of isomorphism.

2. Isomorphism: The Basics

In order to understand the proper role of isomorphism in philosophy, we must first understand it as a mathematical notion, and as one appealed to informally in fields such as psychology, neuroscience, and physics. There are three points to take away from this discussion: (i) isomorphism precisifies the intuitive notion of similarity or *relevant* sameness; (ii) it does so by defining a structure-preserving map between relevant features of two mathematical structures; as such (iii) it subsumes a family of technical notions that differ in their details, as relevant structure differs across subfield of mathematics. A moral is: it is a mistake to think any completely formal definition of isomorphism will be adequate to capture a rich concept such as representation across all domains, as relevant structure differs with context.

"Isomorphism" comes from the Greek: *ἴσος* ("equal") + *μορφή* ("form"). Use of the term to mean sameness of structure in English predates the technical notion in mathematics: in the early nineteenth century, "isomorphism" was used to refer to elements that crystallize into the same, or similar, geometrical structures. The earliest mathematical reference to isomorphism in the OED is 1892. Of course, structure-preserving maps appeared earlier in mathematics: anachronistically, they may be seen latent in the Euclidean notion of congruence by construction. More recently, the idea that correspondence of parts is an appropriate test for similarity or (relative) sameness is present in the mid-nineteenth century foundations of geometry (Riemann, 1868; Helmholtz, 1868) and of set theory (Cantor, 1874; Dedekind, 1888), under the guise of notions such as *Abbildung* ("mapping") or *Zuordnung* ("coordination," Ryckman, 1991).

In these early examples, one begins from the intuitive notions that (a) sameness (or similarity) means sameness of *structure*, and (b) this sameness of structure may be tested by establishing some kind of correspondence between parts. Exactly what structure is involved, and exactly what standard should be used to assess successful correspondence, differs across mathematical context. For instance, consider the claim above, that the natural and rational numbers are isomorphic—there is indeed a cardinality-preserving map between them, so they are isomorphic by the standards of set theory. However, natural and rational numbers have very different ordering properties—rationals are densely ordered (between any two there is another), naturals are not (there is no natural number between 2 and 3)—so they are not *order isomorphic*.

The paradigmatic notion of isomorphism in mathematics is that found in (abstract) algebra, which considers sets that are not only ordered, but also have operations defined over them that induce distinguished elements. The additive group of natural numbers ($\langle \mathbb{N}, + \rangle$) and that of positive rationals ($\langle \mathbb{Q}^+, + \rangle$) are not isomorphic in the algebraic sense, i.e. there is no function from one to the other that satisfies the condition:

> **(Algebraic) Isomorphism**: for sets with binary operations $\langle \mathcal{X}, * \rangle$ and $\langle \mathcal{Y}, \otimes \rangle$, an isomorphism is a bijective function $f: \mathcal{X} \longrightarrow \mathcal{Y}$ such that for any $x_1, x_2 \in \mathcal{X}$, $f(x_1 * x_2) = f(x_1) \otimes f(x_2)$.

In the algebraic case, the structures at issue are typically groups, or, when two operations are defined, rings. If we loosen the notion to include not merely operations, but an arbitrary number of n-ary relations, then we have a very general notion of isomorphism. Set-theoretic and order isomorphisms are just particular instantiations of this more general concept (where the set of relations is empty, or includes only one that is transitive and asymmetric).

The point here is just that the notion of isomorphism is underspecified until one has stipulated the relevant structure to be preserved. Strictly speaking it does not even make sense to ask if two structures are isomorphic if one does not already know at least the number and arity of relations at stake in the putative structure-preserving map. This is why (on some presentations) finite model theory, with the rigor characteristic of the logician, does not even accept isomorphism claims to be well-defined (and thus apt for truth value assessment) except between sets of the same *signature*, i.e. number of relations of some specified arity (e.g. Libkin, 2004). It is also the reason that many areas of mathematics have distinguished words for the particular flavor of structure-preserving map that identifies sameness of structure in the sense relevant to their domain, e.g. diffeomorphisms between smooth manifolds, homeomorphisms between topological spaces, or bisimulations between models for a modal language. In each case, the identity determining function at issue satisfies the general notion of an isomorphism, but a specialized term is used to indicate which structure the map preserves.

Another technical wrinkle that has generated more heat than light in the representation literature is the distinction between homomorphism and isomorphism. Technically, the term isomorphism is used only for bijective functions, functions that are both *one-to-one* (each element in the image is mapped to by only a single element of the domain) and *onto* (every element of the image is mapped to by some element in the domain). Coupled with the definitional stipulation that isomorphisms are *functions* (i.e. are defined for all elements of the domain, and assign each a single value), the restrictions of one-to-one and onto ensure that isomorphisms preserve cardinality. In contrast homomorphisms, though still functions, may be one-to-one, onto, or neither, consequently they do not necessarily preserve cardinality (so: homomorphism is the more general notion, with isomorphism a special case).

While the preservation of cardinality may seem a particularly important requirement for "sameness" if one is coming from the perspective of set theory, there are plenty of contexts in which relevant sameness and strict isomorphism come apart. An example from logic is bisimulation: two K-models may contain different numbers of worlds, yet still be bisimulation invariant, i.e. satisfy the same sentences in a modal logic (Blackburn & van Benthem, 2007). So, the relevant notion of sameness of models of modal languages is weaker than strict isomorphism. Conversely, bijections between continuous structures are not enough to ensure intuitive sameness, as one sees from the well-known result that any intuitive "size" of continuum is of equivalent cardinality to the unit interval (0,1). If one accepts that the formal elements that make up a continuous space have ontological significance as objects, and may be uniquely identified or "chosen" at will (i.e. one accepts the Axiom of Choice), then bizarre consequences arise, for instance, that one may slice a sphere into finitely many pieces and rearrange these pieces into two spheres each of equivalent size to the first (the Banach-Tarski Paradox). This odd example illustrates how irrelevant cardinality invariance is to our intuitive notion of "same size," *a fortiori* sameness or similarity in general.

So, there is a pre-theoretic notion of isomorphism, which predates the mathematical one, and whose meaning is etymologically transparent, namely "sameness of form." It is not a conceptually antecedent condition on sameness of form that there be sameness of number of primitive elements, rather this is an artefact of certain set-theoretic results and the ultimate technical restriction of the term "isomorphism" in mathematics to bijective functions. The crucial insight that engendered isomorphism as an important technical notion in mathematics, however, has only contingent bearing on cardinality. More fundamentally, this insight was that sameness of overall form may be established by placing primitive elements or parts of complex structures into correspondence, and locally checking that any relation holding between parts of one structure is mirrored by some relation holding between corresponding parts of the other. This basic idea has been cashed out in a variety of mathematical domains, differing in each case as to the relevant type of structure, or relation, which is preserved across correspondence, and only some of which satisfy the restriction to bijection.

