Successful Visual Epistemic Representation

Agnes Bolinska

Abstract: In this paper, I characterize *visual epistemic representations* as concrete two- or three-dimensional tools for conveying information about aspects of their target systems or *phenomena of interest*. I outline two features of successful visual epistemic representation: that the vehicle of representation contain sufficiently accurate information about the phenomenon of interest for the user's purpose, and that it convey this information to the user in a manner that makes it readily available to her. I argue that actual epistemic representation may involve tradeoffs between these features and is successful to the extent that they are present.

1. Introduction

How do scientists gain information about a physical system? The most straightforward way would be to examine the system using their unaided senses. But the senses often do not suffice for observation because the system is too small (like a molecule), too distant (like a star), too dispersed (like a population), or otherwise imperceptible. In such cases, scientists may use instruments to facilitate their investigation. But once they have exhausted all observational avenues – once they have done everything possible to gain information from the system itself – they usually have only completed the first phase of their investigation. For scientists are generally not interested in particular measurements, but in generalizations or patterns that can be inferred from them.[1] These patterns, which I will refer to as *phenomena of interest*, are rarely accessible through direct observation.[2] So the question is: once all available observational data has been acquired, how can scientists use a vehicle of representation – an entity physically separate from the system it represents – as a tool for gaining information about the phenomenon of interest?

The answer, I will argue, is by representing the system in a way that makes these patterns perspicuous to the user. I call vehicles of representation that are used as tools for gaining information about phenomena of interest *epistemic representations*. As I will show, these include not only scientific representations, but also other representations that are used in similar ways outside of scientific practice. By bringing the features of interest to the fore, an epistemic representation unlocks for the user a dimension of access to the phenomenon of interest that she wouldn't otherwise have.

There are two sorts of contexts in which epistemic representations may be used: those in which little or nothing is known about the phenomenon of interest, and the representation functions as an *investigative* tool; and those in which the creator of the representation already understands the phenomenon of interest fairly well and uses it as a tool for *conveying* information about this phenomenon via testimony. In this paper, I focus on the latter sort of context and limit my attention to two- and three-dimensional visual representations. I investigate the features in virtue of which such representations convey

1 The point that scientists tend to be interested in patterns or regularities has been nicely articulated by Batterman (2009, pp. 429-30).

2 Why this is so will become clear shortly through the consideration of examples (Section 2).

information to their users about phenomena of interest that they wouldn't otherwise have. That is, I determine the features of *successful* representation for these kinds of cases.[3]

In order to understand what is required for successful visual epistemic representation, two questions must be addressed. First, what kind of information ought the vehicle of representation *contain*, and what is required for it to contain this information? Second, how is this information *effectively conveyed* to the user? I will argue that an epistemic representation is successful to the extent that it contains sufficiently accurate information about the phenomenon of interest for the user's purpose (Section 3) and is able to convey this information to the user in a manner that makes it readily available to her (Section 4). I will show that because visual epistemic representation often involves tradeoffs between these two features, the success of such representation is determined by how well they are balanced (Section 5). But I will begin (Section 2) by outlining three examples of successful visual epistemic representation and showing that they share two general features that inform the more specific features considered in Sections 3 and 4: they are user- and purpose-specific.

2. Examples and general features

User- and purpose-specificity are widely accepted features that comprise a central part of several accounts of scientific representation (cf. Bailer-Jones, 2003; Giere, 2004, 2009; Mäki, 2009; Teller, 2001). I highlight these features here to provide a foundation

3 Many authors (cf. Callender & Cohen, 2006; Contessa, 2007; Suárez, 2004) focus on understanding mere – that is, not necessarily true or accurate – representation. I think that understanding successful epistemic representation (which, as I will show, isn't just true or accurate representation) can also guide us in understanding 'mere' representation. While a full discussion of why this is so is beyond the scope of this paper, it will become apparent once the account has been put forth.

from which the remainder of my analysis may be developed. By taking the user- and purpose-specificity of epistemic representations as a starting point, we may then ask further questions about these features, thereby better coming to understand visual epistemic representation. In which ways, precisely, are such representations user- and purpose-specific? In virtue of what are they so?

