

Scientific Representation Is Representation-As

Roman Frigg and James Nguyen

in Hsiang-Ke Chao and Julian Reiss (eds.): *Philosophy of Science in Practice: Nancy Cartwright and the Nature of Scientific Reasoning*, Synthese Library, Volume 379, Springer: Berlin and New York, 2017, 149-179.

1 Introduction

Nancy Cartwright argues that models rather than theories are the units of science that represent parts or aspects of the world: ‘theories in physics do not generally represent what happens in the world; only models represent in this way, and the models that do so are not already part of any theory’ (Cartwright 1999a, p. 180). Her work has been a driving force behind the now broadly accepted view that models are the primary representational units of science.

This invites two questions. Firstly, what is the relationship between theories and models? Many interesting discussions can be had about this question, and important parts of Cartwright’s work address this issue. Our concern, however, is the second question: in virtue of what do models represent selected parts or aspects of the world (in this context usually referred to as their target systems)? Cartwright admits that she has little to say about the relationship between models and their respective targets, beyond cautioning against thinking of representations in terms of structural isomorphisms and appealing instead to a ‘loose notion of resemblance’ (1999a, p. 192–93, cf. b, p. 261–62). She fully accepts that ‘this is just to point to the problem, or label it, rather than say anything in solution to it’ (*ibid.*). But, as Morrison (2008, p. 70) notes

R. Frigg (✉)

Department of Philosophy, Logic and Scientific Method and Centre for Philosophy of Natural and Social Science (CPNSS), London School of Economics and Political Science, London, UK
e-mail: r.p.frigg@lse.ac.uk

J. Nguyen

History and Philosophy of Science Graduate Program, University of Notre Dame, Notre Dame, IN, USA
e-mail: j.nguyen1@lse.ac.uk

when commenting on Cartwright's theory of models, this is the crux of the problem of representation. So proving an account of in virtue of what models represent fills an important lacuna in Cartwright's account, and this the aim of this paper.¹

A promising account of how to think about the representational relationship between models and the world emerges from the work of Nelson Goodman and Catherine Z. Elgin. Our main contention is that scientific representation ought to be analysed in terms of their notion of *representation-as*. This suggestion has previously been made by Hughes (1997), and more recently by van Fraassen (2008) and Elgin herself (references below). However all of these discussions of scientific models remain by and large programmatic. The aim of this paper is to provide detail to the claim that the representational relationship between models and their targets is one of representation-as. Doing so involves adding specificity to claims and definitions, as well as making a number of (friendly) amendments to the requisite machinery.²

We begin with a discussion of Goodman's and Elgin's views on representation (Sect. 2). We then point out that a number of amendments are needed to transform this view into a theory of how models represent. We offer a statement of such a view, which we call the DEKI account of representation (Sect. 3). Material models and fictional models are the two most important classes of models and we indicate how the DEKI account deals with these models (Sect. 4). We end by briefly pointing out that none of the many criticisms that have been put forward against Goodman's and Elgin's views on representation pose a threat to our line of argument since these are typically specific to pictorial representation, which is not our concern (Sect. 5).

We follow common usage and take 'scientific representation' to refer to the representation relation between models and selected parts or aspects of the world. Understood in a broader sense 'scientific representation' would also refer to other kinds of representations such as scientific graphs, images, and diagrams, but they are not in the scope of our current analysis. Throughout, those who feel discomfort about this use of the term could substitute 'model-representation' for 'scientific representation'.

2 Goodman and Elgin on Representation

In a string of publications both Nelson Goodman and Catherine Z. Elgin have developed an account of representation at the heart of which lies the posit that mimetic accounts are fundamentally at odds with our representational practices both in the

¹We do not suggest that Cartwright herself would agree with our account. But it is worth noting that she does comment approvingly on Hughes' account, which is kindred in spirit to ours.

²An implicit assumption of the current project is that no satisfactory account of representation is currently available, and that therefore an effort to formulate one is not just an idle pastime. For want of space we cannot argue for this premise here and refer the reader to (Frigg and Nguyen forthcoming).

arts and in the sciences: representation does not amount to producing a mirror image of the real thing. When referring to views shared by both authors, we use the acronym ‘GE’ to refer to them jointly.³

In discussing their account we use X to stand for the object that does the representing (the picture, the model, the graph, ...), Y to stand for the target of the representation (the Duke of Wellington, the solar system, the patient’s body temperature at different times, ...), and, where appropriate, Z to denote the genre of a representation (what is meant by the ‘genre’ of a representation is discussed below).

2.1 Reference

At the most basic level, what characterises a representation is ‘aboutness’: a representation of Y is about Y . GE identify *reference* as the rudiment of representation. For X to represent Y it has to refer to Y .⁴ There are two basic modes of reference: denotation and exemplification (CJ, 171). Denotation is the relation between a name and its bearer. Exemplification is reference to a property by a symbol that instantiates that property. Denotation and exemplification are not mutually exclusive. A symbol does not have to be either purely denotational or purely exemplificational. Indeed, some symbols combine denotational and exemplificational functions to procedure different kinds of ‘complex reference’. A particularly important kind of complex reference is *representation-as*, which involves a combination of denotation and exemplification (WRR, 141–2). This kind of reference is crucial in the current context because, as we will see, the claim is that scientific representation is an instance of representation-as.

This sets the agenda. We begin by introducing denotation and exemplification in isolation and then proceed to showing how they can be combined to form representation-as.

2.2 Denotation

Denotation is the two-place relation between a symbol and the object to which it applies. It is the crucial concept for GE because they see denotation as the core of representation in both art and science:

Pictures, equations, graphs, charts, and maps represent their subjects by denoting them. They are representations of the things that they denote. [...] It is in this sense that scientific models represent their target systems: they denote them. (TI, 2)

³Throughout the paper we use the following abbreviations: LA for Goodman (1976), MM for Goodman (1984), WRR for Elgin (1983), CJ for Elgin (1996), TI for Elgin (2010), TE for Elgin (2004), and EIS for Elgin (2009).

⁴Notice that GE do *not* use ‘reference’ as a synonym for ‘denotation’.

So for X to be a representation of Y it is necessary (and sufficient) that X denotes Y because ‘denotation is the core of representation’ (LA, 5). For this reason denotation is ‘representation-of’ (TI, 4).⁵

A number of qualifications need to be added about this use of ‘denotation’. First, denotation is often restricted to language, where a name is understood as denoting its bearer. This restriction is neither essential nor helpful. Signs other than words of a certain language can denote. A portrait can denote its subject, a photograph can denote its motif, and a scientific model can denote its target system. There is nothing intrinsic in the notion of denotation that would restrict it to language (WRR, 19–35; TI, 2).

Second, even within language denotation is often restricted to proper names, expressions denoting a singular object. ‘Big Ben’, for instance, denotes the great bell in tower of the House of Parliament. As such denotation is distinguished from predication, which deals with general terms. This restriction is unnecessary:

A predicate denotes severally the objects in its extension. It does not denote the class that is its extension, but rather each of the members of that class. (WRR, 19; *cf.* LA, 19)

The predicate ‘red’ denotes all red things and a picture of the hydrogen atom denotes all hydrogen atoms.

Thirdly, notice that there can be a number of denotational relationships between a picture and its subject:

What a picture is said to represent may be denoted by the picture as a whole or by a part of it... Consider an ordinary portrait of the Duke and Duchess of Wellington. The picture (as a whole) denotes the couple, and (in part) denotes the Duke. (LA, 28)

Presumably a part of the picture also denotes the Duchess, another part denotes the Duke’s nose, yet another part denotes the Duchess’s dress, and so on. In fact, there may, in principle, be an indefinitely large number of denotational relationships that hold between parts of the picture and parts of the situation it denotes.

The observation generalises. Whilst a picture may denote, as a whole, what it is a picture of, parts of the picture may also denote parts of its subject. Whilst a scientific model, as a whole, may denote a target system, parts of the model may also denote parts of the target, and so on for other kinds of representations. This is not to say that there *must* be part-part denotational relationships to establish the primary one that holds between the picture, or model, and the situation it denotes. Examples from modern art provide plausible instances where there is only one such relation. We can imagine a uniformly red canvas captioned ‘Kierkegaard’s Mood’ which as a whole denotes Kierkegaard’s mood (Danto 1981). It’s hard to consider what it would take for a part of the canvas to denote a part of the philosopher’s mood. So, whether or not there are such part-part relationships, and how many of them there are, can only be established on a case-to-case basis.

Viewing denotation as the core of representation may seem innocuous, but it has important consequences and leads to the introduction of a number of crucial

⁵ We put systematicity above grammatical correctness when we write ‘ X is a representation-of Y ’.

concepts. We discuss and illustrate these with the example of pictorial representation. Nothing depends on this choice; the same points could be made using other kinds of representations. We choose pictures because of their intuitive force and because comparing pictures and scientific models will turn out to be instructive in what follows.

The first consequence of the view that denotation is the core of representation is that not all pictures represent. Pictures of Pickwick or unicorns do not denote anything simply because Pickwick does not exist and nor do unicorns. Such pictures therefore do not represent anything (LA, 21). This observation generalises: whenever a picture portrays something that does not exist then the picture does not represent.

This is counterintuitive and one is tempted to object: if we recognise a picture as portraying a unicorn, then surely it represents something, namely a unicorn. GE get around this objection by drawing a distinction between ‘representing’ and ‘being a representation’. A picture represents if, and only if, it is a representation-of (i.e. if it denotes). However, a picture can be a representation without being a representation-of:

A picture that portrays a griffin, a map that maps the route to Mordor, a chart that records the heights of Hobbits, and a graph that plots the proportion of caloric in different substances are all representations, although they do not represent anything. *To be a representation, a symbol need not itself denote, but it needs to be the sort of symbol that denotes.* Griffin pictures are representations then because they are animal pictures, and some animal pictures denote animals. Middle Earth maps are representations because they are maps and some maps denote real locations. Hobbit height charts are representations because they are charts and some charts denote magnitudes of actual entities. Caloric proportion graphs are representations because they are graphs and some graphs denote relations among real substances. *So whether a symbol is a representation is a question of what kind of symbol it is.* (TI, 2–3, emphasis added; cf. LA, 21)

Whether a picture X is a representation-of depends on whether X denotes something. Whether X is a representation depends on whether it belongs to a class of objects that usually denote. In other words, to be a representation something need not denote; but it needs to be an object of the right kind. To facilitate our discussion in the next section we call this the *k-definition of being a representation* (where ‘ k ’ stands for ‘kind’).