3. Isomorphism as an Allegory for Representation

I claim that appeals to isomorphism in theories of representation are best understood allegorically. On the one hand, this means it is a mistake to assess such appeals by the evaluative standard for any particular mathematical notion of isomorphism, as any such use of "isomorphism" should not (in fact, *cannot*) be understood literally. Conversely, however, the mathematics of isomorphism, including procedures for proving isomorphism, may serve as a kind of blueprint or template for assessing these claims, indicating the various conceptual components, and implied lemmas, needed for them to be coherent or correct.

Typically, we understand by allegory a work of art that has both a surface and a deeper layer of meaning. *Animal Farm* and *Lord of the Flies* are generically novels insofar as they have characters, and a plot with a narrative arc that moves those characters through a sequence of events, but they are also allegories in the sense that these events may be interpreted as symbolizing more abstract claims about the nature of political power and social structure. While one can read and enjoy *Animal Farm* without perceiving this allegorical subtext, recognizing that subtext greatly deepens and clarifies the experience. Aspects of the novel that seem merely fantastical or borderline incoherent—animals that speak, pigs that walk on two legs—are thereby imbued with significance. Yet it is also the case that internal logic binds the novel together, rendering it effective *qua* novel: each animal has a distinctive personality that plausibly motivates its actions, thereby easing the reader through the anthropomorphism and enabling the suspension of disbelief necessary for the surface narrative to engage and compel. Arguably, one reason why the poems and plays that initiated allegory as a narrative form in the medieval period are no longer read much today is their lack of surface coherence—they work more poorly as allegories precisely because they cannot be enjoyed without explicit understanding of their allegorical subtext.[1]

Allegory is of a kind with analogy, yet involves richer and more complex relationships. Just as recognition of an illuminating, simple analogy can help reasoning in a complex domain, so also recognition of a more complex and dynamic allegorical correspondence can aid reasoning, so long as it is sufficiently robust. For instance, consider the use of analogical reasoning in physics, an example analyzed in depth by Mary Hesse (1961). She argued the appeal of the 17th century "mechanical philosophy" was in part that it allowed natural philosophers to reason about the unobservable by analogy with the observable: if matter at a microscopic level behaves just as the cogwheels or billiard balls we are familiar with in everyday life, we can apply our robust intuitions about the everyday to our theorizing about the microscopic, or intangible world. For Hesse, the post-Newtonian rise of mathematical physics is characterized by its appeal to a new kind of analogy, analogies between the natural world and mathematical models, e.g. the Maxwell equations. Like analogies with clockwork and billiard balls, mathematical analogies allow scientists to reason about the world indirectly, by reflecting on something "better known" than the target, in this case, mathematical equations. Unlike billiard balls, we don't necessarily have finely honed intuitions about mathematical equations, but this doesn't matter as mathematics contains within itself, in its very procedures, the means of ensuring internal consistency and probing logical implication. In this sense, we do (or in principle can) know mathematical models "better" than the unobservable world itself.

Likewise, when we claim that a map represents some geographical region because, or in virtue of, an isomorphism between map and region, we are using mathematics as a model to help us reason analogically about a conceptual problem, much as do Hesse's post-Newtonian physicists. I characterize this reasoning as allegorical rather than analogical, however, to highlight several distinctive features. First, the relationship between the mathematics of isomorphism and the concept of representation is more abstract than those in Hesse's case study. This relationship does not hold between some particular mathematical object, say an equation, and the world, but rather between a whole class of mathematical entities, the many possible types of isomorphism, and some generalized class of relationships in the world (*... is a map/depiction/photograph of ...*). Second, Hesse's focus was static mathematical objects, e.g. a particular set of equations,[2] while it is not just isomorphism *qua* static piece of mathematics, but the *establishment* of an isomorphism through *an act of proof*, that can constructively inform our philosophical reasoning about representation.

[1] For a detailed discussion of the history of allegory, and its importance in Medieval literature, see Lewis (1936).

[2] Note that equations such as Maxwell's may be "dynamical" in the sense that they describe a system, one of whose parameters is conventionally referred to as "time," or which is said to describe "change"; however, these are just evocative ways of talking about that which, from a mathematical point of view, is fixed: four lines of symbols unchanging in interpretation or implication from context to context.

Consider, for example, a topographic map of northern California. On some views (Camp, 2007; Rescorla, 2009) the map succeeds in representing the topography of northern California insofar as it stands in an isomorphic relationship to the altitude changes in that area. Unpacking this claim allegorically, we use what we know about isomorphisms as mathematical constructs to make sense of this putative relationship between map and world. Isomorphisms are relationships between structured wholes, so both the map and the world need to be understood as composed of elements that stand in relationships to each other. Furthermore, an isomorphism is successful when the relations of one whole are mirrored in the relations of the other, so we know a claim has been made about the (second order) relationship that holds between the (first order) relations in each structure, namely it is one of mirroring. Finally, we need a reason to think that there is a way to place the elements of the map and the elements of northern California in correspondence with one another, otherwise there is no reason to think that an isomorphism might obtain. These are three qualitative claims about the relationship between map and region—they are not intrinsically mathematical, though they have been inspired by our knowledge of the mathematics.

Determining the primitive elements and relations is particularly easy in the case of maps, as only distinguished symbols, often explicitly identified in the map key and scale, do any semantic work; so, we know exactly which features of the map participate in the putative isomorphism. The answer is not so simple for other imagistic representations, however—what are the primitive elements of a painting? *Color regions? Point color? Brushstrokes?* The surface of an oil painting will have a quite complex topography, with paint layered much more thickly in some areas than others—do these changes in paint "altitude" perform any representational function, i.e. do they participate in the putative isomorphism? The answer may differ from painting to painting—paint thickness my do no semantic work for van Eyck, but much for van Gogh. The allegory with mathematical isomorphism tells us to ask this question; it directs us to identify the qualitative features that do representational work, as a (conceptual) prerequisite for the claim of isomorphic representation to hold.

In addition to primitive elements, we must identify the participant relations. Point colors on a painted surface differ (at least) in position, hue, saturation, and brightness, and all of these differences may be assessed at different granularities. On a typical map, colors do representational work only insofar as sameness or difference of color indicates sameness or difference of category membership, as when trainlines on a subway map or countries on a geopolitical map are distinguished by color. The topographic map is different, as colors may indicate changes in altitude by continuous transitions in hue. In a representational painting, color may do even more elaborate work, and correspondingly more elaborate relations between point colors participate in the putative isomorphism, e.g. two shades of brown may, through proximity in position and hue, depict different aspects of the same surface, such as regions of a table in light and shadow.

However difficult the task of identifying the primitive, isomorphism-relevant elements and relations of a representing vehicle may be, that of identifying this structure in the target is far more subtle and contentious. The significance of this problem for isomorphism accounts we address in detail in the following section. For now it is enough to observe that the claim of isomorphism presupposes that some primitive elements and relations in the target (northern California, the subject of the painting, …) are distinguished, for some such distinguished components are needed to correspond to and mirror the relations of the primitive elements in the map, painting, photograph, etc.

