Synthese (2010) 172:197–213
DOI 10.1007/s11229-009-9502-3

Informational versus functional theories of scientific representation

Anjan Chakravartty

Received: 21 March 2008 / Accepted: 5 March 2009 / Published online: 26 March 2009

Abstract Recent work in the philosophy of science has generated an apparent conflict between theories attempting to explicate the nature of scientific representation. On one side, there are what one might call 'informational' views, which emphasize objective relations (such as similarity, isomorphism, and homomorphism) between representations (theories, models, simulations, diagrams, etc.) and their target systems. On the other side, there are what one might call 'functional' views, which emphasize cognitive activities performed in connection with these targets, such as interpretation and inference. The main sources of the impression of conflict here are arguments by some functionalists to the effect that informational theories are flawed: it is suggested that relations typically championed by informational theories are neither necessary nor sufficient for scientific representation, and that any theory excluding functions is inadequate. In this paper I critically examine these arguments, and contend that, as it turns out, informational and functional theories are importantly complementary.

Keywords Scientific representation · Theory · Model · Target system · Similarity · Isomorphism

1 An apparent dichotomy in accounts of representation

Contemporary philosophy of science has witnessed a provocative movement away from thinking of theories as the primary units of analysis in discussions of scientific knowledge. In more practice-oriented philosophy of science, this move has, no doubt, been in effect for some time, for several reasons. For example, in practice, the

A. Chakravartty (✉)
Institute for the History and Philosophy of Science and Technology, Victoria College,
University of Toronto, 91 Charles Street West, Toronto, ON M5S 1K7, Canada
e-mail: anjan.chakravartty@utoronto.ca

term ‘theory’ may not correspond uniquely to any very well-defined concept. Is the Newtonian theory of mechanics a set of mathematical equations, or a collection of models pertaining to different classes of mechanical phenomena, or something else, perhaps? Some theories are identified with problem-solving exemplars that are, as it happens, mutually inconsistent, but we do not generally conclude thereby that such theories are inconsistent. As a consequence of these and other revelations, many have focused more determinedly on scientific practice to gain better insight into the forms that scientific knowledge takes, and this attention suggests that it is probably a mistake to think that the term ‘scientific theory’ labels any one, definitive sort of entity. When scientists themselves talk about theories, they often refer to different things in different contexts: very general, mathematical laws; very specific, problem-solving techniques; and so on. Precisely the same sort of promiscuity, I think, applies to the now fashionable term ‘model’.

Of course, ‘theory’ and ‘model’ are indispensable terms of art, and it would be wrong to suggest that perspicuous discussions of scientific knowledge can do without them, suitably defined. In an effort to gain deeper insight into the nature of this knowledge, however, many philosophers of science have now shifted their attention to the question of how these things, variously defined, *represent* aspects of the world. The deeper question seems to concern the nature of representation. But what is a scientific representation? ‘What is that thing’ questions are requests for clarification regarding ontology, and there are different ways of shedding light on the ontological nature of a thing. One way is to say something about the category or categories to which the thing belongs, as when people argue that scientific representations comprise abstract entities (such as theoretical models), concrete objects (such as diagrams, graphs, and illustrations), and processes (such as computer simulations). Another way to shed light on the ontological nature of a thing is to say something about its properties, and it is this latter sort of aim that is my focus here. What, one might ask, are the “essential” properties of a scientific representation?[1]

Recent discussions of scientific representation offer what may appear to be two broad and conflicting approaches to this question. On one hand, there are theories that emphasize what I will call *information*. The idea here is that a scientific representation is something that bears an objective relation to the thing it represents, on the basis of which it contains information regarding that aspect of the world. By ‘objective’ I simply mean that, although generally the fact that such relations obtain can be grasped only by means of representational conventions, they are nonetheless mind-independent, in a sense to be clarified (in Sect. 2). And while ‘information’ is a technical concept in various sub-specialties of philosophy, cognitive science, and computer science, here I intend it simply in its everyday or colloquial sense, as whatever it is that we learn when we gain knowledge of something. The most general version of the informational approach appeals to relations of similarity. As Giere (1988, Chap. 3; 1999) puts it, scientific representations are similar to their target systems in certain specified respects, and to certain degrees. The generic relation of similarity admits of several species championed by other proponents of the informational view, including

1 Frigg (2006, p. 50) describes this—‘in virtue of what is a model a representation of something else?’—as a semantic issue: a question regarding what it means to call something a scientific representation.

relations of isomorphism, partial isomorphism, and homomorphism. But all share an emphasis on objective relations bearing information.[2]

In contrast, the other broad approach to scientific representation comprises theories that emphasize the *functions* of representations: their uses in cognitive activities performed by human agents in connection with their targets. The idea here is that a scientific representation is something that facilitates these sorts of activities, which likewise fall into several categories championed by different proponents of the view. Some appeal to the demonstrations and interpretations of target systems that representations allow, and others to the inferences they permit concerning aspects of world.[3] Elgin (2004, p. 124), for example, inspired by Goodman (1976), emphasizes the notion of 'exemplification': 'the device by which samples and examples [i.e. representations] highlight, exhibit, display, or otherwise make manifest some of their features', while other features are in various ways downplayed or ignored. Exemplification depends crucially on the cognitive activities of human agents, wherein representations are mentally processed in such a way as to render certain of their features epistemically accessible, which may require specific background assumptions, knowledge, or expertise on the part of the thinking subject. More generally, all versions of the functional approach emphasize the centrality of human agents in giving an account of scientific representation, in apparent contrast to the mind-independence of the informational approach.