But how can a picture be a representation without being a representation-of something? GE point out that we are misled by ordinary language into believing that something is a representation only if there is something in the world that it represents:

What tends to mislead us is that such locutions as “picture of” and “represents” have the appearance of mannerly two-place predicates and can sometimes be so interpreted. But “picture of Pickwick” and “represents a unicorn” are better considered unbreakable one-place predicates, or class terms, like “desk” and “table”. [...] From the fact that P is a picture of or represents a unicorn we cannot infer that there is something that P is a picture of or represents. [...] *Saying that a picture represents a soandso is thus highly ambiguous between saying that the picture denotes and saying what kind of picture it is.* Some confusion can be avoided if in the latter case we speak rather of a ‘Pickwick-representing-picture’

of a 'unicorn-representing-picture' [...] or, for short, of a 'Pickwick-picture' or 'unicorn-picture' [...] *Obviously a picture cannot, barring equivocation, both represent Pickwick and represent nothing. But a picture maybe of a certain kind – be a Pickwick-picture [...] – without representing anything.* (LA, 21–2, emphasis added; cf. TI, 3)

This leads to the introduction of the notion of a Z-representation: *X* is a Z-representation if it portrays a Z. The crucial point is that this does not presuppose that *X* be a representation-of a Z; indeed *X* can be a Z-representation without representing anything. A picture must denote a man to be a representation-of a man. But it need not denote anything to be a man-representation (LA, 25). There is no presupposition that the *k* and the Z be identical. A picture can be a griffin-picture (hence Z = griffin) while it qualifies as the kind of symbol that typically denotes because it belongs to the family of animal representations (hence *k* = animal representation).

How does the classification of pictures into different Z-representations work in cases in which there are no Zs? If there are no griffins, what is the basis for sorting pictures into ones that are griffin-representations and ones that are not? GE respond to this question by introducing the notion of a genre:

Such an objection supposes that the only basis for classifying representations is by appeal to an antecedent classification of their referents. This is just false. We readily classify pictures as landscapes without any acquaintance with the real estate – if any – that they represent. I suggest that each class of [Z]-representations constitutes a small genre, a genre composed of all and only representations with a common ostensible subject matter [...] And we learn to classify representations as belonging to such genres as we study those representations and the fields of inquiry that devise and deploy them. (TI, 3)

These genres are habitual ways of classifying and as such they are neither sharp nor historically stable, and they typically resist exact codification (LA, 23). This, however, does not detract from their importance and usefulness in understanding representations.

The next vexing question is where denotation comes from: what makes it the case that a given *X* denotes something rather than nothing, and what determines its denotatum *Y*? There is pervasive intuition that resemblance is the source of denotation. A man-picture represents men, or a particular man, because it resembles, or looks like, men, or the particular man. This is wrong. Denotation is independent of resemblance or similarity (we use the terms interchangeably). This is obvious enough in the case of language where words do not resemble the things they stand for (at least not in any obvious way), and the observation carries over to pictures, which would seem to be a natural fit for the similarity view. Similarity is not sufficient because it has the wrong logical properties: similarity is symmetrical and reflexive while denotation is neither (LA, 5). Whether or not similarity is necessary for denotation is a more subtle matter. As Goodman points out, everything is similar to everything else in some sense (1972). So all denoting *X*s will be similar to their denotatums. Thus, if similarity is broadly construed in this way then it is necessary for denotation, but vacuously so.

However, in our appreciation of art we do distinguish between relevant and irrelevant similarities in a way that allows us to conclude, for instance, that the portrait of the Duke and Duchess of Wellington is similar to the Duke and Duchess in a way that

Picasso's portrait of Dora Maar is not similar to Dora herself. And yet the portrait still denotes her. The same is true of countless modern pieces of art, but the point applies elsewhere as well. Henry VIII agreed to marry Anne of Cleves after having seen only Holbein's portrait of her, which depicted her as an attractive young woman. The real person was so unlike what he saw on the portrait that Henry decided to annul the marriage immediately. Despite the lack of resemblance, Holbein's portrait did denote Anne. And in still life painting denotation and what is depicted come apart entirely. In Dutch still life symbolism, a snail-picture denotes the humility of everyday life, a jug-of-beer-picture denotes pride in the homeland and a butterfly-picture denotes the transformations of the soul. So, depending on how similarity is construed, it is either unnecessary for denotation, or necessary but vacuously so. Either way, similarity is irrelevant to denotation in any important sense.

The lesson we learn from these examples is that denotation can be achieved by an act of volition: 'Representation-of can be achieved by fiat. We simply stipulate: let $[X]$ represent $[Y]$ and $[X]$ thereby becomes a representation of $[Y]$. This is what we do in baptizing an individual or a kind' (TI, 4). In some cases a stipulation may be a simple ostensive definition (pointing to Y and say 'let X denote this'). Whether or not stipulation suffices in general for establishing that X denotes Y is a question that we cannot address in detail here.⁶ But it is worth noting that in many cases these stipulations have to be mediated by more or less elaborate conventions, which are familiar to a certain audience. In other cases denotation is established by simply captioning the painting. Sometimes appeal to causal chains needs to be made. And so on. The sources of denotation are varied and complicated and much can be said about how reference is established in particular cases (see Chap. 3 of MM for a discussion). Yet the recognition of this diversity leaves the main point untouched: denotation is independent from similarity.

The upshot of our discussion is that there is a complete disconnect between the sort of representation X is and what, if anything, X denotes: 'the denotation of a picture no more determines its kind than the kind of picture determines the denotation. Not every man-picture represents a man, and conversely not every picture that represents a man is a man-picture' (LA, 26; cf. LA, 31).

2.3 Exemplification

An item exemplifies a property if it at once instantiates the property and refers to it: 'Exemplification is possession plus reference. To have without symbolising is merely to possess, while to symbolise without having is to refer in some other way than by exemplifying.' (LA, 53). An item that exemplifies a property is an *exemplar* (CJ, 171). The paradigmatic example of an exemplar is a sample. The swatches of cloth in a tailor's booklet of fabrics (LA, 53), the chip of paint on a manufacturer's

⁶ See Frigg and Nguyen ([forthcoming](#)) for further discussion about the relationship between stipulation and denotation.

sample card (WRR, 71), and the bottle of shampoo we receive as promotional gift (*ibid.*) both refer to relevant properties – a pattern, a colour, and a particular hair treatment – and instantiate them.⁷

The formula ‘exemplification is possession plus reference’ stands in need of qualification. The point to emphasise is that the ‘plus’ ought not to be read literally. Recall from Sect. 2.1 that denotation and exemplification are basic modes of reference. Reference is thus seen as a determinable for which denotation and exemplification serve as determinants. If so, then exemplification cannot literally be reference with something else added to it. Rather, exemplification is the kind of reference that employs instantiation to achieve reference. This can be encapsulated by the altering the formula as follows: An item exemplifies property *P* if it instantiates *P* and *thereby* refers to *P*. This formulation makes it clear that exemplification (like denotation) is a mode, or kind, of reference: the exemplification of a property *P* by an object *X* just is a way for *X* to refer to *P* that involves instantiation.

Exemplification *requires* instantiation: an item can exemplify a property only if it instantiates it (CJ, 172). Therefore, unlike denotation, exemplification cannot be brought about by mere stipulation (TI, 6). Only something that is red can exemplify redness. But the converse does not hold: not every property that is instantiated is also exemplified. Exemplification is selective (TI, 6). An exemplar typically instantiates a host of properties but it exemplifies only a few of them. Consider the example of a chip of paint:

This particular chip is blue, one-half inch long, one-quarter inch wide, and rectangular in shape. It is the third chip on the left on the top row of a card manufactured in Baltimore on a Tuesday. The chip then instantiates each of these predicates in the previous two sentences, and many others as well. But it clearly isn’t a sample of all of them. Under the standard interpretation, it is a sample of “blue”, but not of such predicates as “rectangular” and “made in Baltimore”. (WRR, 71)⁸

Which properties are exemplified and which properties are merely instantiated is not dictated by the object itself: ‘nothing in the nature of things makes some features inherently more worthy of selection than others’ (CJ, 172). In particular, being conspicuous does not, by itself, turn an instantiated property into an exemplified one. A can of paint spilled on the carpet is a vivid instance of the paint’s viscosity yet it does not exemplify viscosity, or indeed anything else (CJ, 174). Turning an instantiated property into an exemplified one requires interpretation (CJ, 175). An interpretation is carried out ‘against the constellation of background assumptions. An interpreter ignorant of those assumptions may be incapable of interpreting or even recognising the symbols’, and ‘[w]ith a change in background assumptions a symbol can come to exemplify new features’ (CJ, 176). The specific details of how this works varies from case to case. Different plays will be exercised and different

⁷Throughout this paper we impose no restriction on what qualifies as a property. An item can exemplify one-place properties, multi-place properties (i.e. relations), and higher order, structural, properties.

⁸For further examples of selectiveness see LA, 53–4, WRR, 72–3, and TI, 5.

interpretational schemes used to render properties salient in Dutch still life painting and in electro-dynamical modelling. But for the purpose of a general theory nothing depends on knowing these details, and one can leave the further elucidation of the details to a study of disciplinary conventions and practices without detriment.