In this section, I will present three examples of successful visual epistemic representation, in each case identifying the phenomenon of interest, indicating the sense in which it is inaccessible to its user(s) through direct examination of the physical system in question, and highlighting the ways in which the representation is successful only for certain users and for specific purposes.

The first example is adapted from Suárez (2004). Suppose we want to represent a system consisting of two ships travelling along the sea using two pens and a piece of paper. Let us further assume that the representation is used in the context of a conversation between the captain of one of the ships and her friend. To help recount the highlights of her last voyage, the captain might move the two pens along the paper to demonstrate, for instance, a maneuver she had to perform to avoid colliding with another ship that had strayed off course. In this example, the phenomenon of interest consists of the relative trajectories of the ships. It is inaccessible to the captain's friend because he was not present to witness her collision-avoiding maneuver. While the intended user of the pens-on-paper system is the captain's friend, it would be suitable for many other users as well, since little background knowledge is required to understand the relative motions of the ships. The purpose of the pens-on-paper system is to help the captain relay certain parts of her voyage to her friend.

A second example of a successful epistemic representation is the iconic map of the London Underground transit system.[4] Originally designed by draftsman Harry Beck in 1933 and modelled after a circuit diagram, the network of railway lines and stations that comprise the Underground is depicted as an orderly array of intertwined coloured lines, along which lie evenly-spaced marks labelled with station names, with white circles replacing these marks to designate interchange stations. Included with the map are keys that tell users how to interpret each of its features (Figure 1). With the aid of this map, approximately three million daily users of the Underground are able to navigate this expansive system. The phenomenon of interest varies between users: for each, it is the set of the possible routes connecting the stations between which she wishes to travel. These routes are inaccessible to her, since while she could in principle ride the Underground in various directions to determine which one connects her to her destination station, this would be extremely impractical. Because the map of the London Underground is intended for use by a broad range of people with a variety of backgrounds and cognitive capacities, it is designed to cater universally to human users.[5] The purpose of the map is to determine the most efficient route from one station to another.

4 This example is also discussed by Contessa (2007), Hoover (2012) and Bolinska (2013).
5 This is not to say that no training whatsoever is required – users must be familiar with the conventions involved in reading maps – but that the level of training required is relatively low and fairly universally held in many parts of the world.