In this paper, I will argue that the dichotomy between emphasizing informational relations and functions suggested by several recent discussions of scientific representation is, in fact, a false dichotomy. The primary sources of the view that these accounts stand opposed are various arguments by some proponents of the functional approach, to the effect that informational theories are problematic. There are three main avenues of criticism here, and I will consider each in turn. The first is that the relations between representations and their targets typically championed by informational theories are not necessary for scientific representation. The second is that such relations are not sufficient. The third is that certain functions are essential to representation in this context, and that any theory excluding them is thus inadequate. These arguments, I believe, require some scrutiny, for they fuel a mistaken impression of rival accounts of representation. Informational and functional theories are in fact complementary, both contributing to a general understanding of scientific representation. In conclusion, I will briefly consider the likelihood of a definitive account of the essential features

[2] Regarding isomorphism, see van Fraassen (1980, Chap. 3; 1989, Chap. 9), and French (2003); regarding partial isomorphism, see da Costa and French (2003, Chap. 3); regarding homomorphism, see Bartels (2006). Bartels (2006, p. 17) actually denies that his account is a version of the similarity theory, since similarity is a symmetrical relation, and homomorphism is not. He is correct of course that homomorphism is not a symmetrical mapping, but this misses the point, I think: if there is a homomorphism from A onto B, then A and B have some structure in common, and are thus similar (indeed, exactly similar) in that respect. Similarly, Suárez and Solé (2006, p. 44) contend that since similarity and isomorphism are distinct relations, and both are means of scientific representation, such means are 'irreducibly plural'. But isomorphism, like homomorphism, can be described in terms of similarity.

[3] Regarding demonstration and interpretation, see Hughes (1997); regarding interpretation as a foundation for inference, see Contessa (2007); regarding inference, see Suárez (2004).

of representation in the sciences, and speculate as to why certain misunderstandings surrounding these issues have occurred.

2 The first charge: non-necessity

Two main considerations are commonly cited as evidence for the claim that the sorts of relations invoked by informational theories are not necessary for scientific representation. The first can be viewed as targeting the idea of similarity generally, and thus applying both to this generic formulation of the informational view, as well as its more specific versions. The second targets particular species of similarity, such as isomorphism. Let us examine these two considerations in turn.

Regarding the first, let me focus on the most general formulation of the informational account, on the understanding that the objection here, if successful, would of course apply to all of its species as well. Thus, consider the view that a scientific representation is similar in some way, and to some extent, to the things it represents. Similarity ultimately admits of different analyses depending on one's finer-grained metaphysical commitments (to universals, tropes, or nominalism), but a noncommittal understanding of it will suffice for present purposes: A and B are similar iff they have some subset of their properties in common.[4] Now, there is an innocuous sense in which similarity thus defined is clearly unnecessary for representation in many contexts, because representation is something that is often established merely by *fiat*. While on sabbatical in Barcelona I may send a postcard of a sumptuous Gaudi building to my colleagues in Toronto, and they may place it on my chair at the next faculty meeting and take it to represent me, if they wish. It is true, of course, that the postcard and I are similar in some respects, since any two things are similar in *some* respects, but the properties we happen to share, such as having mass, or being such that there is no present king of France, are not ones that are relevant to the representation. I will assume henceforth that similarities apparently playing no role in given cases of representation do, in fact, play no role. Similarly, I may take the word 'cats' to represent cats, simply by fiat of linguistic convention, even though the word 'cats' may not bear any relevant or interesting similarities to cats. Thus, similarity is clearly unnecessary for representation in many cases.

This points to an important disanalogy between cases of representation by fiat and the kinds of scientific representation at issue here. It would be a mistake to suggest that the former is absent from scientific domains, since scientists name entities and processes—'quasar', 'Krebs cycle'—in just the same or in similar sorts of ways as those used to name things like cats more generally. The representations at issue here, however, are not names or nouns, but entities such as theories, models, diagrams, simulations, graphs, illustrations, and so on. In the sciences, something more than

4 This definition may seem unduly restrictive, but I suspect that on a liberal enough view of properties, it accommodates all manners of similarity. For example, one might worry that two planets could have similar masses without this being analyzable in terms of a common property, but presumably both would share the property of having mass $m \pm \delta$ kg, where m and δ are an appropriate magnitude and error tolerance respectively. The same goes, *mutatis mutandis*, for representational and target system parameters, which may share properties of the form: having magnitude $x \pm \delta$.

merely wishing it were so, or deciding it is so, is involved in making things such as these into representations of their target systems. In debates concerning accounts of scientific representation, it is these latter kinds of things that are intended, and I will use the term ‘scientific representation’ to refer to them exclusively henceforth. (One might, of course, maintain that even in these cases, mere fiat *can* establish representation, and that something further is required only in cases of successful or accurate representation. If one wishes, one may take the qualification as given, but I will return to this point at the end of Sect. 4.)

