

Maps, Models, and Representation

James Nguyen¹ and Roman Frigg²

Forthcoming in Kareem Khalifa, Insa Lawer, and Elay Shech (eds.): *Scientific Understanding and Representation: Modeling in the Physical Sciences*. London: Routledge

Abstract

Maps are often invoked as a way to understanding scientific modelling: a model represents its target as a map represents its territory. However, without an account of how maps represent this analogy is suggestive at best. We reverse the direction of explanation and show that maps represent like models. Utilising the DEKI account of representation, we provide an account of cartographic representation. This shows that maps and models indeed represent in the same manner, and it provides insight into two areas of philosophical inquiry, namely the nature of representational accuracy and the purpose relativity and historical situatedness of representations.

1. The Map Analogy

Scientists from across different fields construct, investigate, and draw conclusions from scientific models. Such models form the basis of much of our scientific knowledge. An immediate philosophical question then is: how do models perform the function that they do, the function of providing information about the parts and aspects of the world (their target systems) we are ultimately interested in? We call a representation that licences inferences about its target in this way an *epistemic representation*, and we have argued that models perform their function by being epistemic representations of their targets. Of course this just pushes the question

¹ james.nguyen@sas.ac.uk. Institute of Philosophy, School of Advanced Study, University of London and Centre for Philosophy of Natural and Social Science, London School of Economics and Political Science.

² r.p.frigg@lse.ac.uk. Department of Philosophy, Logic and Scientific Method, London School of Economics and Political Science.

back a level: what does it mean for a model to be an epistemic representation of a target?

In the philosophy of science it has become popular to draw on the idea that cartographic *maps* provide the appropriate analogy for understanding how scientific models, and indeed science more generally, work. In his 1994 presidential address to the Philosophy of Science Association, Giere argues that if we are to understand scientific representation, he “would suggest beginning with maps, e.g., a standard road map. Maps have many of the representational features we need for understanding how scientists represent the world” (1994, 11). Kitcher devotes Chapter 5 of his *Science, Truth, and Democracy* to the analogy and writes: “I want to clarify the picture of the sciences I have been developing by looking at the core field, the academically rather unfashionable discipline of cartography” (2001, 55). Winther notes that both a scientific theory and a scientific model can be seen as “a map of the world” (2020, 29 and 46). And a few decades earlier Toulmin noted that “the problems of method facing the physicist and the cartographer are logically similar in important respects, and so are the techniques of representation they employ to deal with them” (1953, 105).³

This suggests that both models and maps are epistemic representations and that models represent their targets as maps represent their territories. We call this the *models-as-maps analogy*. This analogy suggests that philosophers of science interested in representation can turn to cartography to provide them with a worked-out account of representation, ready to be used in the sciences. Cartographers seem rather unconvinced about the viability of this purported lateral knowledge transfer. In their seminal *The Nature of Maps*, cartographers Robinson and Petchenik note with dismay that “while some cartographers and geographers have cast about for

³ Similar observations have also been made by Boesch’s (2019), Bolisnka’s (2013), Contessa’s (2007), Giere’s (1999), and Sismondo and Chrisman’s (2001). There is an interesting question whether a similar relation holds between maps and linguistic representations. This is, unfortunately, beyond the scope of this paper. For recent discussion of this question see Aguilera’s (2021).

things to which they can liken the map [...] scholars in other fields tend to use the map as the fundamental analogy” (1976, 2), and they add “maps clearly are involved in communication, and it would seem much could be learned from other analyses of other types of communication” (*ibid.*, 3). This suggests that there’s no clear account of how maps represent and that the direction of knowledge transfer ought to be from other domains to cartography, rather than vice versa.

If nothing else, this diversity of opinion highlights that there is a question about the work that the models-as-maps analogy does, and about how far the analogy reaches. We suggest that this question is productively addressed by distinguishing two levels at which the analogy can be seen as operating. At the first level, the analogy is seen as providing an account of how models represent. At the second level, the analogy is taken to illuminate features of representation, in particular the nature of accuracy and the purpose relativity and historical situatedness of representations, which have wide-ranging philosophical implications.

Our claim is that the analogy fails at the first level. In line with what Robinson and Petchenik suggest, we turn the models-as-maps analogy on its head. Rather than attempting to use maps to learn about scientific representation, we explore how our preferred account of scientific representation (the “DEKI account”) can be used to help us understand how maps work. This is not just an exercise in the philosophy of cartography; it also further develops the DEKI account demonstrating how its conditions work. By design, these conditions are skeletal in the sense that they need to be filled in, or concretised, in any particular instance of epistemic representation. Thus, maps provide an illustration of one way in which this can be done. By contrast, we believe that the analogy works productively at the second level. By understanding how maps represent, we can deepen our understanding of how representations function more generally and draw interesting conclusions about some of their features.

We begin by discussing a concrete example of using a map and illustrate the mistakes that are made if the map is read naïvely. This, we submit, shows that maps

presuppose rather than provide an account of representation (Section 2). This leaves open what account that is. We submit that the DEKI account fits the bill. We introduce the account (Section 3) and show how it works in the case of maps (Section 4). We then revisit claims that have been built on the models-as-maps analogy and assess their validity (Section 5).