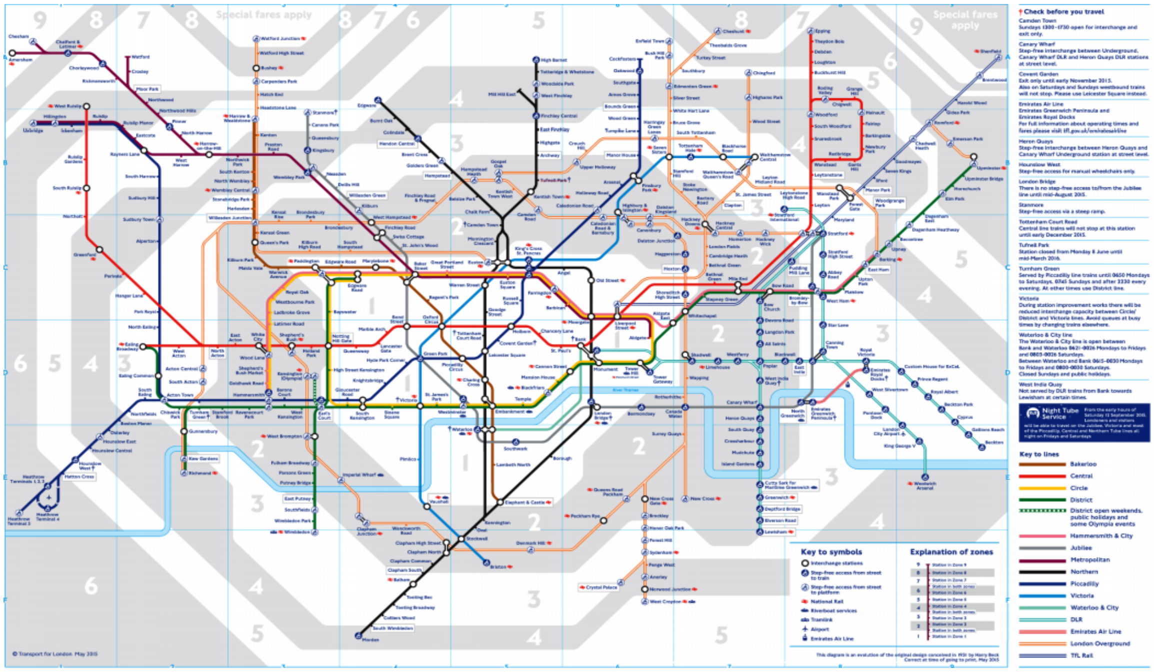

Figure 1: The complete London Underground map. Note the even spacing of labelled stations, colouring of railway lines (with key in the bottom right corner) and white circles indicating interchange stations.

Finally, a third example of a successful epistemic representation is a three-dimensional model of a macromolecule like DNA or protein. The phenomenon of interest in this case is the structure of the molecule, *viz.* its three-dimensional shape, including bond types, lengths and angles between constituent atoms. This structure is otherwise inaccessible to the user: it is not directly observable even using techniques like x-ray diffraction photography, since the images produced using such techniques must be interpreted to yield putative structures, and the process of interpretation can yield results that are often ambiguous or misleading.[6] Unlike the map of the London Underground, a molecular model's key is implicit, so users must be told which features of the model correspond to which features of a molecule, e.g. that white balls stand of hydrogen atoms, black for carbon, etc. While any user who understands this convention may grasp the

6 The problem of determining molecular structure from x-ray diffraction photographs in the mid-twentieth century was notoriously difficult. See Olby (1974) and Judson (1996).

structure as a whole, molecular models are most useful for one with training in molecular biology in the pursuit of further aims. For instance, such a user may rely on the knowledge of the structure she gains from the model to determine the function of the molecule or how it will interact with other molecules. Thus, the purpose for which the model is used often extends beyond simply learning about molecular structure.

3. Containing sufficiently accurate information

With these examples in hand, we may now turn to the more specific features of successful epistemic representation, each of which depends on the phenomenon of interest, the user, and the purpose for which the representation is used. The aim of employing epistemic representations is to learn about an aspect of the target system, the phenomenon of interest. But as I showed in the previous section, users often don't seek to learn about this phenomenon for its own sake, but rather to use what they learn for some further purpose. How accurate the information they gain need be depends on what this purpose is.

3.1 When is information sufficiently accurate?

In the pens-on-paper example the user is interested in the ships' relative trajectories as an aid to understanding the relevant parts of the captain's voyage. Thus, only very general information about the trajectories need be contained in the pens-on-paper representation. For instance, let us assume that one ship was on a head-on collision course with the other, and the other veered off to the right just in time to avoid the collision. This information should be contained in the vehicle of representation, but the vehicle need not accurately portray the precise angle of approach between the ships' trajectories, nor the angle at which the second ship veered off course to avoid the collision. So long as the more general information is contained in the vehicle of representation, the vehicle may serve its function as an aid in the telling of the story.

The map of the London Underground is used with the aim of determining the most efficient journey between stations selected by its users; thus, it should be possible to determine the most efficient route between *any* two stations in the network using the map. In order to be used for this purpose, the map must contain information about the relative locations of stations along each line, including information about the locations of interchange stations at which users may switch lines. It need not contain information about other features of the network, such as its position with respect to above-ground streets or the distances between stations to represent the network faithfully for its users' purposes.[7] As in the ships-on-sea example, being more accurate in this respect does not contribute to the success of the epistemic representation. In fact, as we will see in Section 5, it could even detract from its success if increased accuracy comes at the price of a reduction in the ease with which this information may be conveyed to the user.