The conceptual/theoretical level of an isomorphism claim instructs us to (informally) identify the elements and relations in vehicle and target, as it suggests that representational success is equivalent to the existence of a certain type of correspondence between these respective items. Since, at this level, our goal is to analyze the concept of representation, we may not need to specify any specific type of mathematical isomorphism as our primary analog. Already, by allegorical reasoning, we know we need a prior specification of the elements and relations, we know they must exhibit structural similarity between vehicle and target, and moreover that there must be some procedure or process by which this structural similarity may be verified—these are features present in *any* mathematical isomorphism, even under the broad construal above.

Nevertheless, any particular example of representation by isomorphism may also be assessed as a "literal" claim about the relationship between two, specific mathematical structures. While neither a map, nor a painting, nor northern California are themselves mathematical, all three may be *modeled* mathematically. Once, say, northern California and the map are both modeled mathematically, one may ask of these models whether they are isomorphic to each other. This project is distinct from that of conceptual analysis, but it may complement and inform it.[3] If a modeling project like this is to be successful, then it may involve proof, and should certainly be held to the very specific standards of precision that mathematics demands. Nevertheless, at some point the details of this project will cease to serve an allegorical function. Consider again the case of maps: a project to formally model the isomorphic relationship between a topographical map and Northern California, where continuous gradations of color and distance are preserved across the representational mapping, and one to formally model the isomorphic relationship between the London Underground and its famous map, where only topological relations and brute (color) category membership are preserved, will look totally different, involving radically different definitions. Nevertheless, both projects, if successful, would allegorically support the conceptual claim that map-like representation succeeds in virtue of an isomorphism between target and map.

Finally, as mentioned above, I think that proof strategies for establishing isomorphism in mathematics should be recognized as a critical part of the allegory of isomorphism. Such proofs are often dynamical and constructive, as in model-checking algorithms for bisimulation or proofs by construction in geometry, and this dynamical character should inform

[3] For an example, see Suppes, Perreau-Guimaraes, and Wong (2009), where mathematical modeling and conceptual analysis explicitly inform each other in an analysis of the isomorphism between neural and perceptual representations of language.

our theory of how and where representation occurs. Consider, for instance, projective geometry and linear perspective. Linear perspective is one mode of pictorial representation, whereby depicted objects are foreshortened as they would be if projected from a point vantage onto a 2-dimensional surface. This relation is typically a many-to-one map that preserves *betweenness*, and as such is an instance of homomorphism, or in our general sense above, an isomorphism.[4] When assessing the representational relationship between a photo, or a painting in linear perspective, and its target, it is not just the mere existence of a mapping between the two which inspires our intuitions, but our understanding of the *mechanism* by which that mapping has been established, precisely because this mechanism appears to physically realize the mathematical procedure for calculating projections of geometrical objects. In a camera, for instance, the straightness of rays of light plays the same role in the physical mechanism of image formation as the stipulated straightness of geometrical rays, implemented by stylus and straightedge when calculating projections on paper.

In fact, in many cases, we believe representation to be grounded in isomorphism because we understand the mechanism for establishing a representational map, even if we don't (yet) have access to the primitive elements and relations of the target, or even vehicle. Consider, for instance, the human visual system. We know from the mechanics of the lens that a scene is projected onto the retina, and from the neural wiring of the retina and beyond that the activation of cones and rods is projected in parallel through the LGN to higher visual processing areas V1 through V5. The nature of the wiring may convince us that this is a structure-preserving map of some sort, and thus that neural representation of a visual scene is grounded in isomorphism. Nevertheless, there is still great uncertainty about exactly which features are extracted in higher areas of visual cortex and, correspondingly, exactly which features in the target visual scene they represent. Here, it is as if we have a "proof" of isomorphism, since we have a map-establishing mechanism, but we don't yet have a full qualitative understanding of the structural relations preserved across this mapping.

So, the allegory of isomorphism directs us to establish a number of qualitative ingredients for any putative representation grounded in isomorphism. These ingredients include the primitive elements and relations that characterize both representing vehicle and target, as well as the mechanism that establishes their correspondence. We may also be able to establish a literal, formal isomorphism between mathematical models of vehicle and target. However, for broad, conceptual claims about representation, such specific models are not only unnecessary, but potentially misleading. Rather, it is a whole class of mathematical mapping relationships, and their proof strategies, that do the work to motivate a theoretical account of representation as isomorphism. This is the justification for calling the role of

[4] Projective maps preserve structure, e.g. collinearity of points and tangency of surfaces, and are thus an instance of a structure-preserving map, which we have been calling "isomorphism" in a loose, qualitative sense. Of course, they are not isomorphisms in the strict, bijective sense, or even isometries (preserving distance), but the fundamental fact that they systematically preserve structure is nevertheless important, as it grounds human and computer vision (e.g. Hartley & Zisserman, 2004), and this I think warrants treating projective maps as a special case of the general concept of isomorphism. I stress this point because some have argued that linear projection is an alternative to isomorphism rather than a paradigmatic instance of it (Greenberg, 2013).

mathematics here allegorical (rather than analogical)—general strategies for modeling and proof do the inspirational work, not specific sets of equations.

4. Structure in the World, neither Given nor Trivial

Suppose we accept that isomorphism claims are allegorical, what work can this do for us? I think it should change our perspective on some of the key debates concerning isomorphism and representation. On the critical side, perhaps the most well-known objection to isomorphism accounts is the triviality objection: isomorphisms come too cheap to do any explanatory work. I take this objection to be deeply confused, and in particular to commit the error of importing mathematical details of isomorphism into the conceptual theory that do no allegorical work, and are thus worse than irrelevant: they are misleading. To see why, we must first return to the problem of the target, and look at some examples of how it is resolved in practice.

The most subtle conceptual problem for isomorphism accounts is the question of the structure of the target. Targets in the world—subjects of paintings, mapped geographical regions, phenomenon described by scientific theories—do not come to us parsed already into primitive elements and relations. Nevertheless, our talk of isomorphism presupposes some such structural analysis of these targets, for it is only if a target is conceptualized as some set of elements standing in distinguished relations that the claim of isomorphism is meaningful. I believe the allegorical interpretation helps make progress on this problem.

On the allegorical account, we recognize that the question of structure in the target should be approached on two levels. At the conceptual/theoretical level, we assert that some distinguished structural analysis of the target is possible, such that the components of this analysis are mirrored in those of the vehicle. At the formal level, we assert that this structural description may be realized in a mathematical model, a set of elements standing in precise relations, and that this model may be put in an isomorphic correspondence to a model of the vehicle. Crucially, the allegory of isomorphism does not depend for its success on the details of this modeling endeavor, since its role is to direct us to seek qualitative elements and relations distinguished in the target and vehicle, and a mechanism that induces a mapping between them. Thus, we may make sense of the claim of isomorphic representation, even when we do not (yet) have an account of the distinguished elements and relations of the target that are preserved across a representational map. Yes, some target structure is presupposed by the assertion of isomorphism, but in practice, this structure might be discovered only after this assertion has been made.