Just as parts of a picture can denote parts of its subject whilst the picture denotes the subject as a whole, different parts of an exemplar can exemplify different properties, all of which may be distinct from those exemplified by the exemplar as a whole. For example, the part of the portrait of the Duke and Duchess of Wellington that denotes the Duke may exemplify ferocity and candour; whilst the part of the portrait that denotes the Duchess may exemplify astuteness and wisdom.

A crucial feature of exemplars is that they provide epistemic access to the properties they exemplify: from an exemplar we can learn about its exemplified properties (WRR, 93). This is because they instantiate the properties they exemplify in a way that makes them salient. The paint chip makes a particular shade of blue salient and thereby acquaints those using the chip with that shade of blue. An exemplar is therefore not merely an instance of a property but a telling instance (CJ, 173; TI, 5): it presents the exemplified properties in a context that is designed to render them salient and make them known to those engaging with the symbol. This is indeed a necessary condition for an item to exemplify a property: if it does not present the property in a way that makes it epistemically accessible, then it cannot exemplify it even if it does instantiate it. The beam of a flashlight instantiates the speed of light but it does not exemplify it because it affords no epistemic access to it (CJ, 174).

Exemplars do not belong to a special category of objects. Anything can in principle become an exemplar simply if it serves as an example (TI, 6). Even items that are not usually used as symbols can be turned into exemplars simply by being used as an example. The front door of a building turns into an exemplar if someone uses it to explain to the workers that all the other doors in the building have to be painted in the same shade of red.

Summing up, we can give the following definition of exemplification: *X* exemplifies *P* if and only if *X* instantiates *P* and *thereby* refers to *P*, and it does so in a way that both makes *P* salient and provides epistemic access to *P*. Bear mind that exemplification is by definition a mode of reference – so this condition in effect singles out a specific way in which an *X* can refer to a property *P*.

There is a final point to clear up before we turn to using exemplification to define representation-as. So far we have used a realistic idiom to talk about properties and their instantiation. This is an expedient that carries no metaphysical commitments. One could provide a nominalist translation for all property-talk, and the notions of exemplification that drop out can be used in the manner discussed below regardless of the metaphysical position adopted. In fact GE prefer a nominalist view of properties (see LA, 4–55, for a nominalist formulation of exemplification). For the purpose of this paper nothing hangs on what stance one takes on the question of the metaphysics of properties and we remain neutral on the matter.

2.4 Representation-As

Many representations represent a thing as something else. Caricatures are a paradigm example: Churchill is represented as bulldog; Thatcher is represented as a boxer; the Olympic Stadium is portrayed as a UFO; etc. These are cases of *representation-as*.

Representation-as can be analysed in terms of representation-of and exemplification (LA, 27–31; TI, 3–10). As we have seen in Sect. 2.2, X is a representation-of Y if X denotes Y . And whether or not this is the case is entirely independent of whether or not X is a Y -picture or not. This makes room for X to be a Z -representation, and denote Y (even where $Y \neq Z$). For instance, X can be a bulldog-picture and denote Churchill. Thus, one might be tempted to define representation-as in the following manner: X represents Y as Z if and only if X denotes Y and X is a Z -representation.

But having a bulldog-picture denote Churchill is not sufficient to represent Churchill *as* a bulldog. Representing Y as a Z involves more than having a Z -representation denote Y :

Evidently, it takes more than being represented by a tree-picture to be represented as a tree. Some philosophy departments can be represented as trees. But to bring about such representation-as is not to arbitrarily stipulate that a tree picture shall denote the department. (TI, 4)

What is lacking in an arbitrary stipulation is a relevant connection between Z and Y . There is a temptation to invoke similarity to bridge the gap between Z and Y and say that X represents Y as Z if, and only if, X is Z -picture and Z is similar to Y . GE deny that this is a solution: everything is similar to everything else in some respect and therefore the requirement that Z be similar to Y is always trivially true and every Z -picture represents every Y as Z . Every case of representation is then *ipso facto* also a case of representation-as (TI, 4).⁹ This renders the notion representation-as useless.

Furthermore, and indeed more importantly, what matters when representing Y as Z is not that Y actually is similar to Z ; what matters is that certain features of Z are imputed to Y . If we represent Thatcher as a boxer we impute certain properties of a boxer such as strength, relentlessness, and mercilessness to Thatcher. Whether or not she actually is similar to a boxer in these regards is immaterial to X 's status as a representation-as; if X represents her as a boxer, X represents her as having these properties irrespective of whether she actually does.

A representation-as not only represents Y as having certain features of Z ; it does so by affording epistemic access to these features. The caricature not only shows Thatcher as a boxer; it also shows us the properties that are attributed to her. That is the crucial difference between a caricature showing Thatcher with boxing gloves and the sentence 'Thatcher is like a boxer'. This observation points the way for the final analysis of representation-as in terms of exemplification:

⁹For a discussion of similarity see also Goodman (1972).

I said earlier that when $[X]$ represents $[Y]$ as $[Z]$, $[X]$ is a $[Z]$ -representation that as such denotes $[Y]$. We are now in a position to cash out the “as such”. It is because $[X]$ is a $[Z]$ -representation that $[X]$ denotes $[Y]$ as it does. $[X]$ does not merely denote $[Y]$ and happen to be a $[Z]$ -representation. Rather in being a $[Z]$ -representation, $[X]$ exemplifies certain properties and imputes those properties or related ones to $[Y]$. [...] The properties exemplified in the $[Z]$ -representation thus serve as a bridge that connects $[X]$ to $[Y]$. This enables $[X]$ to provide an orientation to its target that affords epistemic access to the properties in question. (TI, 10)

This gives a name to the step that was missing in the above example of the philosophy department: imputation. The tree-picture is not a representation-as of the philosophy department because no properties of a tree are imputed to the philosophy department. Like representation-of, imputation can be analysed in terms of stipulation. Although the tree-picture exemplifies certain properties, when Elgin uses it to represent the philosophy department, she does not further impute these properties onto the department. Although this further act of stipulation may appear to make the notion of representation-as relatively easy to come by, it pays to bear in mind that for an agent to impute properties, P_1, \dots, P_n , of X onto Y it does not suffice to simply stipulate that Y has P_1, \dots, P_n , X must also exemplify them. And as we have discussed previously, this is not a trivial matter of stipulation.

We will elaborate further on this definition in the next section. For now we want to add a qualification concerning the role of Z . We speak of X representing Y as Z , which might suggest that the Z is the crucial ingredient. This is not quite the case. What bears the *semantic* weight in a representation-as are the properties exemplified by X *itself*. In a more precise idiom we would say that X represents Y as having certain properties P_1, \dots, P_n , and these properties are instantiated by X . But this does not render Z otiose. As we have seen above, in order to establish that X exemplifies P_1, \dots, P_n , one has to turn to features outside of X itself and this is where Z is crucial. When Z is a boxer-picture, the properties that are exemplified are those that we typically associate with boxers.¹⁰ And it is these properties that are imputed to Thatcher when the picture represents her as a boxer.

If, however, the *same* drawing is interpreted not as a boxer-picture but as a Peter-Buckley-picture, then it would make properties like being a looser salient, and these would be imputed on Thatcher.¹¹ So if we say ‘ X exemplifies P_1, \dots, P_n ’ then Z is eliminable. If we want to explain *why* M exemplifies P_1, \dots, P_n rather than P_m, \dots, P_k , then we can appeal to the fact that M is a Z -representation, and P_1, \dots, P_n are the properties we typically associate with Z . In this way, talk of Z -representations is an ellipse for conveying effectively which properties of X are salient and are therefore exemplified by X .

¹⁰This is an instance of metaphorical exemplification. We return to this issue in Sect. 3.3.

¹¹Buckley is the world’s worst boxer in that he has lost more fights than any other boxer.

3 Entering the Arena of Science

3.1 *Scientific Representation Is Representation-As*

Representation-as is not only the modus operandi of many pictures; it is also claimed to account for how scientific models represent their target systems. Elgin explicitly refers to scientific models repeatedly throughout TI and EIS; scientific examples are mentioned alongside other representations in CJ; and Hughes (1997) and van Fraassen (2008) claim that representation-as is central to the way in which models function in science.

We agree with these authors that representation-as is the basic relation between models and their targets: models are symbols that refer to their targets, exemplify certain features, and represent their targets as exhibiting those features. But the discussions we find in the above-mentioned sources are signposts indicating the way to the inn rather than the inn itself. They offer suggestive remarks, but they do not provide a nuts-and-bolts account of how models represent their target systems. The aim of this section is to provide such an account. We begin by tightening up the definition of representation-as to better fit the way that scientific models function representationally. We argue that certain aspects of the above definition of representation-as do not sit well with scientific models. We then provide a reformulation of the definition that eliminates mismatches and tensions, at least with respect to scientific representation.

3.2 ‘Or related ones’

Paradigm examples of representation-as are ones where a property exemplified by X is identical to the property imputed to Y . We ponder a certain X because the property of interest in Y is no different from the one exemplified by X (TI, 8). However, in some instances of representation-as this is not the case. The crucial clause in the above definition of representation-as presents X as exemplifying certain properties and imputing ‘those properties *or related ones*’ to Y (TI, 10, emphasis added). In fact Elgin emphasises the importance of ‘related ones’:

“Or related ones” is crucial. A caricature that exaggerates the size of its subject’s nose, need not impute an enormous nose to its subject. By exemplifying the size of the nose, it focuses attention, thereby orienting its audience to the way the subject’s nose dominates his face or the way his nosiness dominates his character. (TI, 10)

How are we to understand this qualification? The observation that the properties exemplified by X and the properties imputed to Y need not be identical is exactly right. In fact, few, if any, models in science portray their targets as exhibiting exactly the same features as the model itself. The problem with invoking ‘related’ properties is not its correctness, but its lack of specificity. Any property can be related to any

other property in some way or another, and as long as nothing is said about what this way is it remains unclear what properties X ascribes to Y .