Why not, then, take similarity to be a necessary feature of scientific representation? The first non-necessity argument is straightforward, and runs as follows: the informational view holds that similarity is necessary, but there are indisputably genuine scientific representations that are not similar to their targets; therefore, similarity is not necessary for scientific representation. One might reasonably wonder, however, how there *could* be cases of genuine scientific representation in which no similarity relations obtain. According to functional theories, a scientific representation is something that facilitates practices such as interpretation and inference with respect to its target system. And how, one might wonder, could such practices be facilitated successfully, were it not for some sort of similarity between the representation and the thing it represents—is it a miracle? The first indication that something has gone wrong in the rejection of similarity as a necessary condition for scientific representation, is that without it, the success of the very functions that functional accounts take to be central would appear to be inexplicable.[5]

No doubt, in some cases, the respects of similarity may be few. Models may prove instrumentally useful, for instance, even in cases where we believe that their internal features bear little or no resemblance to their targets, or where we are agnostic about such resemblance. But even in these cases, specified parameter values taken by such models must be similar to those taken by their targets—if not, we would not adopt them as instrumentally useful representations. It is thus no surprise, I suspect, that putative cases of scientific representation in the absence of similarity are few and far between. Concrete representations, such as Watson and Crick’s demonstration model of the DNA molecule, and abstract representations, such as the model of the simple pendulum, clearly have certain properties in common with their targets—structural features, for instance. So let me consider the one case that does seem to present a *prima facie* difficulty for an analysis of scientific representation in terms of similarity: linguistic representation. And recalling the demarcation I established earlier regarding the forms of representation relevant here, by ‘linguistic representation’, let me refer to scientific descriptions associated with devices such as sentences and mathematical expressions, as opposed to terms associated merely with naming and the like.

On their face, linguistic descriptions do not appear similar to their targets in any interesting sense, but nevertheless represent them. It is thus no surprise, perhaps, that most advocates of the informational approach *also* subscribe to the semantic view of theories, according to which scientific theories are simply families of models—that is, non-linguistic entities—as opposed to linguistic devices. This, I believe, contains

[5] For work detailing connections between some specific similarities and functions, see Swoyer (1991).

a clue regarding the correct response to the charge that similarity is unnecessary for scientific representation: the semantic *content* of a (successful) linguistic description bears some similarity to its target, even though the superficial means by which that content is expressed (likely) does not. To borrow an example from Suárez (2003, pp. 231–232), the quantum state diffusion equation for a particle subject to a localization measurement, as written in a textbook or on a blackboard, does not appear similar to the properties of any particle, to be sure. But this, surely, is to see the equation in a superficial way, as merely blotches of ink on a page, or intricately-shaped trails of chalk dust. Anyone who sees only this when viewing such a thing would presumably have no grounds for thinking that it was a representation at all. On the other hand, having learned the languages of mathematics and physics, one may view the content of these blotches and trails as trajectories in phase space, and these most certainly have features in common with the states of particles subject to localization measurements.

The idea of grasping the semantic content of a linguistic expression is important to a clarification of the sense in which, as I mentioned earlier, informational approaches take similarity relations (or more specifically, relations of isomorphism, etc.) to be objective. On the informational view, such relations are mind-independent in the sense that they obtain between scientific representations and their targets quite independently of the various cognitive activities emphasized by the functional approach. Some care is required here, though, for there is another sense in which these relations are clearly mind-dependent. Aspects of the content of a linguistic description can only be viewed as similar to aspects of a target system insofar as one has grasped the semantics of the language employed. Learning a language is part of the more general phenomenon of learning the representational conventions according to which one thing is taken to represent something else, and I will have more to say about this later. The important point here is simply that, given an appropriate semantics and conventions of representation, there is an objective fact of the matter about whether the content of a linguistic description is or is not similar in specified respects and degrees to its target. Thus, it is too quick to say that these kinds of representations do not furnish similarity relations in connection with their targets.

None of this should be taken to suggest, of course, that the notion of content here is entirely clear—far from it. The question of how we are able to learn languages and grasp the content of well-formed expressions in them is infamously thorny. One strategy for beginning to answer this question in the scientific context is to adopt the semantic view, and thereby analyze the idea of grasping content in terms of acquiring some sort of familiarity with the abstract models with which linguistic descriptions are associated; other possibilities may appeal to different sorts of entities, such as propositions. But these are just beginnings, and in any case, it is not my intention to endorse any particular strategy here. However philosophers of language ultimately answer this question, the moral for present purposes is the same: the content of linguistic representations in the sciences is informative with respect to their targets, precisely because it bears specifiable relations of similarity to those targets. Mathematical equations, for example, generally contain variables and express relations between the properties and quantities these variables represent. To the extent that the variables refer, and the relations expressed obtain, the informational view is satisfied.

Indeed, it is precisely *because* the informational view is satisfied in this way that cognitive activities such as interpretations and inferences regarding target systems are successful in the first place. Earlier I suggested that in the absence of substantive relations of similarity (in some form or other) between scientific representations and their targets, it would be something of a mystery how these devices represent things in scientifically interesting ways at all. As even Suárez (2003, p. 229) admits: if A represents B, 'then A must hold some particular relationship to B that allows us to infer some features of B by investigating A'. That relationship, I suggest, whatever it may be, will be an instance of similarity (see n. 2). Certainly, before one can interpret a scientific representation in connection with its target, or make an inference about that system on the basis of the representation, one must first understand the languages and representational conventions employed in the construction of that representation. One must learn the relevant bits of the language of mathematics, for example, in order to interpret trails of chalk dust as trajectories in phase space, prior to any interpretations or inferences involving subatomic particles. But having understood this, our mystery is resolved: similarity is a necessary condition of scientific representation. Lines of chalk are not by themselves representations; they *become* representations when we learn how to interpret the relevant languages and conventions, and grasp the relations of similarity they express.