2. Reading Maps Naïvely

There is a naïve view that maps are somehow “natural” representations that show their territories as they “really are”. While none of the authors mentioned in the previous section holds such a view, showing what’s wrong with it leads the way to a better understanding of maps.

Consider the following imaginary scenario. Like other European countries, there have been strong separatists movements in Sweden. Eventually a referendum is called, and the proposal to split the country is successful. Specifically, the decision is to create the two independent states of North Sweden and South Sweden. All parties agree that the border should be drawn on a purely geographic basis by dividing the country in the middle along the north-south axis. Asbjörn is the government minister tasked with drawing the border. To do so he reaches for his map of Europe (shown in Figure 1) and sets out to determine Sweden’s north-south midpoint.



Figure 1: Map of Europe

Inspecting the patch of the map marked “Sweden”, Asbjörn finds that the point marked “Trekörset”, with Euclidean coordinates (x_1, y_1) on the map’s surface, is the furthest to the top and hence represents the northernmost point of Sweden. Similarly, he finds that the point marked “Smygehuk”, with the Euclidean coordinate (x_2, y_2) on the map’s surface, is the closest to the bottom and hence represents the southernmost point in Sweden. He then determines that the border should be given by the horizontal line through $\frac{1}{2}(y_1 + y_2)$, which is the solid line seen in Figure 2. This seems like a natural way of completing his task of dividing the country half way on the north-south axis. Points to the top of the map represent locations to the north; points to the south represent locations to the south; and distances on the map correspond to distances in the world. So, surely the midpoint between the northern

and the southern tip must be in the middle between the points on the map that represent the northmost and southmost locations.

Natural as this may seem, Asbjörn's procedure draws the border in the wrong place. His technique for dividing Sweden would have been correct with a map whose projection preserves the even spacing between east-west parallels. But the map in Figure 1 does not have this feature. It has been made with the Mercator projection, which preserves bearings (i.e. angles) but distorts distances.⁴ In particular, as we approach the top of the map, the same distances on the top-bottom axis represent ever smaller distances on the north-south axis on the globe.

How bad is Asbjörn's mistake? Are we quibbling about epsilons? To answer this question, let us consider another minister, Berit, who knows about the Mercator projection MP :

$$x = sR(\lambda - \lambda_0), \quad y = sR \ln \tan\left(\frac{\pi}{4} + \frac{\varphi}{2}\right), \quad (1.1)$$

and its inverse MP^{-1} :

$$\lambda = \frac{x}{sR} + \lambda_0, \quad \varphi = 2 \tan^{-1}\left(e^{\frac{y}{sR}}\right) - \frac{1}{2}\pi, \quad (1.2)$$

where λ is longitude (in radians), λ_0 is a central meridian (in radians), which in our case is Greenwich and so $\lambda_0 = 0$, φ is latitude (in radians), and the coordinate system on the globe is chosen such that an increase of λ involves a shift east, and an increase of φ involves a shift north, R is the radius of the globe, and s is a scale factor.

⁴ For a survey of projections see Monmonier's (1991, Ch. 2); for an in-depth discussion see Pearson's (1990). For a mathematical definition of the Mercator projection see Pearson's (1990, Ch 5.VII), and for a discussion of its history Winther's (2020, Ch. 4).

Berit takes the coordinates of the northernmost and the southernmost points, (x_1, y_1) and (x_2, y_2) respectively, and then uses the inverse projection to determine the radian-valued coordinates (λ_1, φ_1) of the northernmost point and (λ_2, φ_2) of the southernmost point. The concrete values she finds, converted to degrees, are 69.06°N 20.55°E and 55.34°N 13.36°E for Treriksröset and Smygehuk respectively. This means that the midpoint on the north-south axis between them (on the globe) lies at $\frac{1}{2}(\varphi_1 + \varphi_2) = 62.20^\circ\text{N}$. This is where the border should be according to the agreement reached between the parties.

Feeding this value into the projection for the y coordinate on the map, Berit determines that the mid-point is the dotted line shown in Figure 2. As we can see, these lines do not coincide with one another. In fact, the solid line (which goes through the midpoint on the map itself) corresponds to a latitude of roughly 63°N. Calculating the distance between them shows that Asbjörn's naïve interpretation of the map, which failed to take into account that the relevant projection is the Mercator projection, resulted in a border roughly 89km too far to the north! The error is significant, even for a country that is roughly 1600km from north to south.



Figure 2: Europe according to the Mercator-projection. The solid line is the midpoint on the map, the dotted line is the midpoint as calculated using the details of the projection

The problem is that the map of Europe doesn't warn Asbjörn that his attempt to divide a territory in the middle by dividing a distance on the map in the middle will result in a grave error. The map's projection does not, as it were, jump off the page when you look at it. *Per se*, the map is piece of paper with certain shapes drawn on it, and you have to know what projection has been used to produce the map in order to use it correctly.

This is not just a toy example cooked up to illustrate our point. Sismondo and Chrisman report that errors of this kind are common:

At least half of a sample of 137 international maritime boundaries appear to have been plotted as equidistant lines on the chart without accounting for differences in scale [...]. Even in the

relatively equatorial situation of Australia and Indonesia, the agreement specifies positions that are 4 nautical miles (7.4 km) south of the actual line of equidistance [...]. This amount of error is large enough for a sizable oil platform or two. (2001, S42).