In the context of science, the relation between the properties exemplified and the ones ascribed to the system is sometimes described as one of simplification (CJ, 184), idealisation (CJ, 184), or approximation (TI, 11). This could suggest that ‘related ones’ means ‘idealised’, at least in the context of science (we are not attributing this claim to Elgin; we are merely considering the option). But shifting from ‘related’ to ‘idealised’ (or any of its cognates) makes things worse rather than better. For one, ‘idealisation’ can mean very different things in different contexts and hence describing the relation between two properties as ‘idealisation’ adds little specificity (see, for instance, Jones (2005) and Weisberg (2007) for discussions of different kinds of idealisation). For another, while some representations are idealisations of their targets, others are not. A map of the world exemplifies a distance of 29 cm between the two points labelled ‘Paris’ and ‘New York’; the distance between the two cities is 5,800 km; but 29 cm is not an idealisation of 5,800 km. A scale model of a ship being towed through water is not an idealization of an actual ship, at least not in any obvious way. Or in standard representations of Mandelbrot sets the colour of a point indicates the speed of divergence of an iterative function for certain parameter value associated with that point, but colour is not an idealisation of divergence speed.

One could put faith into context and argue that no further specifications are needed at a general level, and that context determines what properties are imputed onto the target. Just as the context in which a caricature is presented makes it clear that we oughtn’t impute a large nose to its subject but see the caricature as drawing attention to the subject’s nosy character, the context of a scientific model makes it clear what properties it imputes to its target. While it may well be true that context determines the interpretation of a model, it is important to make explicit in every case of modelling what that interpretation is. Indeed, to understand a model, it is crucial to know exactly what properties it imputes on its target. We therefore prefer to write an explicit specification of the relation between the two sets of properties into the definition of representation-as. Let P_1, \dots, P_n be the properties exemplified by X , and let Q_1, \dots, Q_m be the ‘related’ properties that X imputes on Y .¹² Then the representation X must come with a key K that specifies how exactly P_1, \dots, P_n are converted into Q_1, \dots, Q_m ; in fact, K has to provide such a specification for all properties P_1, \dots, P_n that X exemplifies under a certain interpretation in a certain context. Borrowing notation from algebra (somewhat loosely) we can write $K(P_1, \dots, P_n) = Q_1, \dots, Q_m$. K can, but need not, be the identity function; any rule that associates a unique set Q_1, \dots, Q_m with P_1, \dots, P_n is admissible. The relevant clause in the definition of representation-as then becomes: X exemplifies P_1, \dots, P_n and imputes properties Q_1, \dots, Q_m to Y where the two sets of properties are connected to each other by a key K .

The idea of a key comes from maps, which serve as a paradigm to understanding scientific representation (Frigg 2010). The above examples illustrate what we have

¹²In this n and m are positive natural numbers; it need not be the case that $n = m$.

in mind. Let us begin with the map itself. P is a measured distance on the map between the point labelled ‘New York’ and the point labelled ‘Paris’; Q is the distance between New York and Paris in the world; and K is the scale of the map (in the above case, 1:20,000,000). So the key allows us to translate a property of the map (the 29 cm distance) into a property of the world (that New York and Paris are 5,800 km apart). The key involved in the scale model of the ship is more complicated. The P in this instance is the resistance the model ship faces when moved through the water in a tank. But this doesn’t translate into the resistance faced by the actual ship in the same way in which distances in a map translate into distances in reality. In fact, the relation between the resistance of the model and the resistance of the real ship stand in a complicated non-linear relationship because smaller models encounter disproportionate effects due to the viscosity of the fluid. The exact form of the key is often highly non-trivial and emerges as the result of a thoroughgoing study of the situation.¹³ In the representation of the Mandelbrot set in Argyris et al. (1994, p. 660), a key is used that translates colour into divergence speed (see *ibid.*, 695). The square shown is a segment of the complex plane and each point represents a complex number. This number is used as a parameter value for an iterative function. If the function converges for a number c , then the point in the plane representing c is coloured black. If the function diverges, then a shading from yellow over green to blue is used to indicate the speed of divergence, where yellow is slow, green is in the middle, and blue is fast. None of these keys are obvious or trivial. Determining how to move from properties exemplified by models to properties of their target systems can be a significant task, and should not go unrecognized in an account of scientific representation.

In general K is a blank to be filled. What key a given representation is based on depends on a myriad of factors: the scientific discipline, the context, the aims and purposes for which X is used, the theoretical backdrop against which X operates, etc. Building K into the definition of representation-as does not prejudice the nature of K , much less single out a particular key as the correct one. The requirement merely is that there must be *some* key for X to qualify as a representation-as. The above examples also show that introducing keys does not amount to smuggling in a mimetic conception of representation via the back door. On the contrary, keys can be as conventional as we like (correlating, for instance, colour and divergence speed).

Representations can be right or wrong. As we have seen above, that X portrays Y as having properties Q_1, \dots, Q_m does not prejudice the question whether Y really has those properties. There is no guarantee that Y conforms to what is imputed (TI, 10): a target may or may not have the properties that the representation ascribes to it. This does not call into question the status of X as representation. Truth is no requirement for something to be a representation-as.

¹³ See Sterrett (2006) for an illuminating discussion of this example.

3.3 *The Base of a Representation*

The definition of representation-as in Sect. 2 requires X to be a Z -representation. Recall the problem to which the notion of a Z -representation provides an answer. Pictures portray something and so they are obviously representations. Some pictures portray inexistent objects: griffins, elves, unicorns, the edge of the world, and so on. A mimetic theory explains representation in terms of being a mirror image of the real thing. But there cannot be a mirror image of something inexistent. So we face the paradox that there are representations that do not represent. GE resolve this paradox by offering an alternative analysis of representation: a picture showing a Z (a griffin, say) is a representation because it is the sort of object that denotes (k -definition) and it portrays a Z because it belongs to the genre of Z -representations. In this section we carry over this account to the case of scientific modelling. As we will see, this requires extensions and reformulations in a number of places.

On the face of it there seems to be a mismatch between scientific models and pictorial representations. The Schelling model represents social segregation with a checkerboard; billiard balls are used to represent molecules; the Phillips-Newlyn model uses a system of pipes and reservoirs to represent the flow of money through an economy; and the worm *Caenorhabditis elegans* is used as model of other organisms. But neither checkerboards, billiard balls, pipes, or worms seem to belong to classes of objects that typically denote, as would be required by the k -definition. And neither do scientific fictions such as elastically colliding point particles, frictionless planes, utility-maximising agents, and chains of perfectly elastic springs, as well as the mathematical objects used in science. Matrices, tensors, curvilinear geometries, Hilbert spaces, symmetry transformations, and Lebesgue integrals have been studied as purely mathematical objects long before they became important in the sciences, and their representational use is not grounded in their membership in a class of representational objects (there are also objects such as quaternion groups which are similar to the ones just mentioned and therefore belong to the same classes but which are not representations).

Similar worries arise in connection with the notion of a Z -representation. Models are not classified as Z -models because they show or portray a Z . The pipe model of an economy is not an object that portrays pipes (a pipe-representation) that refers to an economy; it *is* a system of pipes. The billiard ball model of gas is not an object that portrays billiard balls (a billiard-ball-representation) that denotes a gas; it is a collection of billiard balls. A checkerboard model of segregation is not representation of a checkerboard that denotes segregation; it just is a checkerboard that is used to represent segregation. And the point can be repeated for the chain model of a polymer, the lattice model of a crystal, the cellular automaton model of a granular medium, the worm model of cell division – or indeed any other model.

So neither the k -definition nor the notion of a Z -representation seem to sit well with how models work. We now argue that mismatch is only apparent, but removing the air of incongruity requires work and will result in an extension of the framework.

Let us begin with the k -definition. In the current context, the problem with this definition is that checkerboards, worms, point particles and matrices aren't members of a class of typically denoting objects in the same obvious way as pictures are. Most worms are just organisms and most checkerboards are just checked structures, and so neither worms nor checkerboards *as such* belong to class of objects that typically denote. The qualification 'as such' is crucial. Some worms are special in that they are chosen by someone involved in a scientific investigation to serve as a scientific model. In contrast with worms plain and simple, worms *so chosen* do belong to a class of objects that denote. But they do belong to this class not for what they are intrinsically, but for the use they are put to. So what turns worms into representations is the fact of being used representationally by someone. This point has been made by many, and we agree.¹⁴ So at least in the context of scientific modelling the k in the k -definition is the (trivial) condition that an item is chosen by someone to serve as a representation, and any object can be so chosen.¹⁵ For this reason anything can, in principle, be used as a representation.

The perplexities surrounding Z-representation are more recalcitrant. The problem, as we have seen, is that models, unlike pictures, do not seem to fall into classes according to what they portray. Constable's *Salisbury Cathedral from the Meadows* is a cathedral-picture that denotes a particular cathedral in the English countryside. But the Phillips-Newlyn model is not a pipe-representation that denotes an economy; it *is* a system of pipes. What has gone wrong? The way out of this jumble is to realise that there is a third layer that has gone unmentioned so far: the substratum of representation. Constable's *Salisbury Cathedral from the Meadows* is not only a cathedral-picture; it is also a canvass mounted on a frame covered with oil and pigments of a certain chemical composition. In fact, saying that it is a cathedral-picture is a shorthand for saying (something like) the following: *Salisbury Cathedral from the Meadows* is a canvass covered with paint, which, under normal visual conditions, is recognised by normal spectators as portraying a cathedral.

The somewhat obvious yet crucial point is that every representation has a material substratum, and that this substratum ought to be recognised in a theory. We call this substratum the *base of the representation*; *base* for short.¹⁶ The base is seen by onlookers as portraying certain motif Z . For reasons that will become clear soon we call the process of seeing, say, a cathedral in a configuration of pigments an *interpretation*. We then submit that the right analysis of cases like the Phillips-Newlyn model is the following. First appearances notwithstanding, the system of pipes is indeed a Z-representation, but Z does not stand for 'pipes' but for 'economy': the

¹⁴ See for example: Teller (2001), Giere (2004, 2010), Callender and Cohen (2006), Suarez (2004), and van Fraassen (2008).