Let me now turn to the second main consideration that one might take to undermine the idea that the sorts of relations invoked by informational theories are necessary for scientific representation. Recall that this objection focuses on more specific relations between representations and targets than mere similarity, isomorphism constituting a prime example. Here I will be brief, for it seems undeniable that most if not all cases of scientific representation fall short of the rigorous standards suggested by precise mathematical relations such as isomorphism and homomorphism. The fact that theories, models, and other scientific representations are generally idealized is widely appreciated, and given the ubiquity of idealization, the relevant structures of many if not most representations will not stand in precisely these sorts of relations to the corresponding structures characterizing their targets (or models of the data concerning those targets). It is in part because of this that some advocates of the informational view have aimed to describe relations that allow for more leeway in the comparison of structures.[6]

The fact that some proponents of the informational approach have described less stringent relations between scientific representations and their targets is interesting in its own right, but regardless, the ubiquity of idealization and the failings of strict relations such as isomorphism are not by themselves sufficient to undermine these versions of the informational view. Consider an analogy: scientific realism is often described as the view that our best scientific representations are true, or yield truths about the world. Yet most realists happily agree that most scientific representations are in fact false, strictly speaking. Similarly, empiricists may hold that such

[6] For example, see the treatment of isomorphism in Mundy (1986), and of partial isomorphism in da Costa and French (2003). It is an important qualification here that strictly speaking, these morphisms hold between mathematical objects such as representational models and models of the data (as opposed to target systems per se). I take this as given in discussing relations between scientific representations and "the world".

representations are empirically adequate, yet happily admit that this is often not the case, strictly speaking. Are these positions rather obviously internally inconsistent? That would seem an uncharitable interpretation. Claims regarding the truth or empirical adequacy of scientific representations in these contexts, not to mention the presence of relations such as isomorphism, come with an implicit caveat: they must be understood as approximations. Indeed, they are idealizations, reasonably employed in the course of epistemological theorizing. No doubt questions regarding whether such idealizations are justified, and what the relevant accounts of approximation might look like, are important and demand attention—but these are separate matters.

I have argued elsewhere (2004; 2007), for independent reasons, that scientific knowledge might plausibly be thought to comprise more than knowledge of the sorts of abstract mathematical properties that can be inferred from relations such as isomorphism and homomorphism alone, and indeed, there is some controversy as to whether accounts of knowledge exclusively in terms of such properties are trivial.[7] But these are concerns about whether such relations are sufficient, not whether they are necessary, for an account of scientific representation. By itself, the perhaps ubiquitous failure of strictly defined mathematical similarities between representations and their targets tells us nothing about whether such similarities generally obtain, not strictly, but within reasonable bounds of approximation.

3 The second charge: non-sufficiency

One of the most puzzling worries suggested in connection with informational theories of scientific representation is the claim that the kinds of relations these accounts invoke are not sufficient for scientific representation. It should be noted immediately that the reason this worry is puzzling is not that it is incorrect to say that relations such as similarity are insufficient. Goodman (1976, pp. 3–4) was surely right when he maintained that ‘plainly, resemblance in any degree is no sufficient condition for representation’. Indeed, the truth of Goodman’s dictum would seem to be overdetermined by a number of compelling observations. It is widely noted, for example, that if one takes relational properties (such as the property of being temporally located after the Big Bang) and mere-Cambridge properties (such as being such that José Luis Rodríguez Zapatero is the Prime Minister of Spain) into account, any two things are guaranteed to bear an infinite number of similarity relations to one another; but it is clearly not the case that all pairs of things stand in representational relationships.

Many have noted that even if one were to exclude relational and mere-Cambridge properties, relations such as similarity would still not be sufficient for representation. Goodman (1976, pp. 35–36) illustrates this with a nice example:

> Consider a realistic picture, painted in ordinary perspective and normal colour, and a second picture just like the first except that the perspective is reversed and each colour is replaced by its complementary. The second picture, appropriately interpreted, yields exactly the same information as the first. And any

[7] This is the so-called ‘Newman problem’. For an introduction, see Demopoulos and Friedman (1985), Worrall and Zahar (2001), Ketland (2004), and Melia and Saatsi (2006).

> number of other drastic but information preserving transformations are possible. Obviously, realistic and unrealistic pictures may be equally informative; informational yield is no test of realism. … The two pictures just described are equally correct, equally faithful to what they represent, provide the same and hence equally true information; yet they are not equally realistic or literal. … Just here, I think, lies the touchstone of realism: not in quantity of information but in how easily it issues. And this depends upon how stereotyped the mode of representation is, upon how commonplace the labels and their uses have become.

To paraphrase, as I suggested in Sect. 2, scientific representation is achieved only in circumstances in which agents know or have otherwise mastered the system of representation being used to encode information about whatever it is that is represented. That is why relations such as similarity cannot do the job on their own; these relations only serve the goal of representation subject to the internalization of the semantics of their forms of expression and relevant representational conventions by their users, either by means of hard-wired cognitive responses, scientific or other training, or both.

This stand on the non-sufficiency of informational relations for representation has been augmented by a number of authors recently. Suárez (2003), for example, argues convincingly that relations such as similarity and isomorphism cannot be sufficient for scientific representation, for such relations are symmetric and reflexive, whereas representation is clearly non-symmetric and non-reflexive. Something more than these relations is certainly required to establish the essential directionality characteristic of scientific representation. In a supporting vein, Elgin (2006) argues that although such representations are informative because they convey information about certain properties of their target systems (or related ones, in cases of idealization) on the basis of similarities, similarity relations by themselves are not sufficient, since in addition to bearing similarities to their targets, scientific representations must exemplify similarities of interest in given contexts of investigation. Exemplification will generally involve the application of representational conventions by human agents, in order to highlight certain properties at the expense of others. Yet again, we arrive at the conclusion that the kinds of relations invoked by informational theories are not sufficient for scientific representation.