Consider again the example of the visual system, in particular the well-known Hubel and Wiesel project to uncover the functional architecture of V1 (Wurtz, 2009). Let us (re)describe their first experiments in the late 1950's from the perspective of the allegory of isomorphism. Hubel and Wiesel began from an assumption of isomorphism, one motivated by their understanding of visual physiology. They knew there was a mechanism that mapped the structure of the visual field through the retina and into the cells of V1. Likewise, by the assumptions of contemporary neuroscience, they took the firing rate of individual cells to be the primitive elements of the representing vehicle. What they didn't know were the primitive elements and relations that defined the target, i.e. the distal visual scene. They inserted electrodes into feline V1 in order to make single cell recordings, thereby attempting to reverse engineer the assumed map. For each such neuron (representing element), there should be a primitive element in the target to which it corresponds, a correspondence measurable by neural activation, or rate of fire. As related in interviews, Hubel and Wiesel tried to stimulate cat V1 cells with a variety of simple stimuli, i.e. what they took to be candidate elements in the target: spots of light, spots of dark, randomly chosen images, etc., without success. Only by accident one day, when a slide was pushed too far within the projector, initiating a moving line across the cat's visual field, did a monitored neuron fire, serendipitously revealing the actual elements of the target represented in V1: edges moving at specific orientations (Hubel & Wiesel, 1959; Hubel, 1996).

Contrast this with a second example, the studies of Roger Shepard on mental object rotation, which initiated a fruitful research program on mental imagery (Shepard & Cooper, 1982). In the case of Shepard, there is no need to *re*-describe his early advances in terms of "isomorphism," as he explicitly championed use of the term in psychology (Shepard & Chipman, 1970). Shepard performed a series of studies demonstrating an isomorphism between the rotation of target objects in space, and corresponding transformations over mental representations of those objects. He called this a "second-order" isomorphism, to emphasize that a structural correspondence obtains between the dynamic pattern of transformations over a mental image and the pattern of spatial rotations over their target—in contrast to any first-order isomorphism that may or may not obtain between the static shape of a target object and the static features of a structured mental representation. A key early datum in these studies was the lag time it took a subject to judge whether two target (two-dimensional) images depict different rotations of the same three-dimensional object. Shepard demonstrated that the *time* it took a subject to make the judgment varied linearly with the *angular distance* of rotation between the two images, i.e. that relative distance maps to relative time in a sequence of mental representations of spatial objects (Shepard & Metzler, 1971).[5]

In both these examples, only part of the isomorphism story is initially available. In the Hubel and Wiesel case, the mechanism inducing a mapping is understood, as are the primitive elements of the representational vehicle, while the elements and relations in the target are only gradually discovered through trial and error. In the Shepard case, there's no theory of a mapping mechanism, nor of the primitive elements that define the vehicle. But the design of the experiment proceeds on the assumption that spatial orientation is a distinctive feature of the target, and demonstrates that a temporal relation in the vehicle structurally corresponds to the degree of target rotation. This empirical result holds even without a

[5] Suppes et al. (2009) exhibits features similar to this example, as their result, a structure-preserving map between perceptual and neural representations, does not depend on an account of the primitive elements in either domain, but merely orders of (dis)similarity over stimuli.

theory of the primitive elements of the vehicle, other than that they are related somehow by transformations extended in time. In both cases, the allegory of isomorphism provides an explanatory framework that demands certain details be filled in, and its value is confirmed by the productivity of this filling-in project, not by mere possession of any particular, static set of isomorphic models.[6]

This discussion should help provide a new perspective on the most influential line of criticism against isomorphism accounts of representation: the triviality worry. The line of argument goes something like this: isomorphism is cheap; (just about) any two things are isomorphic; so, to say that *A* represents *B* because *A* is isomorphic to *B* is to say something trivial, by which *A* represents not only *B*, but (just about) everything else. Often, these worries emphasize cardinality: *A* is isomorphic to anything of the same cardinality (of which, presumably, there are many). I have in mind here a variety of arguments, not necessarily targeted at theories of "representation" *per se*, but still addressing the kind of semantic question a theory of representation purports to answer.

An early instance of this worry is Newman's (1928) objection to Russell's structuralist theory of knowledge in *The Analysis of Matter*. Russell claims, since perceptual experience stands in a structural correspondence to nature, all we can know of the external world is its "structure." Newman points out that, without any restriction on the relations obtaining between parts of the distal world, i.e. if we truly cannot know what any of those relations are intrinsically, then in fact all we can know of the unobservable is its cardinality. This follows because sameness of structure is isomorphism; relations are defined extensionally; and so a 1-to-1 map may be defined between any two sets of same cardinality, as sufficient "structure" in the second corresponding to that of the first may always be defined:

> Any collection of things can be organised so as to have the structure *W*, provided there are the right number of them. Hence the doctrine that only structure is known involves the doctrine that nothing can be known that is not logically deducible from the mere fact of existence, except ("theoretically") the number of constituting objects. (p. 144)

One may see this as a problem for a theory of representation, if one takes Russell to have claimed that perceptual experience represents only the structure of the world.

Two similar arguments are due to Putnam. Putnam (1980) employs the Löwenheim-Skolem Theorem to argue that our set-theoretic language, *a fortiori* all language, fails to pick out its target. The Löwenheim-Skolem Theorem tells us that any model of a first-order language may be arbitrarily extended to one of larger cardinality, thereby ensuring the language fails to uniquely capture its intended interpretation. By extension, any theoretical language fails to pick out its intended model, because "there will be many models—models differing on the extension of every term […]—which satisfy the entire theory" (p. 477). Again, mere assumption of isomorphism doesn't say much at all, because of the vast number of models isomorphic to any intended target. Once one moves to the infinite

[6] See Shepard (1981) or Smith & Kosslyn (1980) for a detailed defense of this type of methodology for investigating mental imagery.

domain, the one relevant for the semantics of mathematics, then Löwenheim-Skolem shows even cardinality is not uniquely picked out by any description.

Putnam (1988) applies similar reasoning to a more concrete case, arguing that any physical system is a model of any finite state automaton. The "proof" proceeds by assuming some continuity in the behavior of the system, then *ex post facto* carving this behavior into a sequence of discrete states, labeled to match the states of the desired FSA—a move guaranteed possible by the assumption of continuity. If one takes FSA's and other formal descriptions of computation to represent particular sorts of devices (computers) or processes (computations) on the basis of mere structural correspondence, then his argument shows those claims to be trivial—computation is cheap, so (for instance) identifying cognitive processes by mere computational description fails to uniquely identify the desired target. Likewise his 1980 argument: if one takes a theory, or set of true statements, to represent the world up to isomorphism, then Putnam has shown that such representations (*true theories!*) are satisfied so easily they fail to identify the intended state of the world.

Is the allegorical understanding of isomorphism-based representation susceptible to this type of criticism? The quick answer is *no*, precisely because the allegory of isomorphism already endorses two fundamental lessons that follow from this line of reasoning: (1) the claim of isomorphism is only meaningful if one assumes some antecedent structure, and (2) representation rests not on mere, "in principle" existence of an isomorphism, but on some *reason* to expect that isomorphism to obtain, e.g. a map-inducing mechanism. However, this answer is far too glib, for the allegory of isomorphism does countenance the idea that (3) all that may be represented (and thus for, say, a mental image, all that may be known from it) is structure. But the arguments of Newman and Putnam purport to show (1), (2), and *not* (3), or even that (3) is nonsensical.