¹⁵ We focus here on how models meet the k -definition, Elgin reaches a similar conclusion regarding how something functions as an exemplar. She emphasises that '[a]ny item can serve as an exemplar simply by being *used* as an example. So items that ordinarily are not symbols can come to function symbolically simply by serving as examples' (TI, 6, emphasis added cf. WRR, 72, 80).

¹⁶ We assume that our use of the term 'base' coincides with Suarez's 'source' (2004, p. 767) and Contessa's 'vehicle' (2007, p. 48) to denote the object or system that is used to represent a target.

machine is an economy-representation just as Constable's canvas is a cathedral-representation. The base of the representation is the pipe system, which is the analogue to the canvass (the base of Constable's painting).

So the mistake we made above was to conflate what a representation shows (Z) with what the painting is as an object. This mistake was engendered by the fact that the focus is on different places in art and in science. Looking at Constable's painting, we could specify the thickness of the layer of paint, we could say what the chemical constitution of the paint is, and so on. There are all kinds of things one can say about the physical entity that constitutes the painting. Often, however, there is not much interest in such considerations (with the exception of conservators or auctioneers who might have to restore a painting or prove its originality). In science, by contrast, the base is often a matter of great concern. Indeed, models are often classified according to what they are rather than to what they represent: we speak of checkerboard models, worm models, pipe models, etc. To capture this idea we introduce a term of art: O -objects. As used here, ' O ' is simply a specification of what kind of thing an object is. A checkerboard is a checkerboard-object; a set of hard balls is a hard-ball-object etc.¹⁷

Classifying something as O -object says nothing about whether, and if so in what way, the object functions symbolically. Anything that exhibits the pattern of a checkerboard is a checkerboard-object. Some of these objects are chessboards, some are marble floors, and some are elaborately baked cookies. An O -object can, but need not, function symbolically. A 'mere' O -object can be turned into the base of that representation by an act of stipulation. Characterising an object as a base is to provide a *functional* characterisation: whether something is the base of a representation depends on whether it is *used* as such. Hence the characterisation of something as a particular O -object is independent from its characterisation as a base (and vice versa). Anything can be used representationally, yet nothing is necessarily or intrinsically a representation.

A base object becomes a Z -representation through an act of interpretation. In the case of pictures, the nature of this interpretation has been the centre of attention for a good while: how one sees a canvass covered with paint as showing a cathedral is regarded by many as one of the most important problems of aesthetics. Schier dubbed it the 'enigma of depiction' (1986, p. 1), and an entire body of literature is concerned with it (see Kulvicki 2006b for a review). Although an interesting issue in its own right, how pictures work is not our concern. The point of interest here is how an O -object is turned into a Z -representation in a scientific context: how do we come to see a system of pipes and reservoirs as being an economy-representation,

¹⁷There is no expectation that O be a natural kind. The choice of a certain base can be seen as defining a genre of modelling if there is a tradition of modelling with that base. There are genres of checkerboard modelling, *Caenorhabditis elegans* modelling, cellular automaton modelling, harmonic oscillator modelling, billiard ball modelling, and so on. The items belonging to these genres are all models using the same base to represent a target system. This class can be large, and target systems can be varied. There are cellular automaton models of sand piles, forest fires, earthquakes, traffic jams, and many other systems. Using cellular automata to represent a variety of systems can be seen as constituting the genre of cellular automaton modelling.

or a collection of point particles as a gas-representation? There is no simple and universal answer to this question. How this is done will depend on disciplinary traditions, research interests, background theory, and much more. In fact, ‘interpretation’ is a blank to be filled, and it will be filled differently in different cases. (However, below we discuss a constraint on scientific interpretations).

In an enterprise as diverse as science, this is what one should expect. Sometimes an interpretation is ‘close’ to the *O*-object in that it interprets the object in its ‘own terms’. Scale models are a case in point: the small car is interpreted as car-representation and the small ship is interpreted as a ship-representation. It is important, though, that this interpretation is in no way forced upon us by the object itself. The small car could be interpreted as a quality-of-paint-representation demonstrating to potential customers the quality of paint that a certain production line is able to produce; or the small ship could be made from potassium and be interpreted as a human-bone-model used to study the effects of the presence of radioactive isotopes in bone material. These examples are not as far-fetched as they may seem. The Phillips-Newlyn machine is a system of pipes and reservoirs, but these aren’t interpreted in their own terms, i.e. as pipes and reservoirs. This system becomes an economy-representation by interpreting the quantity of water as the quantity of money, the flow of water through a pipe as the flow of money through a part of the economy, and the amount of water in a certain reservoir as the amount of money in the central bank. Yet the same object could be a water-supply-system-representation under another interpretation.

We equally encounter this freedom of interpretation when dealing with fictional rather than material models. An imaginary string consisting of beads connected by springs is interpreted as polymer-representation, or a collection of harmonic oscillators is interpreted as the wall of a black body. Sometimes the same base is reinterpreted to become a different *Z*-representation altogether. A fictional system of two spinning spheres in otherwise empty space that attract with each other with a $1/r^2$ force was interpreted by Newton as sun-earth-representation and by Bohr as a hydrogen-atom-representation. And a cellular automaton can variously be interpreted as a sand-pile-representation, a forest-fire-representation, or a traffic-jam-representation.¹⁸

The basic insight that a *Z*-representation may, but need not, represent *Z*’s stands, and it helps explaining the symbolic function of models without target. There are certain population models that study populations with more than two sexes and find out how they evolve (see Weisberg 2013, §7.4, for a discussion of these models). Such a model is an *n*-sex-population-representation (for a some natural number $n > 2$), but it is not a representation-of an *n*-sex-population simply because there are no such populations.¹⁹ The Phillips-Newlyn machine as an economy-representation;

¹⁸For an extended discussion of such models see Frigg (2003).

¹⁹As Weisberg notes, there are some recently discovered cases of social insects with three sexes (2013, p. 133). But this takes nothing away from our point. These hadn’t been discovered when three-sex population models were originally considered (so denotation was established later), and there is no principled reason to rule out larger *n*-sex models, which don’t represent any actual population.

but it need not be a representation-of any particular economy. It was originally constructed as a model of a Keynesian economy, and there was no claim the any real-world economy is of this kind. That the machine was later used as a representation-of the economy of Guatemala was in no way built into its status as an economy-representation and establishing denotation in this way constituted an extra and conceptually independent step.

There is a final bridge to cross. Recall that the core of representation-as is exemplification, and something exemplifies a certain property P if it instantiates P and thereby refers to P . Have we not just undermined this very idea by introducing the notion of an interpretation? The Phillips-Newlyn machine instantiates water levels in a reservoir, but it does not seem to instantiate money levels in a central bank. So the problem seems to be that an interpretation introduces properties that the O -object doesn't possess. Admittedly there is a tension. However, nothing in our account depends on instantiation being literal instantiation. What matters is that properties are epistemically accessible and salient, and this can be achieved with what we call *instantiation-under-an-interpretation-I*; *I-instantiation* for short. An economic interpretation of the Phillips-Newlyn machine interprets amounts of water as amounts of money. It does so by introducing a clearly circumscribed rule of proportionality: x litres of water correspond to y millions of the model-economy's currency. This rule is applied without exception when the machine is interpreted as an economy-representation. So we say that under the economic interpretation I_e the machine I_e -instantiates money properties.

In cases like the scale model under a standard interpretation the properties literally instantiated and *I-instantiated* coincide; in cases like the Phillips-Newlyn machine they are different. But even when they are different, the tightness of the interpretation guarantees that epistemic access and saliency are warranted. Water flows are conspicuous in the Phillips-Newlyn machine, and since the interpretation imposes a rule of proportionality that allows for a straightforward conversion of amounts of water into amounts of money, money is equally conspicuous. It can then be regarded as a necessary requirement for an interpretation to be scientific that it be able to establish such a close relation.

In passing we notice that, upon close inspection, the same problem arises in the case of pictures. Taken literally a painting only possesses physical properties (namely those which it possesses qua physical object): having such and such mass, emitting light of a certain spectrum, etc. But typically these are not the properties we are interested in, and they aren't the properties that are associated with the Z of whatever type of Z -representation the picture is. The boxer picture exemplified properties associated with boxers – toughness, tenacity and so on. But like money in the Phillips-Newlyn case, these properties aren't instantiated in the O -object (the canvass). In such cases Goodman and Elgin employ the notion of metaphorical exemplification, a notion that requires metaphorical instantiation (LA, 50–51): 'A painting that literally exemplifies 'dark' may metaphorically exemplify 'disturbing' (WRR, 81). To signal the typically high degree of regimentation in scientific contexts we prefer to talk about *I-instantiation* rather than metaphorical instantiation, but this may well be a matter of taste and nothing of substance hangs on it.

A representation of Thatcher as a boxer imputes properties to her that the picture instantiates metaphorically. In the same way a model can impute properties to the target that it *I*-instantiates. This is to say that the machinery of representation-as remains unaffected by the introduction of *I*-instantiated properties: these properties can be exemplified and then keyed up, just like properties that are instantiated literally. The Phillips-Newlyn machine I_e -instantiates flows of money in some model-currency, and those working in the central bank of Guatemala who took the machine to be a representation-of the economy of their country would have had a key translating facts about flows of model-currency into claims about the flow of Guatemala Quetzales.

3.4 *Expedients*

Mechanical models of constrained motion based on D'Alembert's principle exhibit virtual displacements; some climate models have flux adjustments; and in electrodynamic models of radiation there are advanced potentials with the effect that the future exerts a causal influence on the present. Yet no one imputes virtual displacements, flux adjustments, or advanced potentials to reality. And what is more: no one imputes even anything like them to reality. While advanced potentials could be (and indeed often are) dismissed as an undesirable side effect of the mathematical machinery, virtual displacements and flux adjustments cannot be dismissed in this way. They are crucial features of these models; they are the drivers of the inferential machinery and without them the desired results don't follow. So these features are instantiated and they are salient. Yet they are computational conveniences, expedients that have no counterparts in reality.