Why then did I describe this worry as puzzling? It is puzzling not because the claim regarding non-sufficiency is incorrect, but because it is no part of the informational view that relations such as similarity or isomorphism are sufficient for representation. The two are sometimes suggestively linked, however. Given the emphasis placed by informational accounts on these kinds of relations, there is perhaps an understandable temptation to interpret them this way. Suárez (2003, p. 225) directs his critique against 'theories that attempt to reduce scientific representation to similarity or isomorphism'; that 'aim to radically naturalize the notion of representation, since they treat scientists' purposes and intentions as non-essential to representation'. Most proponents of the informational view, however, make no such attempt and have no such aim.[8]

[8] Suárez (2004, p. 768) recognizes that neither Giere nor van Fraassen, to whom he attributes what I call informational theories, are proper targets of his critique. Both acknowledge the role of human intentions. Suárez (2003, p. 229) suggests that while recently, Giere (1999, 2004) disavows the form of naturalism he

Indeed, perhaps the *least* controversial feature of scientific representation is the idea of intentionality: a representation is something that is *about* something else, and it cannot be a representation unless there is something that it represents. The notion of intentionality is so basic to scientific representation, I suggest, that in the absence of an explicit claim to the contrary, the principle of charity simply demands that it be considered an implicit assumption on the part of *any* theory of scientific representation. And any theory adopting this assumption will view such representation, correctly, as non-symmetrical and non-reflexive.

This is not to say, of course, that the informational approach is inconsistent with a naturalism excluding agents' goals and purposes, even if no one would dispute the importance of intentionality. For it is possible to construe intentionality as a mind-independent feature of representations—this would preserve the non-symmetricality and non-reflexivity of scientific representation, while simultaneously removing human agents from the equation. This does not seem promising as a view of the intentionality of *scientific* representations, however. In a grand tradition of thought experiments involving monkeys, typewriters, ants, and Churchill, French (2003, p. 1473) imagines a scenario in which the wind and sea carve the Lorentz transformations into the sand of a beach. Our intuition, he suggests, is that these grooves represent relativistic phenomena, quite independently of any person's intentions. I am not sure how widespread this intuition might be, but there would appear to be two difficulties with it. To the extent that one does have the nagging suspicion that accidents on a beach may represent relativistic phenomena, it is only (I submit) because one is armed with the prior intention to use markings with these shapes to represent such phenomena. Their representational status, if they have any, is certainly independent of any intentions to *construct* the representation, since in this case no one does, but nevertheless seems dependent on other, previously-formed intentions, and this raises a second worry about French's intuition. If the intentionality of beach equations does not have its source in human intentions, from whence does it come?

Bartels (2006, p. 12) offers a possible answer to this question. In addition to agent-based intentions, he suggests, causal relations may furnish a further and naturalistic source of intentionality, as in the case of photography, where photographic images are the end results of causal processes involving the things they represent. I suspect, however, that just as in the beach case (*mutatis mutandis*), the representational status of photographs is parasitic on agents' intentions with respect to imaging technologies.[9] Like similarity, causal relations do not by themselves establish the intentionality of representations in the scientific context. Perhaps there are other contexts in which causal relations *are* sufficient. Philosophers of perception and biology sometimes speculate that things like perceptual states may acquire their intentionality merely in virtue of certain causal relations between an organism's environment and its

Footnote 8 continued
describes, his earlier work (1988) endorses it, but this is not clear. For further reservations along these lines, see Contessa (2007, p. 53, n. 6).

[9] There are differing views on this, however: Scruton (1983, Chap. 9) argues that photographs do not represent at all; Currie (1995, Chap. 2) thinks they do, but in a way unlike other pictures. See also Walton (1984).

sensory and other modalities. By ‘other modalities’ here, I have in mind physiological or biochemical states not usually associated with the senses. Consider, for example, the ability of some organisms to align their migrations with the earth’s magnetic poles, as a consequence of the effects of the magnetic field on their systems. But these seem like special cases, and whatever our considered view of them, they do not seem analogous to cases of representation in the scientific context, where human intentions with respect to theories, models, and so on, are key.

As a final point in their favour, consider that agent-based intentions also help to explain how *relevant* similarities are picked out between scientific representations and their targets: they are those that users of representations take to be relevant. In the absence of such intentions, it is an open question how naturalistic possibilities such as those suggested by French and Bartels can furnish a criterion of relevance. In any case, concerns about relevance aside, the important point for present purposes is simply that intentionality is a widely-acknowledged, fundamental feature of representation, and a commitment to this idea is perfectly consistent with and accepted by most informational theories of scientific representation, whether explicitly or implicitly. The relations invoked by informational theories are not generally nor plausibly offered as sufficient conditions for scientific representation. As such, the fact that these relations are not sufficient is no argument against the informational view.

4 The third charge: essential functions

In the process of excusing informational theories from the charge of non-sufficiency, I noted that they are perfectly compatible with a further condition that is commonly assumed in connection with scientific representation: the notion of intentionality. By extension, it is not difficult to see that in just the same way, informational accounts are perfectly compatible with the further conditions that functional theories take to be important to representation in this context, such as capacities to support interpretations and inferences regarding target systems. Some may feel, however, that merely pointing out these compatibilities is too forgiving. The informational theories I have mentioned all place their emphasis on what I have described as the information provided by representations by means of specific relations. Aspects of scientific representations are similar in specified ways to aspects of their target systems, and that is all. Given that representational functions such as interpretation and inference are so central to scientific work, however, is it not obvious that any theory of scientific representation that has nothing to say about these functions is rather missing the point? Here we have a third criticism of the informational approach: it has nothing to say about the essential functions of scientific representations, and is therefore defective. I believe that this worry is premised on a confusion, whose resolution may lead ultimately to a dissolution of the apparent dichotomy between informational and functional theories of scientific representation. Let me turn to these issues now.