The point generalizes: the way in which a map represents its terrain cannot be read off the map directly. You could interpret it naïvely, and uniformly scale every measurement on the map to a measurement on the terrain. Alternatively, if you are aware of the distortion introduced by the projection, you can take this into account in your inferences from the map to the terrain. Whilst the former may look more “natural” than the latter, you’d be well served to avoid such an interpretation if you want the results of your map-to-terrain reasoning to be accurate. In order to employ the latter kind of interpretation you need to know the details of the projection used to create the map, and the conventions associated with its use.

This point has been recognised by philosophers writing on maps. Giere, for instance, notes: “Maps require a large background of human convention for their production and use. Without such they are no more than lines on paper” (1994, 11).⁵ This is right as far as it goes, but it leaves important questions open. What are the conventions we must be aware of and in what way does the mode of production of a map matter? Answers to these questions come from an account of representation (and will clearly also involve an investigation into the conventions and practices associated with particular maps, which will in turn be the subjects of such an account) rather than from the map itself (whatever that may mean), which undermines the utility of the models-as maps analogy.

Indeed, as noted in Section 1, we’re going to turn this point on its head and show that our DEKI account, originally offered as an account of how models represent, provides the required understanding of how maps represent.

⁵ See also Sismondo and Chrisman’s (2001, S42), Kitcher’s (2001, 56-57), and Toulmin’s (1953, 108).

3. Analysing Maps with DEKI

The DEKI account is designed to answer the question: in virtue of what is a model (or a map) M an epistemic representation of a target (or territory) T ? It understands M as the ordered pair of an object X and an interpretation I , hence $M = \langle X, I \rangle$, and it postulates that M is an epistemic representation of T if and only if:

- (i) M denotes T (and parts of M may denote parts of T);
- (ii) M exemplifies features F_1, \dots, F_n ;
- (iii) M comes with key K , associates F_1, \dots, F_n with a collection of features Q_1, \dots, Q_m ; and
- (iv) M imputes Q_1, \dots, Q_m to T .

These four elements – Denotation, Exemplification, Keying-up, and Imputation – give the account its name.⁶ We have presented them in detail elsewhere (Frigg and Nguyen (2018; 2020, Chs. 8 and 9)). Our purpose here is to briefly summarise them, and then discuss how they play out in the case of maps.

Let us begin with the internal structure of M . At a basic level, a model is an object X : a system of pipes, two imaginary perfect sphere, or an oval block of wood. *Per se* these things are just objects, like the tables and chairs in our offices. What turns a “mere” object into a model is that it is endowed with an *interpretation I*. A system of pipes becomes a model when the flow of water through pipes is interpreted as the flow of money through an economy (the Phillips-Newlyn model); the two spheres become a model of hydrogen if they are interpreted as the electron and the proton (the Bohr model), and the oval block becomes model when it is interpreted as ship.⁷

⁶ One thing to note here is that we are not assuming that epistemic representation is, in some sense, “mind-independent”, or “naturalisable” (we are grateful to Elay Shech for encouraging us to be explicit about this). As such, the conditions we use to explicate it already include some intentional notions. For more on this see Frigg and Nguyen’s (2020, 39-40).

⁷ For a discussion of these models in the context of the DEKI account see, Frigg and Nguyen’s (2020), (2016), and Nguyen and Frigg’s (2022), respectively.

Hence, a model is the pair of an object X with an interpretation I : $M = \langle X, I \rangle$. Crucially, this does not presuppose that models have targets. Interpretations presuppose a conceptual scheme in terms of which the interpretation is phrased, but they don't presuppose that anything real is singled out. One can interpret a flow of water as the movement of elves, the block of wood as a UFO, and the two spheres as Vulcan and one of its moons. This will still turn the objects into models, albeit targetless models.⁸

This carries over to maps. Per se, a map is a piece of paper (or, increasingly, an image on a computer screen) exhibiting certain lines and shapes. The piece of paper, X , becomes a map only once it's endowed with an interpretation I according to which surfaces enclosed by solid lines are interpreted as countries; the lines themselves as borders; the blue surfaces as water; and so on. Without such an interpretation, a map is merely a coloured piece of paper. Like models, maps need not be maps of a real territory. Maps of Atlantis, the world according to Game of Thrones, Winnie-the-Pooh's Hundred Acre Wood are maps, but not ones that represent a real territory.

If representation is not built into the notion of M , what does it take for an M to be a representation of a T ? In the first instance, we want to know what makes M be about T . Condition (i) addresses this point by appealing to the notion of *denotation*.⁹ Models and maps can denote their targets just as proper names denote their bearers, predicates denote objects in their extension, and photographs denote their subjects. The Phillips-Newlyn model denotes the Guatemalan economy, the Bohr model denotes hydrogen, and the ship model denotes The Queen Mary. In addition to M as a whole denoting T as a whole, parts of M can denote parts of T . The flow on the right in the Phillips-Newlyn model denotes foreign trade; the small sphere in

⁸ In Goodman's (1976) terms, models are Z -representations, where the Z is given by the interpretation.