Cases like these show that some models have salient features that we do not impute to reality. An account of representation must make room for such features. There are a number of ways to deal with them in the current framework, and which way one chooses may well depend on the case at hand. The first is to deny that features like virtual displacements are exemplified. Exemplification involves reference over and above possession and saliency, and virtual displacements refer to nothing. The second option is to note that the notion of a key does not require that all exemplified properties be translated. The equation $K(P_1, \dots, P_n) = Q_1, \dots, Q_m$ makes room for m being smaller than n and some of the P s can be left untranslated. These two options seem desirable for features that are deemed pure expedients, which makes any translation into a real-world counterpart futile from the outset.

Not all instrumental parts of models need to be of that kind. Advanced potentials could be real, no matter how fantastic they seem to us now. In this case a key can be offered, but imputation can be suspended. So there may be Q s which are simply not imputed. In these cases the model does not represent the target as having Q ; but Q may be offered for consideration in case the situation changes (which is different from simply dismissing something as a pure expedient without any 'reality value'). There is no point in figuring out what the real world counterpart of a flux adjustment

would be; we know that there is no such thing in the world; but there might be something like an advanced potential even though we deem this implausible given our current state of knowledge. This difference in attitude can be reflected in the difference between non-keying up and suspension of imputation.

3.5 The DEKI Account of Scientific Representation

We are now in a position to tie the loose ends together and provide a definition of scientific representation:

Consider an Agent A . The agent chooses an O -object as the base of representation and turns it into Z -representation by adopting an interpretation I . Let X refer to the package of the O -object together with the interpretation I that turns it into a Z -representation. X represents Y as Z if, and only if, the following conditions are satisfied:

(RA1) X exemplifies features P_1, \dots, P_n .

(RA2) X comes with a key, K , specifying how P_1, \dots, P_n are translated into a (possibly identical) set of features Q_1, \dots, Q_m : $K(P_1, \dots, P_n) = Q_1, \dots, Q_m$.

(RA3) X denotes Y , and parts of X may denote parts of Y .

(RA4) A imputes at least one of the properties Q_1, \dots, Q_m onto Y .

We call this the DEKI account of representation to highlight its key features: denotation, exemplification, keying-up, and imputation. A representation is *faithful* if Y indeed possesses the properties that X ascribes to it via A 's imputation (LA, 36). That this be the case is *not* built into the notion of representation-as. X can represent Y as possessing properties Q_1, \dots, Q_m and Y cannot instantiate a single of them.²⁰ Figure 1 provides a schematic representation of the DEKI account.

Let us now add a number of qualifications and comments. The first qualification is that this is the general form of an account of representation and as such it needs to be concretised in every particular instance of representation. In every concrete case of a model representing a target one has to specify what the base object is, how it is interpreted, what sort of Z -representation it is, what properties the base object exemplifies, how denotation is established, what translation key is used, and how the imputation is taking place. Depending on what kind of representation we are dealing with, these 'blanks' will be filled differently. But far from being a defect, this degree of abstractness is an advantage. Scientific modelling is an umbrella term covering a vast array of different activities in different fields, and a view that sees representations in fields as diverse as elementary particle physics, evolutionary biology, hydrology, and rational choice theory as working in exactly the same way is either mistaken or too coarse to make important features visible. Our definition

²⁰ It ought to be noted also that faithfulness is orthogonal to the realism versus antirealism question (TI, 13). Agent A can interpret the model so that it either only exemplifies observable properties or, that only observable properties are keyed up. Such an interpretation is antirealist.

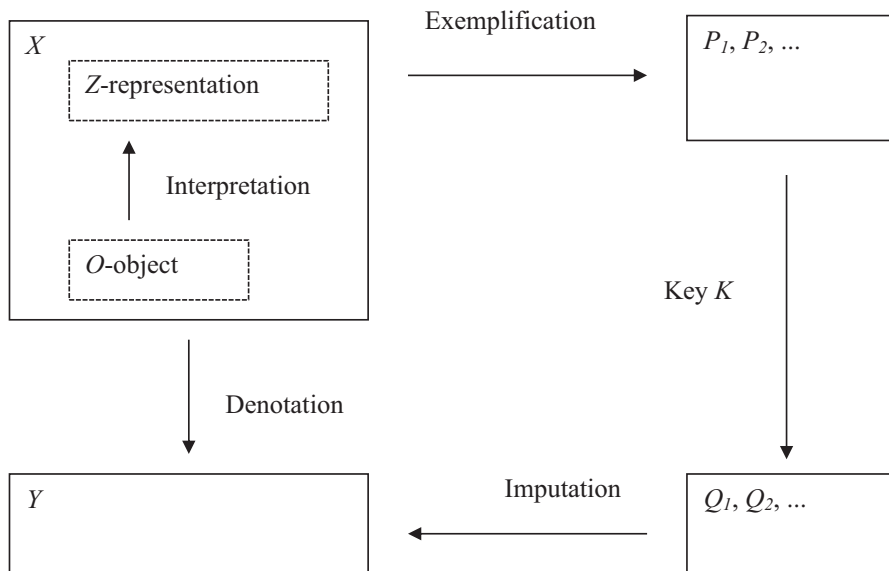


Fig. 1 The DEKI account of representation

occupies the right middle ground: it is general enough to cover a large array of cases and yet it highlights what all instances of scientific representation have in common.

A second qualification is that the ordering of the conditions is not supposed to introduce a temporal element into scientific representation, nor indicate any logical priority. It needn't be the case that a model user first establish denotation, before determining which properties are exemplified by the model and only then translating them and imputing them to the target system. None of the four conditions RA1–4 has to be established prior to the others. A model user could equally well start off with the target system and a set of properties of interest. She could then construct an inverse key associating those properties with ones that we have firmer grasp on in the context of model building.²¹ She could then construct a model that exemplifies those properties, in the appropriate manner against the appropriate interpretational scheme, before taking the model as a representational base and establishing the denotation relation between it and the target. Such a process is not ruled out by our conditions. DEKI does not function as a diachronic account of scientific representation. It is synchronic: as long as the conditions are met, in whatever order, a model represents its target system as Z .

Also notice that our account makes it obvious how scientific models can be used to learn about their target systems. We look at the properties P_1, \dots, P_n exemplified by the model, along with the key, and infer that the target system has the properties

²¹ We are grateful to Alexander Bird for pointing out this possibility.

Q_1, \dots, Q_m that result from the application of the key. It is important though to emphasise that the key often has the character of a hypothesis (Frigg 2010, p. 129). We stipulate that we expect a model to represent in a particular way. This stipulation may be grounded in background knowledge, a theory, or yet something else. There is no guarantee that the target indeed instantiates the imputed properties. Whether or not this is the case does not follow from the model itself and has to be determined by other means. This is the methodological aspect of the above observation that faithfulness is not built into the notion of representation.

Fourthly, note that our previous discussions of the piecemeal nature of denotation and exemplification are relevant here as well. Just as a X can represent Y as having properties Q_1, \dots, Q_m , a part of X can represent a part of Y as having properties Q_1, \dots, Q_m . And this can be explained in an analogous manner to the above. A part X' of X exemplifies a set of certain properties. The key translates those properties into a (possibly identical) set of properties. Assuming that X' denotes a part Y' of Y and the agent imputes those properties onto Y' , then X' represents Y' as having the properties Q_1, \dots, Q_m . Again, whether or not parts of X denote parts of Y , and whether or not the former represent the latter as having properties Q_1, \dots, Q_m is something that has to be determined on a case-by-case basis.

Fifthly, typically the interpretation is chosen so that it interprets the base in terms of the target. If the intended target is an economy, the interpretation would usually construe the base object as an economy. This, however, is a rule of common sense rather than a requirement of semantics. One could construct contrived examples whereby, say, a pipe system is interpreted as an economy and taken to be a representation-of a biological population.

Furthermore, while space constraints prevent us from offering a comprehensive comparison of the DEKI account with other accounts of representation, a few quick remarks along those lines are in order. It is now clear where *purely* stipulative accounts of scientific representation go wrong. We said that a 'mere' object is turned into a representation by being used representationally by someone. This may suggest that *all* it takes to turn an object into a scientific representation is to baptise it as a representation. This view has been advocated by Callender and Cohen, who argue that scientific representation comes down to an act of arbitrary stipulation (2006, p. 73–74). But stipulation is only part of the story. Choosing a particular object as the X for one's representation is an arbitrary act of stipulation. But that object becomes a representation-as only once denotation is established (which may, or may not, involve a further act of stipulation), an interpretation is adopted and properties are exemplified (and since exemplification requires instantiation more than an act of fiat is needed for this to be the case), a key is introduced, and the Q -properties imputed on the target (again, which may involve a further act of stipulation). To do this successfully requires more than pointing to two objects and declaring 'the first shall be a representation of the second'.

Also of interest is the relation between representation-as and the similarity view of representation. Using X to represent Y as Z involves imputing properties of X , or related ones, to Y . Is that not tantamount to claiming that X and Y are similar to one

another, at least in some respects and to some degrees? The grain of truth in this observation is that the base and the target sometimes do share properties. But this need not always be so and the sharing of properties is not sufficient for representation. Properties of the X have to be connected to those imputed to Y with a key K . Sometimes the key can say that the properties in the model are idealisations of the properties of the target (for instance when we model the surface of an ice rink as frictionless). In such a case it is plausible to say that the properties in question are similar. However, keys need not be of this kind; they can correlate properties that are rather dissimilar. Claiming colours are similar to divergence speeds or that distances on a map are similar to distances in the world would stretch the notion of similarity beyond breaking point even on a liberal understanding of similarity. And understanding the differences between the model ship and real ship in their interaction with a viscous liquid was an important step in understanding what forces affect the motion of a body through water.²² So there are cases of representation-as that are not cases of similarity. And even resemblance in the relevant respects and degrees is not sufficient. As Elgin (TI, 11–12) points out, it is not enough for the similarity to be there, the representation must make that similarity manifest. Hidden similarities don't ground representation, but making them salient requires exemplification.