It may be objected here that it's not the case that the map of the London Underground represents features such as distance between stations inaccurately: distances between stations are *omitted* from the map, not misrepresented, since the map's content is determined by its intended use. This objection raises an important point about

7 We may assume that the distances between stations are sufficiently similar to one another that in each case the most direct route between stations – that requiring the fewest number of stops and changes between lines – will be the quickest.

the relationship between how much information is contained in a representation and how accurate that information is.

In general, the fact that one representation contains less information than another does not render it less accurate. 'Smith lives in Canada' is less informative than 'Smith lives in Toronto', yet both sentences may be equally accurate, e.g. if Smith lives in Toronto, Canada, or the former may be more accurate than the latter, e.g. if Smith lives in Mississauga, a suburb of Toronto. But in the case of epistemic representation, we are concerned not with accuracy *simpliciter*, but in *accuracy-for-a-purpose*. And accuracy-for-a-purpose is closely related to the amount of information a representation contains. Recall that epistemic representations are used to convey information not about the target system in general, but about the *phenomenon of interest*; moreover, users are interested in gaining this information *for a particular purpose*.

If a representation is not sufficiently information-rich – if it omits bits of information that are necessary for understanding the phenomenon of interest for that purpose – then it cannot represent the phenomenon of interest sufficiently accurately for that purpose. The map of the London Underground omits information about distances between stations – it represents these stations as evenly spaced, but users are not meant to infer from this feature of the map that distances between stations are, in fact, equal. Suppose the map also omitted information about where the lines intersect, simply depicting them in parallel with one another. The omission of this crucial information would render this map not only less informative, but also less accurate, given a user's interest in getting from any one station in the system to any other. Key pieces of information, such as the places in which the lines intersect, must be contained in the map in order for it to accurately portray routes between stations for this purpose. So the accuracy of an epistemic representation – that is, its accuracy-*for-a-purpose* – depends in part on it containing a sufficient amount of information to accurately represent the phenomenon of interest for the user's purpose.

Finally, in the examples of molecular modelling, the target systems of the representations are molecules like protein or DNA, while the phenomena of interest are the structures of these molecules. The purpose of gaining information about these structures is often to enable further research into molecular function. For instance, knowing the structure of DNA readily suggested a copying mechanism for this molecule of inheritance. Thus, unlike in the previous two examples, where only general and relative positions of the relevant objects are required for users' purposes, fairly specific information about atomic positions and bond lengths and angles needs to be contained in the vehicle of representation in these cases to enable further investigation. Without such specific information, one cannot, for instance, determine what accounts for DNA's stability or how it may be 'unzipped' and replicated.

3.2 How does a vehicle contain information?

We've now established *which* information need be contained in a vehicle of visual representation in order for that vehicle to epistemically represent its target system: information about the phenomenon of interest that is sufficiently accurate for its user's purpose. We will now turn to another question: how does a vehicle come to contain that information? While the term 'information' may be used differently in other contexts,[8] here we are concerned with *semantic* information, the sense in which this term is perhaps most commonly used. This is the sense in which a newspaper contains information about current events, a map contains information about the layout of a city, and a model contains information about the physical system it represents.

Semantic information may be either natural or non-natural information.[9] Natural information is a product of certain naturally occurring correlations. Examples include smoke carrying the information that there is fire and rings on a tree carrying information about the tree's age. We will generally not be concerned with natural information, since most vehicles of epistemic representation, and particularly scientific models, are human constructs. Thus, 'information' will here be used to refer to non-natural information.