⁹ Thus, when we say " M is about T " we mean it in the minimal sense that M denotes T . And following, e.g. Goodman (1976), by "denotation" we mean the bare relation between a symbol and what is symbolised, without invoking the "meaning" or "descriptive content" which may (or may not) be associated with the symbol. See Salis, Frigg, and Nguyen's (2020) for further discussion of this point. Thanks to Elay Shech for encouraging us to be explicit about this.

Bohr's model denotes an electron, and so on. The same is true of maps. The map in Figure 1 denotes Europe. In addition, every point on the map denotes a point on the globe, namely the point specified by the inverse projection, given in Equation (1.2). Some of these points are given names. If, say, a point is labelled "Stockholm" this means that the point denotes the place on the globe where the city of Stockholm is located.¹⁰

Denotation is necessary but insufficient for epistemic representation. It is necessary because it establishes the bare sense in which M is about T , and parts of M are about parts of T . It is insufficient because denotation alone is too weak to ground epistemic representation: the fact that M denotes T doesn't enable us to use M to generate claims about T (one cannot draw inferences about London from investigating the features of the syntactic object "London").

Explaining how a representation can function epistemically proceeds in several steps. The first step, condition (ii), involves the concept of *exemplification*. Exemplification is a mode of reference that occurs when an object refers to a feature it instantiates. This is established relative to a context. We can define it as follows: M exemplifies a certain feature F in a certain context iff M instantiates F and the context highlights F , where a feature is highlighted if it is identified in the context as relevant and epistemically accessible to users of M . An item that exemplifies a feature is an *exemplar*. Standard examples include samples (the beer you try at the brewery exemplifies its flavour) and swatches (the colour swatch in the paint shop exemplifies its colour).¹¹ Exemplification is selective; the colour swatch doesn't exemplify

¹⁰ Although note that Stockholm is of course a region, rather than a point, on the globe, so points on the map denote locations in Stockholm, and these point-wise denotation relations then associate a region on the map with Stockholm, considered as a region on the globe. Thanks to Mark Risjord and Jared Millson for highlighting this.

¹¹ Note that exemplification is a semantic, rather than epistemological notion. If, for example, the beer sample was from the bottom of the bottle and thus exemplified consisting of a large amount of yeast slurry, this doesn't entail that all of the beer in the bottle will exhibit this feature. Exemplification may

rectangularity and even though it instantiates it. Only selected features are exemplified, and which features are selected depends on the context.

Exemplars provide epistemic access to the features they exemplify. This is because they instantiate the features they exemplify in a way that makes them salient, which in part depends on the context in which they are embedded. The paint chip makes a particular shade of blue salient and thereby acquaints those using the chip with that shade of blue because in that context the chip's colour is salient and accessible to an observer. Epistemic representations exemplify certain features. The Phillips-Newlyn model exemplifies a certain level of taxation, the Bohr model exemplifies certain energy levels, and the ship model exemplifies a certain resistance when moving through water. As these examples indicate, instantiation is here understood as instantiation under interpretation I . The Phillips-Newlyn model is system of water-pipes interpreted in terms of economic properties.¹² Under this interpretation, the model can instantiate, and hence exemplify, economic properties like having a low-tax fiscal regime. We're not committed to restricting instantiated, and exemplified, properties to properties possessed by X as a "bare" object, which have the absurd consequence that the Phillips-Newlyn model could only exemplify water-and-pipe properties. The point is pertinent to maps. The map in Figure 1, under the standard interpretation of a map, does not exemplify lines and colours; it exemplifies there being borders between countries and landmasses having a coast lines; and it exemplifies certain points being at a certain distance from each other. In general, then, M will exemplify certain features F_1, \dots, F_n .

That M denotes T and exemplifies features F_1, \dots, F_n is still not sufficient to make M an epistemic representation of T . To get to that point, two further steps are needed. For M to represent T as being such and such, a user of the representation has to *impute* features to T (by this we just mean the user ascribes features to the target). But

allow us to successfully infer features of the sampled from features of the sample, but it doesn't guarantee it. We are grateful to Elay Shech for encouraging us to be explicit about this.

¹² One might worry that non-concrete models like the Bohr model cannot strictly speaking instantiate features. This worry is addressed in Frigg and Nguyen's (2020, Ch. 9).

which features are these? A natural suggestion here is that these features are simply the ones that are exemplified by M in the representational context. Indeed this is how some kinds of epistemic representation work: taken as model of your new front door, the colour swatch in the paint shop exemplifies a certain shade of red and imputes exactly this shade of red to your door. But at least in scientific contexts, this is the exception rather than the rule. It's rarely the case that epistemic representations represent their targets as having *precisely* the features that the former exemplify. Someone who knows how to use a mechanics model from physics won't conclude that a real skier will have the particular trajectory that it follows in the model that assumes air resistance and friction to be absent: they know that the skier is subject to friction and air resistance, and they can take this into account when they use the model to reason about the target.