Finally let us turn to the notion of a scientific model. Even though 'scientific model' and 'scientific representation' have often been used almost interchangeably, models and representation have an uneasy relationship. Trivially not all representations are models. And vice versa not all models are representations, as the examples of multi-sex populations and Yang-Mills particles illustrate (see Weisberg (2013, §7.4) and Hartmann (1995) respectively for further discussion of these target-less models). There are also the cases of models such as Maxwell's model of the ether that were, at some point in time, thought to denote actual target systems but then turned out to fail to do so. A full examination of this relationship is a task for another day. The aim of our paper is not to equate models and representations. Indeed we put forward only the more modest conditional claim: if a model represents a target, then it does so in the sense of the DEKI. This, however, leaves open the question of how to define a model, and it involves no claim that all models represent.

4 A Glimpse at Cases

In this section we briefly indicate how the theory we have outlined applies to different types of scientific models. We present sketches of a number of cases; for want of space we leave in-depth case studies for another occasion.

²²One can of course retrospectively call the relevant scaling relation 'similarity', no matter how complicated it turns out to be. This is not wrong; but it is useless. 'Similarity' is a success term, a label attached to a relation *ex post facto*, and it contributes nothing to either discovering or understanding that relation.

4.1 *Material Models*

Material models are models whose base is a material object. The wooden model of car, the model of the new Airbus A380 in the travel agent's window, and the ball-and-stick models used in chemistry classes are standard examples. Architectural models of buildings belong to this group, and so do *Caenorhabditis elegans*, the Phillips-Newlyn machine, the hydraulic model of the San Francisco Bay, and insect robots.²³ Some of them are ready-mades in that they occur in nature and have been selected to function as models because they instantiate properties under scrutiny; others have been carefully crafted to instantiate certain features of interest. This illustrates the point that exemplification requires instantiation. If all that mattered was denotation then one could have one's saltshaker represent the San Francisco Bay. This won't work because typically saltshakers don't instantiate properties that can be keyed up with ones of the Bay. Scientists carefully choose their models because they are useful only if they instantiate the right properties.

This does not imply, however, that all properties that are instantiated matter; nor does it imply that a model imputes its own properties one-to-one on the target. Exemplification is selective. A wooden model of a car and architectural model of house exemplify their shape but not their material constitution. To infer from such models that the car has a wooden body or that the balcony of the house is made from PVC is misunderstanding what properties the model exemplifies. Even if models are given the right interpretation and the exemplified properties are identified correctly, properties are not imputed on the target using a 'same same' rule, nor is the relation always simple scaling. The real car's air resistance will not be ten times the 1:10 model car's air resistance. Understanding the key of the model is crucial. In cases such as architectural models this is easy; in cases such as the car's air resistance it is far from trivial.

In fact spelling out the key of such a model is nothing short of a science in its own right. Sterrett (2002) provides a vivid description of the complex interplay of dimensional analysis, scaling relations, and theoretical tools like the Mach number, the Reynolds number and the Richardson number that are involved in using a material model to learn about the mechanical properties of a target system. And similar stories can be told about almost every material model. The details of these stories don't matter for now; what matters is that they fit the mould of our account of representation: the material model exemplifies properties that are keyed up with other properties in a potentially complicated and nontrivial way. These properties are then imputed to the object that is denoted by the model.

²³ For further discussions of these models see: Ankeny (1998) for *Caenorhabditis elegans*; Morgan (2012, Chap. 5) for the Phillips-Newlyn model; Weisberg (2013) for the San Francisco bay model; and Webb (2001) for robotic models of insect behaviour.

4.2 *Fictional Models*

Fictional models are models whose base is a non-existent object. The stock examples of mechanical modelling fall into this category: a perfect cylinder rolling down an inclined frictionless Euclidean plane, a small perfect sphere orbiting around a large perfect sphere where the only force acting on either is the mutual gravitational attraction, etc. Models involving point particles, three sex populations, perfectly rational agents, and markets in equilibrium also belong to this family of models. It is readily acknowledged that such models are abstract, in the sense that they are not spatio-temporally located. But for our account to apply to these models we have to make sense of how they can exemplify properties. In particular, since exemplification requires instantiation, we have to make sense of how such models can instantiate properties.

For some properties this is relatively straightforward. A three-sex model instantiates the property of being a biological model. A two-body model instantiates the property of being a mechanical model. But these properties do not exhaust the properties exemplified by such models. At least some of the properties exemplified by them seem to be *physical* properties. We use them to learn about how reproduction rates vary with the number of sexes in a population, or the specific physical orbit of a body under the influence of gravitational attraction of a larger mass.

In the context of discussing how the similarity-based account of scientific representation accounts for the purported similarity between the ideal pendulum and an actual pendulum, Martin Thomson-Jones points out that:

[N]o non-spatiotemporal object can have the properties ascribed to the simple pendulum, for no object which has, for example, a length, and behaves in the way the simple pendulum is said to behave in descriptions of it – moving through space over time in a particular way – can be non-spatiotemporal. (2010, p. 291)

The observation generalises to all fictional models (TI, 8–9). If such models cannot instantiate physical properties, then they cannot exemplify them; and if they cannot exemplify them, then they cannot represent a target as possessing them. So no fictional model can represent a planet as having mass or a population as consisting of rabbits simply because no fictional object can have a certain mass or consist of rabbits. This, so one might object, leads to the breakdown of our account because these models clearly represent actual target systems as having such properties.

Let us briefly sketch the main line of argument of two replies. The first appeals to the fiction view of models, which regards models as fictional scenarios of the same kind as, say, Tolkien's middle Earth (Godfrey-Smith 2006). Far from opening the flood-gates to arbitrariness, this view carefully regulates what properties model-objects possesses. It is true in *Lord of the Rings* that hobbits are diminutive human-

oids; it is false that they are giants. In the same way it is true in the Newtonian model of the solar system that planets move in stable elliptical orbits; it is false that they move in square orbits. In other words, according to the relevant fictions hobbits instantiate humanoidity while they do not instantiate giantness, and orbits of model planets instantiate ellipticality while they do not instantiate squareness. Different versions of the fiction approach explain fictional truth and fictional instantiation differently, but no account does without it.

One of us has offered a version of this view in which models are identified with acts of imagination (Frigg 2010). Walton's pretence theory is used to explain the status of imagined objects, and both an account of truth in fiction, as well as an epistemology for models thus understood, is offered. So on this view there is a clear distinction between properties that a fictional object has and ones that it doesn't have, and the model affords epistemic access to the properties of interest. Making a property salient is no different in a fictional model than in a material one; and neither is the establishment of a denotational relationship. So we have found all the ingredients needed for exemplification, and we see no reason why one should not say that fictional scenario exemplifies properties as long as it is understood that instantiation is explained in as an act of pretence. Both models, and parts thereof, *fictionally* instantiate physical properties just in case we are prescribed to imagine that the system, or parts thereof, instantiates those properties. And we can simply generalise the notion of exemplification to allow for fictional instantiation: X exemplifies P if and only if X (fictionally or otherwise) instantiates P and thereby refers to it.

An alternative response to Thomson-Jones' objection is to point out that the account of scientific representation proposed here, does not require that scientific models literally instantiate physical properties. This observation can be utilized in response to the objection in two distinct ways. Firstly, just as the notion of an interpretation allows us to make sense of physical O -objects I -instantiating physical properties that they don't literally possess, it also allows us to make sense of abstract O -objects I -instantiating physical properties that they cannot possess *qua* abstract object. Secondly, note that all that is required by RA1 is a scientific model exemplify a set of *some* properties P_1, \dots, P_n which are translated with the help of key into a set of properties, Q_1, \dots, Q_m , (which need not be identical to P_1, \dots, P_n) which can then be imputed onto the target. There is no presumption that the P_1, \dots, P_n be physical properties. Elgin (TI, 8–9) suggests that abstract models can instantiate magnitudes of physical properties such as mass, and that this is sufficient to do the work we need it to do:

[...] not being a material object, the model does not have mass. So it cannot exemplify the mass of the pendulum. This is true. Strictly, the model does not exemplify mass. Rather it exemplifies an abstract mathematical property, the magnitude of the pendulum's mass. [...] It does no harm to say that they exemplify physical magnitudes. But this is to speak loosely. Strictly speaking, they exemplify mathematical (or other abstract) properties that can be instantiated physically. (TI, 8–9)

Thus, scientific models can represent their targets as having physical properties without literally instantiating them.

These replies to Thomson-Jones are not mutually exclusive. One may be a fictionalist with respect to most scientific models, whilst still admitting that *purely mathematical* models are, ontologically speaking, mathematical structures (either those directly presented, or those structures satisfying certain mathematical equations), and that these structures either *I*-instantiate physical properties, or that the key that links them to their target systems takes mathematical properties to physical ones. The ontological requirements of our account of scientific representation are minimal.

5 Conclusion

We have developed an account of scientific representation that provides detail to the claim that the representational relationship between models and their targets is one of representation-as. In doing so we have relied heavily in GE's account of representation-as in art. One might object that we have bet on the wrong horse. GE's account of representation is not uncontroversial in the realm of pictorial representation. Indeed Douglas Arrell writes:

Goodman's contention that representation is a form of denotation has achieved notoriety rather than acceptance. A survey of some forty of the articles and reviews which appeared in the wake of *Languages of Art* reveals that in about three-quarters of them this theory was a major topic of concern, and that overwhelmingly, the concern was to refute it; indeed, it is hard to find a clear-cut case of someone agreeing with it. (1987, p. 41)

If the account does not work in aesthetics, why should we expect it to work in the context of scientific representation? In answer to this, it is worth briefly outlining why aestheticians have found GE's discussion of pictorial representation so unsatisfactory, and why we think the objections – even if assumed to be sound – have no bite in the current context.