An artefact may contain non-natural information by convention. For example, in a weather forecast, a picture of a cloud and rain labelled '80%' contains the information that there is an 80% chance of showers, and a sign that says 'Gift shop downstairs' contains the information that the gift shop is downstairs. In these examples, the conventions via which the vehicles contain information are known to most people who encounter them. But this need not be the case. Before arriving at a party, I might establish with my companion that either of us uttering the sentence 'This wine is fantastic' indicates his or her desire to leave. In this case, 'This wine is fantastic' contains the information that its utterer would like to leave the party (provided that the utterer is either my companion or I).

8 For instance, Fisher information, Shannon information, Kolmogorov complexity and quantum information each refer to different notions of information. For a survey of how these notions are used and related to one another, see Adriaans (2013).

9 Grice (1957) made the distinction between natural and non-natural *meaning*; Piccinini and Scarantino (2010) have extended this distinction to *information*.

The information contained in a vehicle of representation need not be true. A sign that says 'Gift shop downstairs' contains the information that the gift shop is downstairs even if, as a matter of fact, it has been relocated to the second floor; in this case, the sign contains false information.[10] Similarly, suppose I really like the wine being served at the party, and carelessly remark, 'This wine is fantastic', even though I'm having a good time and don't yet want to leave. In this case, the sentence nevertheless contains the information that I want to leave the party, but this information is false.

Because a vehicle of representation may carry information by convention, is it the case that anything may carry information about anything else, just as anything may denote anything else? As noted earlier, information is a multifarious concept; here we are restricting our attention to the sense in which information can be contained in a vehicle of representation about an aspect of a target system, not, for instance, the sense in which information is carried in our genetic code, or that in which it is stored in computer databases. Neither the information contained in our genetic code nor that stored in computer databases is directly accessible to a user who consults these things directly. Rather, a complex, multi-step process is necessary for its extraction. A DNA molecule must first be sequenced; the sequence must then be decoded by someone with the requisite training and experience. A computer's hard drive doesn't issue information about its contents directly, but must be accessed using particular software and hardware.

10 Many authors (Barwise & Seligman, 1997; Floridi, 2004, 2005; Graham, 1999; Grice, 1989) take truth to be a requirement for information. For instance, Dretske contends that "*false* information and *mis*-information are not kinds of information – any more than decoy ducks and rubber ducks are kinds of ducks" (1981, p. 45). However, this requirement fails to accommodate many of the ways in which 'information' is used. Moreover, the authors who hold this position have failed to give a good argument for why it should be adopted (Scarantino & Piccinini, 2010).

In contrast, it only makes sense to speak of a vehicle of epistemic representation – a tool for conveying information – as containing information if that information is available to the user when she consults the vehicle. The sentence 'This wine is fantastic' contains the information that its utterer wants to leave the party: once the convention has been established, the person who hears this sentence may grasp that its utterer wants to leave the party. 'This wine is fantastic' does not contain the information that its utterer wants to leave the party *because he or she is tired*, since this further bit of information was not included in the convention when it was established. If my companion utters 'This wine is fantastic' with a yawn and vacant stare, I might infer that he wants to leave the party because he is tired, but this inference would be based on the information the sentence contains *together with* information obtained from the yawn and vacant stare accompanying its utterance.

Of course, if my companion and I were to establish such a convention at the outset, 'This wine is fantastic' would contain the information that the utterer of the sentence wants to leave the party because he or she is tired. In this case, if I heard my companion say 'This wine is fantastic', I could immediately infer that he wants to leave the party because he is tired, even absent the yawn and vacant stare. This sentence could also be made to contain more information than this, such as that its utterer wants to leave the party because he or she is tired and has become bored with the conversation. But the amount of information that 'This wine is fantastic' can contain is constrained by the limitations of its users' memories. So, for instance, it cannot contain information about the names and birthdays of each of the party's guests: even if I stipulate that 'This wine is fantastic' contains the information that John Smith's birthday is on February 24th and Anne

MacDonald's is on June 2nd and so on for each of the party's guests, the only people familiar with the convention – my companion and I – are unable to extract that information from the sentence.[11] Thus, there is no meaningful sense in which this vehicle really does contain that information.