I think Newman and Putnam go wrong by confusing the two levels of the allegory of isomorphism. We must always keep the level of mathematical models distinct from that of conceptual analysis, lest we become confused about which aspects of the mathematics do allegorical work, and which do not—remember: *eating Turkish delight is not always a sin!*

For instance, Putnam's argument that every physical system implements an FSA depends on treating a particular mathematical model of a generic physical system as equivalent to the system itself, i.e. on drawing a conclusion about the conceptual level from a specific result at the mathematical level. Putnam does not start from an assumption that some structure in the system is isomorphic to an FSA, and then seek to *discover* what it may be (a project that would have been analogous to that of Hubel and Wiesel—though, one suspects, less successful). Rather, he defines a model that *by stipulation* realizes the desired states (at the cost that their physical analogs are only accessible to "a Laplacian supermind" [1988, p. 122]). This definition ensures he can prove an isomorphism, but the strategy of this argument is disanalogous to typical isomorphism proofs in mathematics. For, as emphasized above, isomorphism claims are only meaningful in reference to *antecedent* structure. Consider again the question of the order-isomorphism of the natural and rational numbers. This question is understood with reference to the antecedently defined orderings $\leq$ over N and Q^+.

While one *could*, Putnam-style, define some new order over Q^+ that would be preserved over a mapping into N, that process wouldn't convince anyone Q^+ is order-isomorphic to N, as it ignores rather than answers the question at issue. Likewise, Hubel and Wiesel are not at liberty to stipulate structure in the visual scene isomorphic to firing in V1, rather they are constrained to discover this structure amongst only those features of the scene that are in causal contact with V1 through the mechanism of visual wiring. Putnam has mistaken the stipulated features of one possible mathematical model of a physical system for general, discoverable features of all such systems. The allegory of isomorphism indeed demands that such physical features be susceptible to mathematical modeling, but this does not license an inference from contingent features of one particular model back to broad qualitative theory. Putnam has got his allegorical reasoning backward.

This response to Putnam asserts that only some structure is pertinent to isomorphism claims, and as such has an affinity with the influential views of Lewis (1984). "Yes," Lewis admits in response to Putnam (1980),

> your formal reasoning is correct, but it assumes that all relations are equal—for it is only then that you can claim so vast a set of targets will satisfy our representation. However, this is not what we (realists) suppose the world to be like, for we take some relations to be privileged over others; these are the "natural" relations, and only these are "eligible" to define plausible models. (p. 227)[7]

The followers of Lewis have taken him to imply here that some metaphysical property, a "reference magnetism," makes these eligible relations attractors to our representations (Sider, 2009; Schwarz, 2014; Azzouni, 2017). This view, that we need both an account of distinguished relations in the world *and* a reason to think we have latched on to those relations has appeared repeatedly in response to the triviality worry. Russell himself, for instance, wrote a letter to Newman in 1928, acknowledging the force of his criticism, and offering "co-punctuality" as "a relation which might exist among percepts and is itself perceptible," i.e. may be known not only structurally, but also for what it is in the world (1968, p. 260). Likewise Carnap (1928, §154) combats a Newman-like triviality worry by insisting that some "experienceable, 'natural' relations" are "founded," meaning that they are the same as experienced and as they are in the world (Demopoulis & Friedman, 1985; Bueno, 2017). The Lewis of his acolytes, Russell, and Carnap all capitulate to Newman by granting that the claim that *only* structure is represented or known is indeed incoherent.

The answer from the allegory of isomorphism is subtly different, however. In particular, the relations that the allegorical view asserts are distinguished need not have a special metaphysical status. Rather, they are taken to have a critical *procedural* status, in virtue of their role in constraining our project to fill out the explanatory framework suggested by mathematical isomorphisms. This procedural constraint then implies an epistemic constraint: only

[7] In philosophical parlance, only some relations "carve nature at the joints," or are, in Goodman's terms, "projectible." Lewis takes this response from Merrill (1980), who calls these relations "objective" (pp. 71–72), and shows that building them into model theory blocks Putnam's proof.

some relations are knowable—not because only these relations are natural, reference magnets, or whatever, but rather because only these relations have been antecedently accepted as candidates for structure preservation.

It is easy to see first how this account works in empirical examples. In the case of Hubel and Wiesel, for instance, their experimental set up—cats constrained to view images projected on a screen—determines a restricted set of possible relations which might be discovered in the target. Likewise Shepard restricts his attention to relations in the vehicle manifest in behavior, such as time of response. In general, when we wish to analyze any representation (picture, map, scientific theory, whatever) as isomorphic to its target, we implicitly assume some antecedent structure in both vehicle and target. While we may not have explicit mathematical models in hand, this assumption is typically the result of our expectation about what kind of model is possible. I assume the subway system is topologically connected, or the subject of the painting may be described geometrically, or the target of my theory is structured causally, etc.

What happens when we move beyond this kind of concrete theory building to its philosophical implications—if mental images structurally represent the world, or if successful scientific theories are merely isomorphic to the world, what then can we know about the world? The allegorical view claims: we can only know structure, but we are constrained to represent this structure to ourselves by our antecedent modeling choices. Just as in the experimental examples, when we develop scientific theories or form mental representations, we are procedurally constrained. In the former case, the constraints come from empirical method and available mathematics; in the latter case, the constraints are physiological, ecological, and evolutionary.[8] In both cases, as epistemic agents, we recognize these constraints as defining a space of ways the world can be (possible precisifications into models), and take our representation to pick out one of these. Yet, knowing also that these modeling choices are indeed merely choices, we must acknowledge that all we know of the world *tout court* is structure. So, the claim that all we know of the world is structure is not trivial, because it is indexed by our intended interpretation. Nevertheless, this indexing does not come from the world (in the form of special reference fixing properties, or joints-to-be-carved), but rather from our own practice, and in particular our understanding of our epistemic status as allegorical: as comprehensible by analogy with a pair of isomorphic models—of the world and our representational state, respectively—whether we have these models in hand or not.

5. Mathematical Representations

Most of the debate about Newman-type problems has explicitly concerned mathematical representations; for instance, both models of computation, such as finite state automata, and scientific theories, at least in physics, take the form of sets of equations or rules in a formal

[8] Contra Domopoulos & Friedman (1985), some readings of Carnap's project interpret foundedness along these lines (Leitgeb, 2009; Toader, 2015).

language. These examples seem especially confusing for the allegorical view, as I have argued that there are two levels of structural mapping in an isomorphic representation: theoretical/conceptual and modeling/mathematical. On which of these levels does a mathematical theory live? *Both*, a point which is not so much problematic in itself, but insofar as it has led some to confound mathematical models *of the target* with the target itself, a confusion that clouds several key questions in philosophy of science. We can see these issues in a clearer light if we bear in mind that any mathematical representation is also a model, yet *qua model* necessarily describes the target in a manner at best partial and idealized.