The first point to observe is that majority of objections to GE's discussion of representation focuses on *pictorial* representation rather than the notion of representation-as, or its basis in exemplification. Indeed, Robinson notes that in 'contrast with the reams of pages devoted to Goodman on representation, the concept of exemplification has been relatively neglected' (2000, p. 215). We have been careful to remain silent on the issue of whether representation-as applies to pictorial representation. But it seems clear to us that the most powerful objection against GE with respect to pictorial representation as representation-as are irrelevant here. The objection stems from the observation that we recognise what a picture represents by means of our perceptual experience (or at least something closely related to it) when looking at the picture. Schier calls this the 'natural generativity of pictures' (1986); relatedly Wollheim refers to the perceptual skill of 'seeing-in' which itself

was influenced by Gombrich's account of pictorial representation as being based on the illusion of thinking we are viewing the target of the representation (Gombrich 1960).²⁴

Regardless of how plausible this objection is in the case of pictorial representation, it is irrelevant in the context of scientific representation. There is no 'natural generativity' of scientific models, and no phenomenological experience of 'seeing-in'. If there were, it would make the practice of science much easier than it actually is. That models need to be interpreted, both to determine what is exemplified, and to find an appropriate key for translating such properties into ones imputed on their target systems, is a significant part of scientific practice.

Finally, a substantial part of a critical discussion of GE's account of pictorial representation has focused on how they distinguish it from other kinds of representation. The details of this distinction – involving syntactic and semantic density and repleteness – go beyond our current purpose (see Kulvicki (2006a) for an accessible overview of these notions, and criticisms thereof). Again, these elements play no role in our account and we do not use them to demarcate scientific representation from other kinds of representations. As such, criticisms against them do not bite here.

So, the most pertinent criticisms against GE's account of representation in aesthetics are irrelevant to the position advanced in this paper. We conclude that scientific representation should be seen as representation-as. This claim has been made before, but not with the level of detail needed to fully understand its impact. In this paper we have provided such detail. What falls out is this: denotation, exemplification, translation keys, and imputation are individually necessary and jointly sufficient conditions on scientific representation. As we have discussed, these conditions are abstract in the sense that they need further specification in each instance, but they remain general enough to capture what all cases of scientific representation have in common. The next step is to understand the instances, and the framework we have presented tells us how to do that. When faced with a specific model representing a specific target the questions to ask are the following: How is denotation established? What interpretation is adopted? What properties does the model exemplify and how does it make them salient? What properties does the key translate, and what properties does it translate them to? Which of these are imputed onto the target, and how? Answers to these questions, across multiple instances, will then deliver a rich understanding of how scientific models represent the world.

Acknowledgements We are extremely grateful to Catherine Z. Elgin for extended discussions as well comments on an earlier draft of this paper. Furthermore, Julian Reiss and two anonymous referees provided useful feedback. We are also in debt to audiences at the universities of East Anglia, Toronto, and Bristol for their stimulating comments.

²⁴See Kulvicki (2006b) for an overview of the variety of proposed accounts of pictorial representation.

References

- Ankeny, R. A. (1998). Fashioning descriptive models in biology: Of worms and wiring diagrams. *Philosophy of Science*, 67, S260–S272.
- Argyris, J. H., Faust, G., & Haase, M. (1994). *Die Erforschung Des Chaos: Eine Einführung Für Naturwissenschaftler Und Ingenieure*. Braunschweig: Vieweg+Teubner Verlag.
- Arrell, D. (1987). What Goodman should have said about representation. *The Journal of Aesthetics and Art Criticism*, 46, 41–49.
- Callender, C., & Cohen, J. (2006). There is no special problem about scientific representation. *Theoria*, 55, 7–25.
- Cartwright, N. (1999a). *The dappled world: A study of the boundaries of science*. Cambridge: Cambridge University Press.
- Cartwright, N. (1999b). Models and the limits of theory: Quantum Hamiltonians and the Bcs models of superconductivity. In M. Morgan & M. Morrison (Eds.), *Models as mediators: Perspectives on natural and social science* (pp. 241–281). Cambridge: Cambridge University Press.
- Contessa, G. (2007). Scientific representation, interpretation, and surrogative reasoning. *Philosophy of Science*, 74, 48–68.
- Danto, A. (1981). *Transfiguration of the commonplace: A philosophy of art*. Cambridge, MA/London: Harvard University Press.
- Elgin, C. Z. (1983). *With reference to reference*. Indianapolis: Hackett.
- Elgin, C. Z. (1996). *Considered judgement*. Princeton: Princeton University Press.
- Elgin, C. Z. (2004). True enough. *Philosophical Issues*, 14, 113–131.
- Elgin, C. Z. (2009). Exemplification, idealization, and scientific understanding. In M. Suárez (Ed.), *Fictions in science. Philosophical essays on modeling and idealization* (pp. 77–90). New York/London: Routledge.
- Elgin, C. Z. (2010). Telling instances. In R. Frigg & M. C. Hunter (Eds.), *Beyond mimesis and convention: Representation in art and science* (pp. 1–18). Berlin/New York: Springer.
- Frigg, R. (2003). Self-organised criticality-what it is and what it isn't. *Studies in History and Philosophy of Science*, 34, 613–632.
- Frigg, R. (2010). Fiction and scientific representation. In R. Frigg & M. Hunter (Eds.), *Beyond mimesis and convention: Representation in art and science* (pp. 97–138). Berlin/New York: Springer.
- Frigg, R., & Nguyen, J. (forthcoming). Models and representation. In L. Magnani and T. Bertolotti (eds.) *Springer Handbook of Model-Based Science*.
- Giere, R. N. (2004). How models are used to represent reality. *Philosophy of Science*, 71, 742–752.
- Giere, R. N. (2010). An agent-based conception of models and scientific representation. *Synthese*, 172, 269–281.
- Godfrey-Smith, P. (2006). The strategy of model-based science. *Biology and Philosophy*, 21, 725–740.
- Gombrich, E. (1960). *Art and illusion: A study in the psychology of pictorial representation*. London: Phaidon Press.
- Goodman, N. (1972). Seven strictures on similarity. In N. Goodman (Ed.), *Problems and Projects* (pp. 437–446). Indianapolis: Hackett.
- Goodman, N. (1976). *Languages of art* (2nd ed.). Indianapolis/Cambridge: Hackett.
- Goodman, N. (1984). *Of mind and other matters*. Cambridge, MA: Harvard University Press.
- Hartmann, S. (1995). Models as a tool for theory construction: Some strategies of preliminary physics. In W. E. Herfel, W. Krajewski, I. Niiniluoto, & R. Wojcicki (Eds.), *Theories and models in scientific processes (Poznan studies in the philosophy of science and the humanities 44)* (pp. 49–67). Amsterdam: Rodopi.
- Hughes, R. I. G. (1997). Models and representation. *Philosophy of Science*, 64, S325–S336.
- Jones, M. (2005). Idealization and abstraction: A framework. In M. Jones & N. Cartwright (Eds.), *Idealization Xii: Correcting the model-idealization and abstraction in the sciences*. Amsterdam: Rodopi.

- Kulvicki, J. (2006a). *On images: Their structure and content*. Oxford: Oxford University Press.
- Kulvicki, J. (2006b). Pictorial representation. *Philosophy Compass*, 1, 535–546.
- Morgan, M. (2012). *The world in the model. How economists work and think*. Cambridge: Cambridge University Press.
- Morrison, M. (2008). Models as representational structures. In S. Hartmann, C. Hofer, & L. Bovens (Eds.), *Nancy Cartwright's philosophy of science*. New York: Routledge.
- Robinson, J. (2000). Languages of art at the turn of the century. *The Journal of Aesthetics and Art Criticism*, 58, 213–218.
- Schier, F. (1986). *Deeper in pictures: An essay on pictorial representation*. Cambridge: Cambridge University Press.
- Sterrett, S. G. (2002). Physical models and fundamental laws: Using one piece of the world to tell about another. *Mind and Society*, 5, 51–66.
- Sterrett, S. G. (2006). Models of machines and models of phenomena. *International Studies in the Philosophy of Science*, 20, 69–80.
- Suárez, M. (2004). An inferential conception of scientific representation. *Philosophy of Science (Supplement)*, 71, 767–779.
- Teller, P. (2001). Twilight of the perfect model model. *Erkenntnis*, 55, 393–415.
- Thomson-Jones, M. (2010). Missing systems and face value practice. *Synthese*, 172, 283–299.
- van Fraassen, B. (2008). *Scientific representation: Paradoxes of perspective*. New York: Oxford University Press.
- Webb, B. (2001). Can robots make good models of biological behaviour? *Behavioral and Brain Sciences*, 24, 1033–1050.
- Weisberg, M. (2007). Three kinds of idealization. *The Journal of Philosophy*, 104, 639–659.
- Weisberg, M. (2013). *Simulation and similarity: Using models to understand the world*. Oxford: Oxford University Press.

Roman Frigg is Professor of Philosophy in the Department of Philosophy, Logic and Scientific Method, Director of the Centre for Philosophy of Natural and Social Science (CPNSS), and Co-Director of the Centre for the Analysis of Time Series (CATS) at the London School of Economics and Political Science. He is the winner of the Friedrich Wilhelm Bessel Research Award of the Alexander von Humboldt Foundation. He is a permanent visiting professor in the Munich Centre for Mathematical Philosophy of the Ludwig-Maximilians-University Munich and held various visiting appointments. He is associate editor of the *British Journal for the Philosophy of Science*, member of the steering committee of the European Philosophy of Science Association, and serves on a number of editorial and advisory boards. He holds a PhD in Philosophy from the University of London and masters degrees both in theoretical physics and philosophy from the University of Basel, Switzerland.

James Nguyen is Postdoctoral Research Associate at the University of Notre Dame and Research Associate of the Centre for Philosophy of Natural and Social Science (CPNSS) at LSE. He holds a PhD in Philosophy from the LSE, having received previous degrees from the University of Cambridge and King's College London. His primary area of interest is in the philosophy of science where he specialises on topics surrounding scientific modelling and representation in particular.