In general, the more discernable subcomponents a vehicle has, the more information it may contain. A discernable subcomponent is a part of the vehicle that can be distinguished from other parts by its user, and what constitutes a discernable subcomponent in each case depends on the means of representation and the user. For example, each of the pens and the piece of paper are discernable subcomponents of the pens-on-paper system, but an arbitrarily chosen point on the blank piece of paper is not, since it can't be distinguished by an ordinary user from the rest of the paper.[12]

This is important because it precludes certain vehicles from counting as successful epistemic representations in a given set of circumstances. For instance, consider an alternative to the pens-on-paper representation, the *paper*-on-*pens* system, in which a blank piece of paper represents the ship trajectories and two pens represent the surface of the sea. While such a system may certainly *denote* the ships-on-sea system, it cannot be an *epistemic* representation of that system. For there is no genuine sense in which information about the individual trajectories of the ships can be conveyed to a user

11 John Kulvicki (2010) argues that more complex or detailed information can be contained in a fairly simple vehicle of representation, but may not be 'extractable' from that vehicle. For instance, the information that its target is square is extractable from the representation 'square'; however, 'square' also contains the information that its target has four sides of equal lengths and right internal angles, but this information is available only via an inference. In contrast, I am concerned with information that is contained in a vehicle of representation and available to the user either immediately from the vehicle *or* via this kind of direct inference. Thus, my usage of 'extractable' differs from his more specific sense of the term.
12 There may be more than one way to individuate subcomponents, but this will not be important for our purposes, as will become clear shortly.

consulting the blank piece of paper; as Suárez points out, such a representation “seems counterintuitive and unnatural” (2004, p. 772).

To see why, let us consider in more detail how the captain might use the paper-on-pens system to represent the ships-on-sea system in her recounting of the details of her voyage. The general strategy would be to stipulate that the paper represents the trajectories of the ships and the pens represent the surface of the sea before recounting these events. Then, upon reaching the point in her story where she performs the collision-avoiding maneuver, she could simply gesture towards the paper-on-pens system. In order to stipulate that the paper represents the trajectories of the ships and the pens the surface of the sea, the captain would have to somehow get her friend to understand what these trajectories looked like, either by describing them ('I was heading due north when I noticed a ship coming from port side up ahead…'), or tracing the trajectories with her hands in the air, or even using the pens to trace them on the paper (!) prior to reversing the representational arrangement. Once this was done, she could then use the paper to represent the trajectories of the ships and the pens the surface of the sea. But this would be an odd thing to do at this stage, since all of the work required to convey information about the trajectories would already have been completed in the earlier stipulation of the representational arrangement. When the captain reached the part of her story in which she describes her collision-avoiding maneuver, gesturing towards the paper to represent the trajectories of the ships would be superfluous; the captain could just as easily say 'And that’s when I performed that maneuver I described earlier' instead. Thus, the paper-on-pens system cannot be used as an epistemic representation of the ships-on-sea system because it does not have the resources to individuate the relative trajectories of two ships.

The paper is a single uniform object; at least two distinct objects are required to convey information about these trajectories.

Each of our examples of successful epistemic representation contains information about the phenomenon of interest because it contains a sufficient number of subcomponents, each of which carries more specific information about subcomponents of this phenomenon. In the ships-on-sea system, these are the individual ship trajectories; the trajectory of each pen in the pens-on-paper system and the positions of these trajectories with respect to one another convey to the user information about the phenomenon of interest.

Similarly, in order for the map of the London Underground to be able to carry information about the most efficient routes between stations, it must carry several specific bits of information that are necessary for making this determination. In particular, it must carry information about the relations between lines – where they intersect, allowing users to travel between them – and information about which stations are located along each line and in which order. The map contains features that are responsible for carrying each of these bits of information. The railway lines of the London Underground system are represented as distinctly coloured lines on the map; interchange stations are labelled with a white circle overlapping the lines between which the user may travel; and individual stations are represented by protrusions along each line labelled with station names, the order of the stations along the lines in the map reflecting the order of the stations in the Underground system.