The allegory of isomorphism asserts that the theoretical relationship between vehicle and target is analogous to a mathematical relationship between models of both vehicle and target. When the representing vehicle is itself mathematical, then it may be taken as its own model.[9] This means that when reasoning about the vehicle, it is easy to conflate the two levels of the allegory, and thereby inappropriately treat mathematical models of the target as if they are the target itself. We already saw Putnam (1988) make this error, treating a physical system (the target) as if it were a piece of mathematics, and thereby reasoning fallaciously from features of a specific model to the physical system it attempts to capture. I suspect he was seduced into this line of reasoning by the mathematical nature of the vehicle at issue, a model of computation defined in a formal language.

A broad, general class of representations at risk for this kind of problematic reasoning are scientific representations, which—at least in much-discussed examples from physics—are typically (sets of) equations (Maxwell's equations, Newton's laws of motion, Boyle's law, the Schrödinger equation, etc.). Since the mid-twentieth century, philosophy of the mathematical sciences has been dominated by the "semantic view," which holds that a scientific theory should be identified with the set of models that satisfy it (Suppe, 1977). This movement has emphasized isomorphism from the beginning as the key concept for formalizing philosophical questions about science, initially questions about reduction, equivalence, and other relations amongst mathematical theories (Suppes, 1960, 2002). However, from early on, as adherents of the semantic view turned to questions of realism and representation, they considered isomorphism also a candidate relationship between theory and world (van Fraassen, 1980).

From the allegorical standpoint, this approach to the realism question is ill-conceived, so long as isomorphism is intended in the strict, mathematical sense. Of course, considered as a mathematical object, equations characterizing physical theory (say Newton's laws) should indeed be (strictly) isomorphic to some mathematical model of the world. But if we wish to investigate the relationship between Newton's laws, considered now in general as a scientific representation, and the world, considered as physical reality independent of our interests and practices, then the most we can hope for is a loose "isomorphism" in the

[9] Although it need not be—consider, for instance, the practice in climate modeling of replacing complex or intractable pieces of theory (e.g. the Navier–Stokes equations) with discretized approximations. Here, arguably, the representational relation between theory and world detours through mathematical models that themselves simplify and idealize the full mathematical representation (Lloyd, 2010; Parker, 2010).

sense of a qualitative and intuitive structure-preserving map. To seek a mathematical isomorphism, a precisely defined bijective function, between a mathematical theory and the world is to mix apples and oranges—precisely because, as discussed above, the world does not come to us pre-parsed into elements and relations: it is not itself a mathematical object.

Once we realize this point, we see that some debates on both sides of the semantic view are ill-conceived. Within the semantic view itself, for instance, it is vain to seek the exact formal characterization of the mapping that obtains from theory to world. In particular, an extended literature has attempted to specify that mapping in terms of "partial" maps or structures, culminating in the nuanced view of Bueno, French, and Ladyman (2002) that the relationship between mathematical physics and the empirical world is one of "partial homomorphism." This literature acknowledges that formal mappings may only be established between mathematical theories and models of the world, yet it fails to confront the corollary that any such model is only one of many possible. The typical target on offer is a "model of the data," but a data set is itself only one amongst many possible models of a phenomenon, and as every statistician knows, the same data may be modeled many different ways. This implies first that no formal mapping between some particular theory and some particular model of a data set can possibly exhaust the rich question of the relationship of that theory to reality. Second, it implies as corollary that no general formal characterization of such mappings (as partial homomorphism, or whatever) can capture the full panoply of relevant mappings that might obtain between theories, which may employ arbitrarily diverse formalisms, and the many possible models of their intended targets in the world, which subsume all possible methods for data collection and analysis.

Conversely, criticisms of the semantic view as overly mathematical in method also step wrong, failing to recognize the reciprocity between formal models and qualitative theory. For instance, Giere (1988) and Godfrey-Smith (2006) argue the semantic view is inappropriate as an account of scientific method and overly restrictive in demanding the theory–world relationship conform to a formal notion of isomorphism. They both suggest the representational content a theory bears about the world is grounded in some looser notion of "resemblance" or "similarity." Yet, if their loose notion of resemblance involves structural correspondence of some sort, then it confirms rather than contradicts the isomorphism view, at least as understood allegorically. Moreover, once we recognize that Giere and Godfrey-Smith's arguments apply solely to the conceptual/theoretical level of the allegory, while the formal models developed within the semantic view apply at the modeling/mathematical level, the two projects reveal themselves as complementary (cf. French & Ladyman, 1999).

A final example worth considering is the recent enthusiasm for *ontic structural realism* (OSR) in philosophy of science (Ladyman, 1998; McKenzie, 2017). OSR subsumes a family of positions that take structure to be more fundamental than relata, either claiming merely that the properties of relata supervene on their relations, or that relata may be eliminated altogether. The motivation for these views is typically mathematical physics. On the one hand, OSR appears to offer a structuralist response to Kuhn's worry that theory change brings radical changes to ontology, by asserting that nevertheless the mathematical structure of theories has remained largely constant across such changes (Worrall, 1989). On the other hand, the observation that in contemporary physics, entities (such as electrons or spacetime points) tend to be defined solely in terms of their positions within a complex mathematical structure, and are otherwise "indiscernible," supports the claim that fundamental physics postulates only structure in the world (Muller, 2011; Glick, 2016).[10]

There are many worries about the coherence of OSR as a metaphysical position (Wolff, 2012; McKenzie, 2017). Here, however, I want to question, from the allegorical standpoint, the strategy of reasoning from the mathematics of physics to overly precise claims about the natural world. Of course, mathematical physics, especially its most magnificently well-confirmed corners, such as quantum field theory, must be understood as describing the world precisely by a realist. Moreover, the realist should accept these theories as more than merely empirically adequate, and attribute reality to the entities and relations they posit. In the language of isomorphism, then, a realist should accept that the world contains structure corresponding to that of any well-confirmed theory. What the allegorical story adds, however, is a warning not to import aspects of the mathematical structure into this mapping that do no allegorical work. When proponents of OSR reason about the ontology of the world on the basis of what is *indiscernibl*e or *symmetrical* in mathematical physics, they reason from absences and identities in the mathematical model to a positive claim about absences and identities in the world. But here I think one overstates what can be derived at the qualitative/conceptual level from a successful mathematical representation—for we can demand that structural details of the representation be preserved to a high degree of precision in the target, but not so for absences or equivalences, which demand more than good precision or close fit. Equivalences, identities, and symmetries all depend on the *exactness*, or *complete* precision of mathematical definitions, yet complete precision is not a feature of our access to the empirical world. Measurements, and other attempts at theory testing or model fitting, are always bounded to some threshold—represented by range of variance, or number of significant figures in an assigned value. Features of a mathematical representation that depend on an arbitrarily high degree of precision—precision without threshold—are only meaningful as features of the representation *qua* model. Yet, it is precisely those features distinctive of the representation *qua* model that we should not expect to map to the target itself, but only to our models of the target.