This leads us to the notion of a *key*. In the abstract, a key can be thought of as a function: it takes as inputs the exemplified features F_1, \dots, F_n , and it delivers as outputs the features Q_1, \dots, Q_m that the user of the representation should be willing to impute (or ascribe) to the target system.¹³ As stated this is an abstract notion, and this is by design. We submit that keys are one of the main locations that encode the disciplinary conventions associated with the use of epistemic representations. In the case of modelling, these keys are the sorts of things that students learn when they learn to use their models. In some cases these keys might involve weakening the

¹³ Our use of the term "key" in the DEKI account is inspired by the sorts of representations we're investigating here: maps! However, it should be noted that there may be a difference between the "key" as used in DEKI and the explicit "key", or "legend" (we'll use the latter term to disambiguate between the two), that is literally written down on a map. Many simple maps (such as city maps for tourists) may not come with an explicit legend, and thus the keys that are used to interpret them are implicit in the conventions and practices associated with the maps (we are grateful to Mark Risjord and Jared Millson for pointing this out). However, it's also worth noting that many maps do in fact explicitly contain the sort of information that keys (in the sense of DEKI) require in their legends; for example Ordnance Survey maps designed to guide walkers around regions in the UK are explicit that they are constructed from the Transverse Mercator Projection, and include information relevant for the key, such as the distinction between Magnetic North, True North (on the globe) and Grid North (on the map).

isolated exemplified feature of the model to the claim that the target only has a disposition to behave that way (Nguyen 2020); in others they may result in certain limit-based reasoning: the model feature might be related to the imputed feature via taking a limit (Nguyen and Frigg 2020). One of the important upshots of the DEKI account is the demand that in any particular instance of epistemic representation, we have to understand the key that accompanies it.¹⁴

Like models, maps exemplify certain features. For instance, it exemplifies there being land-boundary between the countries labelled “Norway” and “Sweden”; it exemplifies Norway being to the left of Sweden; it exemplifies the dot labelled “Stockholm” being higher up than as the dot labelled “Gothenburg”; and it exemplifies the dot labelled “Stockholm” and the dot labelled “Umea” being 20 centimetres from each other. We noted that exemplification is selective, and we can see an example of this here. While blue is exemplified (indicating water), the other colours are not: countries are individuated by solid lines indicating borders, and their colour is a mere convenience that does not contribute to the map’s content. Similarly, the texture of the paper on which the map is printed (or the make of the screen on which it is displayed), the typeface of the letters used to label points and areas, and the fact that it’s been printed in Germany are all properties of the map, but they are not exemplified.

¹⁴ In the scientific context, the key may be implicit in the practice of the discipline in which a model is embedded. But it is worth noting that scientists are free to experiment with different keys, and that their choice of key may be the result of an investigative back-and-forth between a model and its target. Under this understanding of keys, whilst some may be more accurate than others, from a semantic point of view, scientists are free to choose them as they see fit. Shech (2015) uses the term “code” to describe a related notion, but as he notes, it has a dual meaning: “[o]n the one hand, a code, understood as a key, legend or guide, is needed in order to make use of a representational vehicle for surrogative reasoning. On the other hand, understood as a cryptogram or cipher, the code of a representation is not always known and so it must be “deciphered,” so to speak” (p. 3469). Our notion of key thus corresponds to the former meaning. For related discussion see footnote 22, our response to Millson and Risjord in this volume, and Frigg and Nguyen’s (2020, Ch. 8).

We should not assume that the properties exemplified by a map are imputed unaltered to the territory represented. Some are; some are not. A place being to the left of another place, or a place being further up than another place, are meaningless when imputed to the world because there is no unique left-right or up-down on the globe, and no one would take the map to say that the cities of Stockholm and Umea are 20cm from each other. Like with models, we need a key to tell us how the transition from the map to the territory works.

The bedrock of the key are the part-part denotations given by the inverse transformation in Equation (1.2), which specifies for each point on the map which point on the globe it stands for.¹⁵ The first element of the key is the rule that when a certain point on the map is singled out by the map's interpretation as having a certain characteristic – being on a border, being on a coast line, being a city centre – then this characteristic is imputed to the globe-point that the map-point stands for. If, say, the interpretation of the map specifies that solid black lines are borders, and a certain point lies on a black line, then the map imputes to the point it denotes the property of being on a border.

The projection also helps us keying up properties like *to the left of* and *higher up than*. *To the left of* corresponds to lower values of x , and the inverse projection for the x -axis, $\lambda = \frac{x}{R} + \lambda_0$, tells us that higher values of x correspond to points further east. This means that the map-property *to the left of* is keyed up with the globe-

¹⁵ We emphasise here that Equation (1.2) associates *points* on the map with *points* on the globe. It doesn't tell us anything about how *features* on the map represent *features* on the globe. To illustrate this consider the fact that "roads" on maps (i.e. coloured lines) are typically wider than they "should" be, given the maps' scales. According to our discussion, this is because the points on the edge of the coloured line (i.e. the points that seem to represent the road as wider than it is) are associated with points on the globe that aren't in fact paved. But this is the result of the colouring on the map, not the point-to-point denotation relations given by Equation (1.2), and moreover the way that map keys associate colours on a map's surface with a road typically don't require that a map user infer that roads are wider than they in fact are (the *width* of the coloured line is neither exemplified or keyed-up). Thanks to Mark Risjord and Jared Millson for encouraging us to be explicit about this.

property *to the west of*. While this is a frequently used convention, it is in no way necessitated by the situation. The mapmaker could have used a projection with the inverse $\lambda = -\frac{x}{R} + \lambda_0$, in which case *to the left of* would be keyed up with *to the east of*. Likewise, the fact the *higher up than* corresponds to higher values of y , together with the fact that the inverse projection for the y -axis, $\varphi = 2 \tan^{-1}(e^{\frac{y}{R}}) - \frac{1}{2}\pi$ is a monotonic function of y , implies that the map-property *higher than* is keyed up with the globe-property *to the north of*.