Frigg (2006, p. 54) formulates the charge I have in mind here in a helpful way. In response to the claim that informational theories are consistent with agent-based intentions, goals, and purposes, and generally assume such things to play a role in scientific representation, either explicitly or implicitly, he remarks:

> Merely tacking on intentions as a further condition is question begging. To say S is turned into a representation because a scientist intends S to represent T is a paraphrase of the problem [of giving an account of scientific representation – of explaining why or how S represents T] rather than a solution.

Now, this is perhaps unwarranted, since 'tacking on intentions' is not intended by anyone to explain what it is to represent in this context, but rather to explain the source of intentionality characteristic of scientific representations. Nevertheless, the pressing idea that somehow the most important parts of scientific representation have been omitted by informational theories is palpable here. Frigg offers the analogy of attempting to give an account of reference. Claiming that speakers intend certain terms to refer to this or that does not tell us what reference is; in order to do the latter, one must say something more. The analogy is instructive, I think, because it illuminates what I take to be a confusion that has entered into some recent discussions of scientific representation. To say that a speaker intends the word 'cat' to refer to that fluffy, graceful, strikingly independent thing over there does not offer much if any insight into the nature of reference, certainly. But what if, in addition to finding out that a speaker intends to use the word 'cat' to refer to this or that, one is also told that referring terms are socially sanctioned strings of alphabetical symbols in languages such as English? One will not have learned much about what reference is, perhaps, but one will certainly have learned something about what a referring term is.

At the heart of the putative dichotomy between informational and functional theories of scientific representation is a conflation—a conflation of means and ends. It is a conflation of thinking about what scientific representations *are*, as a means to realizing their functions, and thinking about what we *do* with them. Informational theories focus primarily on the question 'what are scientific representations?', where representations are conceived as knowledge-bearing entities, such as theories, models, simulations, and diagrams. Functional theories focus primarily on the question 'what is scientific representation?', where 'representation' is conceived as a set of knowledge-exercising practices, constituted by whatever it is that scientists do when engaged in the process of representing things. These are two clearly related, but different questions, and it should come as no surprise that appropriate answers to these questions are clearly related, but different. Naturally, an account of knowledge-bearing entities emphasizes the relations in virtue of which knowledge is borne by those entities. And just as naturally, an account of representational processes emphasizes the various practices in virtue of which that knowledge is exercised. These are complementary questions and answers, both contributing to a general understanding of scientific representation. There is no dichotomy between information and function.

Perhaps an analogy will help to illustrate the point. When metaphysicians describe the natures of causally efficacious properties, they often avail themselves of two interestingly different kinds of descriptions. Categorical descriptions are ones that emphasize the static features of things, such as their dimensions, shapes, and configurations or arrangements; a molecule is tetrahedral, for example, if its atoms have a certain kind of orientation with respect to one another. Dispositional descriptions, on the other hand, are ones that emphasize how things having such properties behave under certain conditions; a substance is soluble if it dissolves when placed in a solvent, *ceteris paribus*. Categorical and dispositional descriptions are often used to illuminate complementary features of the nature of one and the same property. One may describe the mass of an object, for example, in terms of a quantity of mass-units, or in terms of how it is disposed to accelerate under an applied force. Analogously, informational and functional approaches to scientific representation focus on different aspects of one and the same thing: the nature of scientific representation.[10]

Having dissolved the impression of rival accounts, are we in a position, finally, to specify the essential features of scientific representation? This question, I believe, requires more consideration than I can give here, but perhaps it is worth sketching some of the issues that will require more detailed scrutiny before a thoroughly satisfying answer is forthcoming. For now it will seem tempting, perhaps, to identify the features of representation presented by informational and functional theories together as constituting necessary, and perhaps even jointly sufficient conditions for representation in the sciences. Such a proposal may well have a broad and intuitive appeal, prima facie: intentionality, relations of similarity (or more specific versions), and capacities to facilitate interpretations and inferences regarding target systems may well sound like plausible necessary and jointly sufficient conditions.

There is at least one good reason for caution, however, regarding the prospect of any straightforward prescription here. It is unclear whether in the sciences, it is appropriate to distinguish between mere representation and successful or accurate representation, where success or accuracy is understood to admit of degrees, from moderate to extraordinary. One might hold, for example, that while grossly false representations such as those associated with the humoural theory of disease or the vortex theory of gravitation are not, it turns out, successful or accurate representations of their intended target systems (since as it happens, these systems do not exist), they are nonetheless scientific representations. Conversely, one might hold that while such things are uncontroversially *scientific*, in the sense that they featured in past scientific investigations, and clearly constitute theories, models, and so on, our subsequent discoveries that they are grossly inaccurate should lead us to conclude that we were mistaken to say that they were *representations*, for as it turns out, they were not. We discover that a concrete model of the elastic solid ether, for example, while no doubt a scientific model, is not a representation after all, upon discovering that there is no such thing as the ether. One of the two relata of the intended representational relation is absent in this case.