These considerations are really just a special case of a point well-known, even foundational, in the modeling literature: models are inherently partial, idealized, or even fictionalized (Levins, 1966; Bokulich, 2011; Weisberg, 2013). The allegory of isomorphism fully endorses this general point, for it recognizes that the mathematical level of description sacrifices generality and concreteness for precision and idealization. Only at the idealized level of mathematical models may we prove precise isomorphism theorems. The corollary, however, is that the qualitatively understood vehicles and targets that models precisify are necessarily more nuanced and rich. Consequently, we may never reason from mathematical

[10] Note: fundamental physics is not the only route to OSR (see Frigg & Votsis, 2011), but it is the only one under consideration here.

models to the actual world on the assumption that the model is ontologically complete, or precise to an arbitrary degree, for it is exactly these features that are abandoned as a matter of principle by the very act of modeling.

6. Conclusion

The allegory of isomorphism is a view about how the mathematics of isomorphism should inform a theory of representation. Etymologically and historically, an “isomorphism” is a sameness of form; the mathematics of isomorphism precisifies this idea by showing how form may treated as “structure,” or set of elements standing in relations, and so sameness of form may be assessed by placing the elements of two structures in correspondence and checking if relations in one mirror those in the other. A qualitative theory of representation takes from this allegory three morals: (1) representational vehicles and their targets must both be susceptible to analysis into sets of primitive elements standing in relations; (2) representational success requires that (some distinguished set of) the relations in a vehicle mirror those of the target; (3) we should only accept isomorphisms that hold for some reason, for instance if there is a mechanism that induces a structure-preserving map from target to vehicle.

Mathematics plays two distinct roles in the allegory of isomorphism. First, as allegory: we take inspiration from the broad array of isomorphism-type equivalence relations and proof strategies to be found across many areas of mathematics. All of these exhibit, support, and inform morals (1), (2), and (3), yet no particular equivalence relation should dominate our qualitative theory of representation. Second, particular representational vehicles and their targets may be modeled mathematically, and isomorphisms between these models proved formally. This exercise can be useful in multi-domain mathematical sciences, such as neural computation, and in formal philosophy, as when demonstrating one theory reduces to another under the aegis of the semantic view. Nevertheless, we must be careful not to reason backward from particular mathematical models to conceptual claims about representation. This error of reasoning, a reversal of the allegory of isomorphism that mistakes the particular for the general, is behind both the triviality objection to structural representation and the positive project of ontic structural realism.

The allegory of isomorphism is thus an approach both reverential and temperate in its attitude toward mathematics. It takes the mathematics of isomorphism as an exemplar, and mathematical models as the *sine qua non* of any precise and complete theory of a particular representational relationship. Nevertheless, it warns against the temptation to reify these models, for even successful models invariably contain contingent mathematical idiosyncrasies that should not be exported unreflectively to our theory of representation in the world.

Acknowledgements

This project has benefitted from the support of the Alexander von Humboldt Foundation and the Munich Center for Mathematical Philosophy.

References

Azzouni, J. (2017). Reference Magnetism. In *The Rule-Following Paradox and its Implications for Metaphysics* (pp. 55–72). Cham, Switzerland: Springer. https://doi.org/10.1007/978-3-319-49061-8_4

Blackburn, P. & van Benthem, J. (2007). Modal Logic: A Semantic Perspective. In P. Blackburn, J. van Benthem, & F. Wolter (Eds.), *Handbook of Modal Logic* (pp. 1–84). Elsevier. https://doi.org/10.1016/S1570-2464(07)80004-8

Bokulich, A. (2011). How Scientific Models can Explain. *Synthese*, *180*, 33–45. https://doi.org/10.1007/s11229-009-9565-1

Bueno, O. (2017). Overcoming Newman's Objection. In M. Massimi, J.-W. Romeijn, & G. Shurz (Eds.), *EPSA15 Selected Papers* (pp. 3–12). Cham, Switzerland: Springer. https://doi.org/10.1007/978-3-319-53730-6_1

Bueno, O., French, S., & Ladyman, J. (2002). On Representing the Relationship between the Mathematical and the Empirical. *Philosophy of Science, 69*, 497–518. https://doi.org/10.1086/342456

Camp, L. (2007). Thinking with Maps. *Philosophical Perspectives, 21: Philosophy of Mind*, 145–182. https://doi.org/10.1111/j.1520-8583.2007.00124.x

Cantor, G. (1996 [1874]). On a Property of the Set of Real Algebraic Numbers. In W. B. Ewald (Ed.), *From Kant to Hilbert: A Source Book in the Foundations of Mathematics* (Vol. 2, pp. 839–842). Oxford, UK: Clarendon Press.

Carnap, R. (1967). *The Logical Structure of the World*. (R. A. George, Trans.). Berkeley, CA: University of California Press. (Original work published 1928)

Dedekind, R. (1996). Was sind und was zollen die Zahlen? In W. B. Ewald (Ed.), *From Kant to Hilbert: A Source Book in the Foundations of Mathematics* (Vol. 2, pp. 787–832). Oxford, UK: Clarendon Press. (Original work published 1888)

Demopoulos, W., & Friedman, M. (1985). Bertrand Russell's The Analysis of Matter: Its Historical Context and Contemporary Interest. *Philosophy of Science, 52*, 621–639. https://doi.org/10.1086/289281

French, S., & Ladyman, J. (1999). Reinflating the Semantic Approach. *International Studies in the Philosophy of Science, 13*, 103–121. https://doi.org/10.1080/02698599908573612

Frigg, R., & Votsis, I. (2011). Everything You Always Wanted to Know about Structural Realism but were Afraid to Ask. *European Journal for Philosophy of Science, 1*, 227–276. https://doi.org/10.1007/s13194-011-0025-7

Giere, R. (1988). Explaining Science: A Cognitive Approach. Chicago, IL: Chicago University Press. https://doi.org/10.7208/chicago/9780226292038.001.0001

Glick, D. (2016). The Ontology of Quantum Field Theory: Structural Realism Vindicated? *Studies in History and Philosophy of Science, 59*, 78–89. https://doi.org/10.1016/j.shpsa.2016.06.007

Godfrey-Smith, P. (2006). The Strategy of Model-Based Science. *Biology and Philosophy, 21*, 725–740. https://doi.org/10.1007/s10539-006-9054-6

Greenberg, G. (2013). Beyond Resemblance. *Philosophical Review, 122*, 215–287. https://doi.org/10.1215/00318108-1963716

Hartley, R. & Zisserman, A. (2004). *Multiple View Geometry in Computer Vision*. Cambridge, UK: Cambridge UP. https://doi.org/10.1017/CBO9780511811685

Hesse, M. (1961). *Forces and Fields*. London, UK: Nelson.

Hubel, D. H. (1996). In L. R. Squire (Ed.), *The History of Neuroscience in Autobiography* (Vol. 1, pp. 294–317). Washington, DC: Society for Neuroscience.