Things get more involved when we turn to distances. In the idiom of DEKI, Asbjörn adopted a key according which map-distances scale linearly with globe-distances, i.e. a key according to which a map distance d_m is keyed up with the globe distance $d_g = c d_m$, for some scale factor c . As we have seen, this key returns wrong results. In fact, there is no scaling factor for map distances at all! Map distances in Mercator maps are not keyed up with globe-distances. The correct key says that the distance between the map-points (x_1, y_1) and (x_2, y_2) is given by $G[MP^{-1}(x_1, y_1), MP^{-1}(x_2, y_2)]$, where G is the great-circle-distance on a sphere (i.e. the length of line-segment between two points on a great circle drawn through these points).¹⁶

Consideration of the same kind are also needed to determine the surface of an area. We can't simply measure the surface of part of the map and expect it to scale linearly. We will have to project the boundaries of the relevant territory back onto the globe with the inverse transformation, and then determine the measure of the relevant surface on the sphere.

This shows that the map requires a key for its use, and that this key is more than just a trivial identity which says something like "whatever is true in the map is true on the globe". And the key is only part of what is needed to use the map. As we have seen,

¹⁶ For discussion of the great circle distance see Pearson's (1990, Ch. 3).

all elements of the DEKI account do essential work in explaining how maps work: the interpretation turns a pattern of lines and shapes into a territory-representation; denotation turns this territory representation into representation of a particular territory; exemplification singles out relevant features; the key transforms these into other features; and act of imputation says that the territory has these other features.

4. Philosophical Lessons

In Section 1 we have seen that the models-as-maps analogy can be seen as operating at two levels: at the first level, it can be seen as providing an account of how models represent while, at the second level, it is taken to illuminate features of representation, in particular the nature of accuracy, the purpose relativity and historically situatedness of representations, and the possibility of total science. We now assess how well the analogy fares with these.

A time-honoured position in the literature on scientific representation appeals to the notion of similarity: a model M accurately represents its target T in virtue of M and T being similar to one another in the appropriate respects, and to the appropriate degree.¹⁷ Given the popularity of the models-as-maps analogy, it is not surprising to see this being motivated by the idea that maps function in the same way. For example, Giere, who defends a similarity-based account of scientific representation, presents a map of Pavia and argues: “How does this map represent Pavia? The answer is: by being spatially similar to aspects of Pavia” (1999, 45), and then, to illustrate the context sensitivity of what is meant by similarity, further discusses contexts such as using a map of the London Underground to be such where “the important similarities are those between these topological features of the map and of the whole metro system” (*ibid.*, 46).

¹⁷ See Frigg and Nguyen’s (2020) for further discussion of similarly accounts of representation in general; see Winther’s (2020, Ch. 5) for discussion of similarity in the context of maps.

We submit that an appeal to similarity is either misleading or empty, and that the models-as-maps analogy does not do any productive work at the first level (at least in the sense that the first level analogy doesn't motivate an account of representation that appeals to similarity as common to both maps and models). It would seem to be an obvious consequence of the notion that maps are similar to their territories that map-distances are proportional to territory-distances. While this is correct in a standard city map, it leads, as we have seen in Section 2, to significant error in maps like the ones shown in Figure 1. One might now turn around and say that this is not how similarity should be understood: the map being similar to its territory here simply means that the transformation in Equation (1.1) holds.¹⁸ There is a question whether this equation, or indeed other projections, can meaningfully be regarded as a kind of similarity. The vagueness of the notion of similarity makes it difficult to say. Let's set this issue aside. The more pressing problem is that this way of approaching the issue makes similarity otiose. One first has to know the projection and all the conventions used, and only when everything has been spelled out one can turn around and say "see, they are similar".¹⁹ Thus understood, similarity does not work and it becomes a success term that gets attached to a finished product when things work out as envisaged.²⁰ It would seem to more productive to think about maps in

¹⁸ Similar points can be made about the structuralist conception of representation because the Mercator projection does not preserve certain structural features (e.g. distance ratios). This way of thinking about representation is discussed in Frigg and Nguyen's (2020, Ch. 4).

¹⁹ We are not claiming here that the map user needs an explicit philosophical account of representation in order to use the map: rather it's that the map user should adopt the conventions associated with the map when using it, and that these conventions, which don't have to be understood in terms of similarity, are the subject of our account of representation.

²⁰ An advocate of the idea that it is similarity (structural or otherwise) that establishes representation could argue that it's the "keyed-up" map (i.e. the map with the key applied to it) that is supposed to be similar to the target, not the "bare" map itself (thanks to Mark Risjord and Jared Millson for suggesting this possibility). But this would pull the rug from under such an account: the crucial move would be the application of the key, and then claiming that the results of such an application are "similar" to the target is just another way of describing how the outputs of the DEKI account of representation should be compared to the target.

terms of DEKI, which is explicit that the account treats denotation and keys as blanks to be filled on every occasion, and then filling them with the appropriate projection.