The choice between whether or not to regard such theories, models, and so on as genuine representations, it seems, is significant. For the question of whether the relative accuracy of a putative representation should determine whether or not it counts as a representation at all in this context has immediate consequences for what one regards

[10] The idea of shifting emphasis explains how Giere can be both an informationlist (1988) and a functionalist (2004). In the latter mood, he stresses that he is sketching an account of 'the *activity* of representing' (p. 743, emphasis mine). There he notes that '[t]he assumption that scientific theories are sets of statements [commonly] goes along with the view that scientific representation is to be understood as a two-place relationship between statements and the world.' This may be a common assumption , but it is important to note that nothing in the view that theories are statements entails that the *activity* of representing is limited to a two-place relation. This would follow only if one conflates the issue of what a representation is (as a means) with the issue of what the activity of representing is (as a realization of ends)—a conflation I am urging against.

as the necessary (and potentially sufficient) conditions for scientific representation. Some form of intentionality is applicable across the board, and so too, perhaps, are functional capacities, since it is arguable that even very poor interpretations and inferences concerning grossly inaccurate representations are nonetheless interpretations and inferences. Things are more complicated when it comes to considering informational relations such as similarity, however. If one is inclined to accept the distinction between mere and (even moderately) accurate representations, and thus regard the latter as constituting a proper subset of the former, one may exclude informational relations as a necessary feature of mere representation, but insist on them in connection with accurate representation. If one is inclined to reject the distinction, and thus view some threshold of accuracy as demarcating genuine scientific representation, one may insist on informational relations as a condition for scientific representation *simpliciter*.

The question as to whether the central argument of this paper facilitates the project of giving an account of the necessary and sufficient conditions for scientific representation thus hinges on, *inter alia*, one's view regarding how accurate a putative representation must be in order to qualify as a genuine representation in the sciences. But here I see no facts of the matter, and as a consequence, no motivation for legislating intuitions on this point. I suspect that in scientific contexts where putative representations are so poor as to manifest no relevant similarities to their intended targets, many will be tempted to reject them as representational, and that where they manifest many such similarities, everyone will agree that they are. One may be tempted to draw a line between these cases, but it seems unlikely that there is any indefeasible reason for drawing it in any one place. The term 'scientific representation', much like the terms 'theory' and 'model', is a term of art. We may define it as best serves the various philosophical uses to which it is put. And if in the context of some philosophical investigation it makes sense to distinguish mere from successful or accurate scientific representation, this should not be taken as a challenge to informational theories. For in this case, informational theories are theories regarding successful or accurate representation, and nothing is lost in the qualification.

Some may balk at the suggestion that answering the question of whether 'scientific representation' is a success term is simply a matter of convention, but it is well supported, I think, by plausible and conflicting intuitions on either side. If one counts only scientific representations standing in sufficiently good informational relations to their targets as genuinely representational, one will count anything falling below that threshold—severe enough misrepresentations—as no representations at all. This will seem awkward to anyone sharing the intuition that a badly failed representation is a representation nonetheless, albeit a poor one. On the other hand, this usage does justice to the intuition that the intentionality of a putative representation, its "aboutness", is simply lacking in cases where the relata of putative representations turn out not to exist, as in the case of models of the ether. Conversely, if one counts all merely intended scientific representations as genuine representations regardless of whether they stand in informational relations to their intended targets, the intuition that severe misrepresentations are nonetheless representations is satisfied, but at the cost of frustrating the intuition that 'representation' is not a predicate that is sensibly applied to things that do not represent anything. Furthermore, if all merely intended scientific representation is genuine representation, then the term 'scientific representation' connotes nothing distinctive, for 'representation of the ether' would seem to mean nothing more nor less than 'model of the ether' (or 'diagram of the ether', or what have you). One might plausibly maintain, however, that the term 'scientific representation' should connote something other than the mere acknowledgement that something is a model.[11]

5 Postscript: fetishizing practice after positivism

If the diagnosis of the explanatory aims of different approaches to understanding scientific representation given in Sect. 4 is correct, then it seems we may now be in a position to move forward with discussions of these issues in a more irenic manner. On reflection, I suspect that the case of "rival" accounts of scientific representation is an instance of a more general methodological muddle. Let me conclude with a brief meta-philosophical remark concerning the susceptibility of recent philosophies of science to putative oppositions of this sort, one instance of which I have aimed to consider here.

The idea that scientific representations are things that contain information regarding their target systems, and more specifically, that this information can be analyzed neatly in terms of relations defined precisely in mathematics and logic is, arguably, a legacy of logical positivism. Some influential responses to positivism and subsequent work have preserved this legacy; consider, for example, the emphasis on model theory adopted by many proponents of the semantic view of theories, and the interest in formal relations between theories described by philosophers in the structuralist tradition. The retention of an aspect of the spirit of logical positivism is not by itself a bad thing, of course. But when the positivist edifice crumbled, many were tempted by a wholesale rejection of anything apparently tainted by it, and one of the central features of the program suffering this judgment was the positivist obsession with the possibility of scientific knowledge expressed in logico-mathematical terms. In the wake of the historical turn in the philosophy of science in the 1960s, whose central methodological slogan was to privilege descriptive accounts of scientific practice over and above the "prescriptive" analyses that had only recently been overthrown, it became only natural, perhaps, to conflate the idea of representations *as things* with the notion of representation *as practice*. Discussions of scientific representation are still swimming in the wake of that turn today.

With the benefit of hindsight, however, it is hopefully easier today to recognize that there is no great tension between these perspectives after all, at least not so far as the study of scientific representation is concerned. Representations in the sciences are intentional, informational tools, which scientists exploit so as to perform various activities in the process of representing the world. The different aspects of

[11] This point is often missed in the literature. Callender and Cohen (2006), for example, insist that informational relations are unnecessary for scientific representation, but their arguments for this thesis often treat 'scientific representation' as a synonym for 'model' (pp. 75–76, n. 6; p. 80), or even 'meaningful' (p. 72). One may take something to be a scientific model or meaningful, however, without thinking that it is also representational. The authors appear to acknowledge this later (p. 81, n. 11), observing that in cases where models turn out not to represent anything, one might say that they are not, in fact, representations after all.

representation targeted by the informational and functional approaches respectively, viz. the issue of what a representation is, as a means, and the issue of how representations are used, in the end, stand in need of further attention. Regarding the former, interesting questions remain about the nature of representational content, how it is grasped, and the precise manners in which it manifests similarity relations to target systems across the variety of ontological categories to which scientific representations belong.