Hubel, D. H., & Wiesel, T. N. (1959). Receptive Fields of Single Neurones in the Cat's Striate Cortex. *Journal of Physiology, 148,* 574–591. https://doi.org/10.1113/jphysiol.1959.sp006308

Ladyman, J. (1998). What is Structural Realism? *Studies in History and Philosophy of Science, 29*, 409–424. https://doi.org/10.1016/S0039-3681(98)80129-5

Leitgeb, H. (2011). New Life for Carnap's Aufbau. *Synthese, 180*, 265–299. https://doi.org/10.1007/s11229-009-9605-x

Levins, R. (1966). The Strategy of Model Building in Population Biology. *American Scientist, 54*, 421–431.

Lewis, C. S. (1936). *The Allegory of Love*. Oxford, UK: Clardendon Press.

Lewis, D. (1984). Putnam's Paradox. *Australasian Journal of Philosophy, 62*, 221–236. https://doi.org/10.1080/00048408412340013

Libkin, L. (2004). *Elements of Finite Model Theory*. Berlin, Germany: Springer. https://doi.org/10.1007/978-3-662-07003-1

Lloyd, E. (2010). Confirmation and Robustness of Climate Models. *Philosophy of Science, 77*, 971–984. https://doi.org/10.1086/657427

McKenzie, K. (2017). Ontic Structural Realism. *Philosophy Compass, 12*, e12399. https://doi.org/10.1111/phc3.12399

Merrill, G. H. (1980). The Model-Theoretic Argument Against Realism. *Philosophy of Science, 47*, 69–81. https://doi.org/10.1086/288910

Muller, F. A. (2011). Withering Away, Weakly. *Synthese, 180*, 223–233. https://doi.org/10.1007/s11229-009-9609-6

Newman, M. H. A. (1928). M. Russell's 'Causal Theory of Perception.' *Mind, 37*, 137–148. https://doi.org/10.1093/mind/XXXVII.146.137

Parker, W. (2010). Whose Probabilities? Predicting Climate Change with Ensembles of Models. *Philosophy of Science*, *77*, 985–997. https://doi.org/10.1086/656815

Putnam, H. (1980). Models and Reality. *Journal of Symbolic Logic, 45*, 464–482. https://doi.org/10.2307/2273415

Putnam, H. (1988). *Appendix to Representation and Reality* (pp. 121–125). Cambridge, MA: MIT Press.

Rescorla, M. (2009). Predication and Cartographic Representation. *Synthese, 169*, 175–200. https://doi.org/10.1007/s11229-008-9343-5

Riemann, G. F. B. (1996). On the hypotheses which lie at the foundation of geometry. (W. K. Clifford, Trans.). In W. B. Ewald (Ed.), *From Kant to Hilbert: A Source Book in the Foundations of Mathematics* (Vol. 2, pp. 652–661). Oxford, UK: Clarendon Press. (Original work published 1868)

Russell, B. (1927). *The Analysis of Matter*. New York, NY: Harcourt, Brace & Co.

Russell, B. (1968). *The Autobiography of Bertrand Russell: 1914–1944*. Boston, MA: Little, Brown and Co. https://doi.org/10.2307/2218050

Ryckman, T. A. (1991). Conditio Sine Qua Non? Zuordnung in the Early Epistemologies of Cassirer and Schlick. *Synthese, 88*, 57–95. https://doi.org/10.1007/BF00540093

Schwarz, W. (2014). Against Magnetism. *Australasian Journal of Philosophy, 92*, 17–36. https://doi.org/10.1080/00048402.2013.765900

Shepard, R. (1981). Psychophysical Complementarity. In M. Kubovy & J. R. Pomerantz (Eds.), *Perceptual Organization* (pp. 279–341). Hillsdale, NJ: Lawrence Erlbaum Associates. https://doi.org/10.4324/9781315512372-10

Shepard, R. N., & Chipman, S. (1970). Second-Order Isomorphism of Internal Representations: Shapes of States. *Cognitive Psychology, 1*, 1–17. https://doi.org/10.1016/0010-0285(70)90002-2

Shepard, R. N., & Cooper, L. A. (1982). *Mental Images and their Transformations*. Cambridge, MA: MIT Press.

Shepard, R. N., & Metzler, J. (1971). Mental Rotation of Three-Dimensional Objects. *Science, 171*(3972), 701–703. https://doi.org/10.1126/science.171.3972.701

Sider, T. (2009). Ontological Realism. In. D. Chalmers, D. Manley, & R. Wasserman (Eds.), *Metametaphysics* (pp. 384–423). Oxford, UK: Oxford University Press.

Smith, G. E., & Kosslyn, S. M. (1980). An Information-Processing Theory of Mental Imagery: A Case Study in the New Mentalistic Psychology. *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association, 1980, Vol. 2: Symposia and Invited Papers*, 247–266. https://doi.org/10.1086/psaprocbienmeetp.1980.2.192593

Suppe, F. (Ed.). (1977). *The Structure of Scientific Theories* (2nd ed.). Urbana: University of Illinois Press.

Suppes, P. (1960). A Comparison of the Meaning and Uses of Models in Mathematics and the Empirical Sciences, *Synthese, 12*, 287–301. https://doi.org/10.1007/BF00485107

Suppes, P. (2002). Representation and Invariance of Scientific Structures. Stanford, CA: CSLI.

Suppes, P., Perreau-Guimaraes, M., & Wong, D. K. (2009). Partial Orders of Similarity Differences Invariant Between EEG-Recorded Brain and Perceptual Representations of Language. *Neural Computation, 21*, 3228–3269. https://doi.org/10.1162/neco.2009.04-08-764

Toader, I. D. (2015). Objectivity and Understanding: A New Reading of Carnap's Aufbau, *Synthese, 192*, 1543–1557. https://doi.org/10.1007/s11229-014-0648-2

van Fraassen, B. (1980). *The Scientific Image*. Oxford, UK: Oxford University Press. https://doi.org/10.1093/0198244274.001.0001

van Helmholtz, H. (1977). On the Origin and Significance of the Axioms of Geometry. In R. S. Cohen & Y. Elkana (Eds.), *Hermann von Helmholtz: Epistemological Writings. Boston Studies in the Philosophy of Science* (Vol. 37, pp. 1–38). Dordrecht, the Netherlands: D. Reidel. (Original work published 1868) https://doi.org/10.1007/978-94-010-1115-0_1

Weisberg, M. (2013). *Simulation and Similarity*. Oxford, UK: Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199933662.001.0001

Wolff, J. (2012). Do Objects Depend on Structures? *British Journal for Philosophy of Science, 63*, 607–625. https://doi.org/10.1093/bjps/axr041

Worrall, J. (1989). Structural Realism: The Best of Both Worlds? *Dialectica, 43*, 99–124. https://doi.org/10.1111/j.1746-8361.1989.tb00933.x

Wurtz, R. H. (2009). Recounting the Impact of Hubel and Wiesel. *Journal of Physiology, 587*, 2817–2823. https://doi.org/10.1113/jphysiol.2009.170209

The editorial and publishing process of this publication has been financed by the research grant "What is Thinking with Images?", SONATA 10, granted by the National Science Centre, Poland, on the basis of the decision No. 2015/19/D/HS1/02426.