In contrast with the first level, the analogy is largely correct at the second level. In the remainder of this section we discuss some pertinent second level points and explain how they bear out in an analysis of maps based on DEKI.

Let us begin with accuracy. There is temptation to say that a map on which Iceland appears larger than Romania, where in reality Romania has more than twice the surface area of Iceland, is inaccurate. Kitcher, rightly, protests that such verdict would be “foolish” because “[a]ssociated with any map there are conventions that determine which aspects of the visual image are to be taken seriously” (2001, 56). As we have seen in the previous section, maps come with a key (which can be either implicit or explicit, cf. footnote 13) and disregarding the key leads to wrong results. Calling a map with these features inaccurate relies on a naïve reading of the map, and we have seen that such reading is illegitimate. The accuracy of a map has to be judged relative to a key, not relative to visual appearance: a map is accurate when the territory has the features that the key outputs. The same holds true of models, which we should not expect to be “like” their targets in some pre-theoretical and unreflected sense: a model is accurate if attributing features provided by the key, Q_1 , ..., Q_m , results in true statements. This does not involve “looking alike”, or being similar.²¹

Maps are made for a particular purpose, and there is no such thing as a map that is good for everything. If you’re hiking the Scottish highlands, you’d be well advised to use a map that displays the topography of the terrain depicted and the paths that thread up and down the munros. In contrast, if you’re on a scenic drive from Inverness to Fort William, you’d be better off with a map that marks roads and speed limits. As Kitcher puts it: “we understand how maps designed for different purposes pick out different entities within a region or depict those entities rather differently”

²¹ For a detailed discussion of this point see Frigg and Nguyen’s (2021).

(2001, 56). The same is true of scientific models. As Morrison (2011) and Massimi (2018) have shown, scientists often produce different models of the same target, where the models pick out different aspects and features of the target depending on what aims the scientists pursue. Kitcher generalises this point when he notes that “the aim of the sciences is to address the issues that are significant for people at a particular stage in the evolution of human culture” (*ibid.*, 59). DEKI is compatible with this idea. It sees scientists as having complete freedom both in the choice of model-objects and in the choice of keys, and these choices can be seen as inevitably historically contingent and relative to our aims and purposes.²²

If correct, this has important consequences for the project of science as whole. If all representations that science produces are purpose relative and historically situated, then there is no such thing as the perfect map. Borges (1998) reminds us in his notorious story about cartography that a perfect map would have to coincide with its territory point by point, resulting in a map of the Empire that would be as large as the Empire itself. Such a map would be useless, and soon abandoned. Maps are selective in what they represent. There are two ways of thinking about the scientific project corresponding to Borges’ map. First, one might hope that we will eventually discover the “final theory” in fundamental physics, and such a theory would, in principle at least, suffice to provide a complete representation of the fundamental structure of

²² Of course, as the example of Asbjörn’s misuse of the Mercator projection shows, once the choices are made, they do constrain future uses of maps and models (at least if one aims at using them to generate useful knowledge about their targets). Conventions are freely chosen, but once chosen they are relatively fixed (thanks to Mark Risjord and Jared Millson for encouraging us to clarify this). But it should also be noted that there are occasions in the history of science where models can be fruitfully used by adopting new conventions, that weren’t there at their inception. These include cases of model transfer, where scientists use models originally designed to represent some target system (e.g. a two body celestial system) to represent an alternative kind of target system (e.g. the atomic nucleus as per Bohr’s model), as well as cases where model behaviour is interpreted in novel ways (e.g. Dirac’s observation that negative energy solutions to the Dirac equation, once understood as mathematical artefacts, could be interpreted in terms of positrons, cf. Bueno and Colyvan’s (2011, 365-365)).

the world.²³ Second, and more boldly, one might hope that we will eventually discover the “complete theory”, which contains not only a complete representation of the fundamental structure of the world, but everything else too, from molecular bonds to social practices to political systems and everything in-between.

The realisation that every representation, be it geographical or scientific, is purpose relative and historically situated should at least cast doubt on latter project. Giere submits that “[t]here is no such thing as a universal map” (1994, 11). Kitcher goes into more detail in his discussion of the “ideal atlas” (2001, 60). It couldn’t be a single map, since, as noted, nothing could perform this role except for the terrain itself. But perhaps it could be a collection of “fundamental maps” (corresponding to a “final theory”), from which “all spatial information can be generated, and that they collectively provide a unified presentation of the wide diversity of kinds of knowledge drawn from our actual ventures in cartography (and, presumably, projects we might have undertaken)” (*ibid.*) Kitcher argues that a brief glance at the vast diversity of maps produced in human history should make us immediately sceptical about the possibility of such a compendium, notwithstanding the fact that an ideal atlas would also have to encode the information about projects and investigations we haven’t, but might have, embarked on. Thus, he concludes, the models-as-maps analogy should force us to reconsider idea of a complete theory, and, along with it, the idea that our scientific theories and models are converging on it. We agree.