Regarding the latter, we have only just scratched the surface of understanding how properties and relations *other* than those that make things scientific representations (such as different kinds of similarities) affect the pragmatics of representation. Given our cognitive architecture, for example, representations constructed in some ways allow interpretations and inferences regarding target systems more easily than others. Some kinds of information are more accessible to creatures with brains like ours, for example, when presented graphically as opposed to linguistically.[12] Further study of representational functions may also reveal the extent to which scientists rely on information that is not contained within the representations they mean to consider, such as information derived from background theories and auxiliary hypotheses, when engaged in practices such as interpretation and inference. These and other topics remain open to a *unified* study of the informational forms and cognitive functions of scientific representation.

Acknowledgements Parts of this paper were presented at the annual meetings of the Canadian Society for the History and Philosophy of Science and the British Society for the Philosophy of Science, and the Logos Workshop, 'Scientific Models: Semantics and Ontology', at the University of Barcelona. I am grateful to those audiences for stimulating discussions, and to Gabriele Contessa for editing this volume of *Synthese*. I would especially like to thank José Díez, Steven French, Roman Frigg, Ron Giere, and Kate Thomson-Jones for their excellent insight and help.

References

Bartels, A. (2006). Defending the structural concept of representation. *Theoria, 55*, 7–19.

Callender, C., & Cohen, J. (2006). There is no special problem about scientific representation. *Theoria, 55*, 67–85.

Chakravartty, A. (2004). Structuralism as a form of scientific realism. *International Studies in the Philosophy of Science, 18*, 151–171.

Chakravartty, A. (2007). *A metaphysics for scientific realism: Knowing the unobservable*. Cambridge: Cambridge University Press.

Contessa, G. (2007). Scientific representation, interpretation, and surrogative reasoning. *Philosophy of Science, 74*, 48–68.

Currie, G. (1995). *Image and mind: Film, philosophy and cognitive science*. Cambridge: Cambridge University Press.

da Costa, N. C. A., & French, S. (2003). *Science and partial truth: A unitary understanding of models and scientific reasoning*. Oxford: Oxford University Press.

Demopoulos, W., & Friedman, M. (1985). Critical notice: Bertrand Russell's *The Analysis of Matter*: Its historical context and contemporary interest. *Philosophy of Science, 52*, 621–639.

Elgin, C. Z. (2004). True enough. *Philosophical Issues, 14*, 113–131.

[12] Along the same lines, Elgin (2006) argues that model-building practices such as abstraction and idealization facilitate epistemic access to properties that would otherwise not be exemplified.

Elgin, C. Z. (2006). *Exemplification, idealization, and understanding*. Presented at Beyond Mimesis and Nominalism: Representation in Art and Science. Organized by the Courtauld Institute and London School of Economics.
French, S. (2003). A model-theoretic account of representation (or, I don't know much about art…but I know it involves isomorphism). *Philosophy of Science, 70*, 1472–1483.
Frigg, R. (2006). Scientific representation and the semantic view of theories. *Theoria, 50*, 49–65.
Giere, R. N. (1988). *Explaining science: A cognitive approach*. Chicago: University of Chicago Press.
Giere, R. N. (1999). Using models to represent reality. In L. Magnani, N. J. Nersessian, & P. Thagard (Eds.), *Model-based reasoning in scientific discovery* (pp. 41–57). Dordrecht: Kluwer.
Giere, R. N. (2004). How models are used to represent reality. *Philosophy of Science, 71*, 742–752.
Goodman, N. (1976). *Languages of art: An approach to a theory of symbols* (2nd ed.). Indianapolis: Hackett.
Hughes, R. I. G. (1997). Models and representation. *Philosophy of Science, 64*, S325–S336.
Ketland, J. (2004). Empirical adequacy and ramsification. *British Journal for the Philosophy of Science, 55*, 287–300.
Melia, J., & Saatsi, J. (2006). Ramseyfication and theoretical content. *British Journal for the Philosophy of Science, 57*, 561–585.
Mundy, B. (1986). On the general theory of meaningful representation. *Synthese, 67*, 391–437.
Scruton, R. (1983). *The aesthetic understanding: Essays in the philosophy of art and culture*. London: Methuen.
Suárez, M. (2003). Scientific representation: Against similarity and isomorphism. *International Studies in the Philosophy of Science, 17*, 225–244.
Suárez, M. (2004). An inferential conception of scientific representation. *Philosophy of Science, 71*, 767–779.
Suárez, M. & Solé, A. (2006). On the analogy between cognitive representation and truth. *Theoria, 55*, 39–48.
Swoyer, C. (1991). Structural representation and surrogative reasoning. *Synthese, 87*, 449–508.
van Fraassen, B. C. (1980). *The scientific image*. Oxford: Oxford University Press.
van Fraassen, B. C. (1989). *Laws and symmetry*. Oxford: Clarendon.
Walton, K. L. (1984). Transparent pictures: On the nature of photographic realism. *Critical Inquiry, 11*, 246–277.
Worrall, J., & Zahar, E. G. (2001). Ramseyfication and structural realism. In E. G. Zahar (Ed.), *Poincaré's philosophy: From conventionalism to phenomenology* (Appendix IV, pp. 236–251). Chicago: Open Court.