Even if it’s granted that science doesn’t aspire to a complete theory, some may hold onto the possibility of a final theory. Just because we may not be able to represent every fact in the world, doesn’t preclude the possibility of us, eventually at least, representing the fundamental facts on which the others (making some strongly reductive assumptions), ultimately depend. And perhaps the way in which we’ll represent these facts, won’t turn, in any philosophically significant way, on contextual aspects of our representations. We won’t take sides here, beyond noting that even if the hope in a final theory is still alive, this doesn’t tell against the idea

²³ For a vivid account of the vision of a final theory see Weinberg’s (1993).

that we need to construct partial and purpose relative representations to help us explore non-fundamental domains. As Weinberg (1993, 18), an enthusiastic adherent to the goal of a final theory, notes:

“Of course a final theory would not end scientific research, not even pure scientific research, nor even pure research in physics. Wonderful phenomena, from turbulence to thought, will still need explanation whatever final theory is discovered. The discover of a final theory in physics will not necessarily even help very much in making progress in understanding these phenomena (thought it may with some)”.

So, the models-as-maps analogy pours cold water on the dream of a complete theory, even if a final theory remains a live option. Our (non-fundamental) representations will always be partial, and have purpose relative and historically situated aspects. But as we have seen, this doesn't mean they have to be inaccurate, and there is a clear sense in which some interpretations (like Berit's) are to be favoured over others (Asbjörn's). Giving up on a complete theory doesn't mean anything goes.

Acknowledgements

We are extremely grateful to Jared Millson, Mark Risjord, and Elay Shech for comments on a previous draft of this chapter, and for a stimulating dialogue. We also thank Kareem Khalifa, Insa Lawler, and Elay Shech for inviting us to participate in this project.

References

- Aguilera, M. (2021). Heterogeneous inferences with maps. *Synthese, Online First* <https://doi.org/10.1007/s11229-020-02957-w>.
- Boesch, B. (2019). Scientific representation and dissimilarity. *Synthese, Online First*. DOI:10.1007/s11229-019-02417-0.

- Bolinska, A. (2013). Epistemic representation, informativeness and the aim of faithful representation. *Synthese*, 190(2), 219-234.
- Borges, J. L. (1998). On Exactitude in Science. In *Collected Fictions* (pp. 325). New York: Penguin.
- Bueno, O., & Colyvan, M. (2011). An inferential conception of the application of mathematics. *Nous*, 45(2), 345-374.
- Contessa, G. (2007). Scientific representation, interpretation, and surrogate reasoning. *Philosophy of Science*, 74(1), 48–68.
- Frigg, R., & Nguyen, J. (2016). The Fiction View of Models Reloaded. *The Monist*, 99, 225-242.
- Frigg, R., & Nguyen, J. (2018). The turn of the valve: representing with material models. *European Journal for Philosophy of Science*, 8(2), 205-224.
- Frigg, R., & Nguyen, J. (2020). *Modelling Nature. An Opinionated Introduction to Scientific Representation*. Berlin and New York: Springer.
- Frigg, R., & Nguyen, J. (2021). Mirrors without warnings. *Synthese*, 198, 2427-2447.
- Giere, R. N. (1994). Viewing Science. In *Proceedings of the Biennial Meeting of the Philosophy of Science Association, 1994, Volume Two: Symposia and Invited Papers* (pp. 3-16).
- 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.
- Goodman, N. (1976). *Languages of art* (2nd ed.). Indianapolis and Cambridge: Hackett.
- Kitcher, P. (2001). *Science, Truth, and Democracy*. New York: Oxford University Press.
- Massimi, M. (2018). Perspectival modeling. *Philosophy of Science*, 85(3), 335-359.
- Monmonier, M. (1991). *How to lie with maps*. Chicago: University of Chicago Press.
- Morrison, M. (2011). One phenomenon, many models: inconsistency and complementarity. *Studies in History and Philosophy of Science*, 42(2), 342-351.
- Nguyen, J. (2020). It's not a game: accurate representation with toy models. *The British Journal for the Philosophy of Science*, 71(3), 1013–1041.
- Nguyen, J., & Frigg, R. (2020). Unlocking Limits. *Argumenta*, 6(1), 31-45.
- Nguyen, J., & Frigg, R. (2022). *Scientific Representation* (Cambridge Elements). Cambridge: Cambridge University Press.
- Pearson, F. (1990). *Map Projections: Theory and Applications*. Boca Raton (Florida): CRC Press.
- Robinson, A. H., & Petchenik, B. B. (1976). *Nature of Maps. Essays Toward Understanding Maps and Mapping Hardcover*. Chicago: University of Chicago Press.
- Salis, F., Frigg, R., & Nguyen, J. (2020). Models and denotation. In C. Martínez-Vidal, & J. L. Falguera (Eds.), *Abstract objects: for and against* (pp. 197-219). Cham: Springer.
- Shech, E. (2015). Scientific misrepresentation and guides to ontology: the need for representational code and contents. *Synthese*, 192, 3463-3485.
- Sismondo, S., & Chrisman, N. (2001). Deflationary metaphysics and the nature of maps. *Philosophy of Science (Proceedings)*, 68, 38-49.

- Toulmin, S. (1953). *The Philosophy of Science*. London: Hutchinson's University Library.
- Weinberg, S. (1993). *Dreams of a final theory: The search for the fundamental laws of nature*. New York: Vintage.
- Winther, R. G. (2020). *When Maps Become the World* (Studies in History and Philosophy of Science). Chicago: University of Chicago Press.