Finally, in order to convey to the user information about molecular structure as a whole, molecular models must contain specific information about positions of particular atoms and the lengths and angles of the bonds that join them. This information is contained in the pieces representing atoms and bonds, and the ways in which these pieces are connected to one another.

In this section, I have argued that human cognitive and perceptual capacities impose limits on the amount of information a vehicle may contain about its target system: users must be able to extract information from the representation, and they must be able to distinguish between subcomponents of the vehicle which stand for subcomponents of the target system. But while a vehicle's containing information depends in part on the possibility of using it to convey that information, vehicles of representation nevertheless differ with respect to *how well* they convey information about the phenomenon of interest to their users. The extent to which this information is conveyed is a second important feature of successful epistemic representation. In the next section, we will examine the notions of syntactic and semantic salience, the features that make information cognitively available to users, and thus contribute to how easily information is conveyed.

4. Conveying information

In order for the information contained in a vehicle of epistemic representation to be effectively conveyed to its user, it must be readily accessible to her. In this section, following Kulvicki (2010), I identify two features that contribute to this ready access: the syntactic salience of subcomponents and the semantic salience of the phenomenon of interest. Because both come in degrees, information is effectively conveyed to a user to the extent that they are present.

4.1 Syntactic salience of subcomponents

The success of visual epistemic representation depends in part on the extent to which a feature of the vehicle of representation that carries information is perceptually prominent or s*yntactically salient* for a user (Kulvicki, 2010). When the trajectories of the ships in the ships-on-sea system are represented by the trajectories of the pens in the pens-on-paper system, they very clearly stand out against the backdrop of the paper. Similarly, each of the features of the London Underground map necessary for conveying information about the best possible routes between stations is perceptually accessible. For example, individual railway lines of the Underground system are syntactically salient because they are represented in the map as distinctly coloured lines set against a uniformly white background. Had the same colour but different degrees of saturation been used for each of the lines, the distinction between these lines would have been less syntactically salient for the average user of the map, since humans generally do not have the same capacity to recognize differences between degrees of saturation that they do to distinguish between different colours. Had identical colours *and* degrees of saturation been used, the lines representing railway lines would be completely indistinguishable for the user, and thus not syntactically salient at all. The pieces of which molecular models are built are clearly also syntactically salient: they are very easily perceived against the backdrop of the three-dimensional space in which they are constructed.

Note that whether or not an aspect of a representation is syntactically salient depends on the perceptual capacities of the user(s) of that representation. If someone is colourblind, then a map of the London Underground in which the same colour but different degrees of saturation were used for each of the railway lines might be more syntactically salient for him. However, we must keep in mind that each of our examples of successful epistemic representation is designed to make the relevant features syntactically salient for most of its intended users. Thus, since most people aren't colourblind, the different railway lines are represented with different colours, which are easily distinguishable from one another for the average person.

The syntactic salience of the parts of the vehicle that represent subcomponents of the phenomenon of interest contributes to the success of an epistemic representation. But it is possible to represent these subcomponents in a variety of equally syntactically salient ways, not all of which are equally able to give the user the kind of access to the phenomenon of interest she seeks through using an epistemic representation. For example, consider an alternative representation of the London Underground in which the railway lines are depicted in parallel to one another, rather than as a network, with interchange stations labelled with the other line(s) that can be accessed at them in parentheses next to their names.[13] So, for example, along the Bakerloo line, Picadilly Circus Station would be labelled 'Picadilly Circus (Picadilly)', since the Bakerloo and Picadilly lines intersect at Picadilly Circus Station (Figure 2).

13 Note that this is different from the parallel-lines representation mentioned in Section 3, in which interchange stations were not labelled.

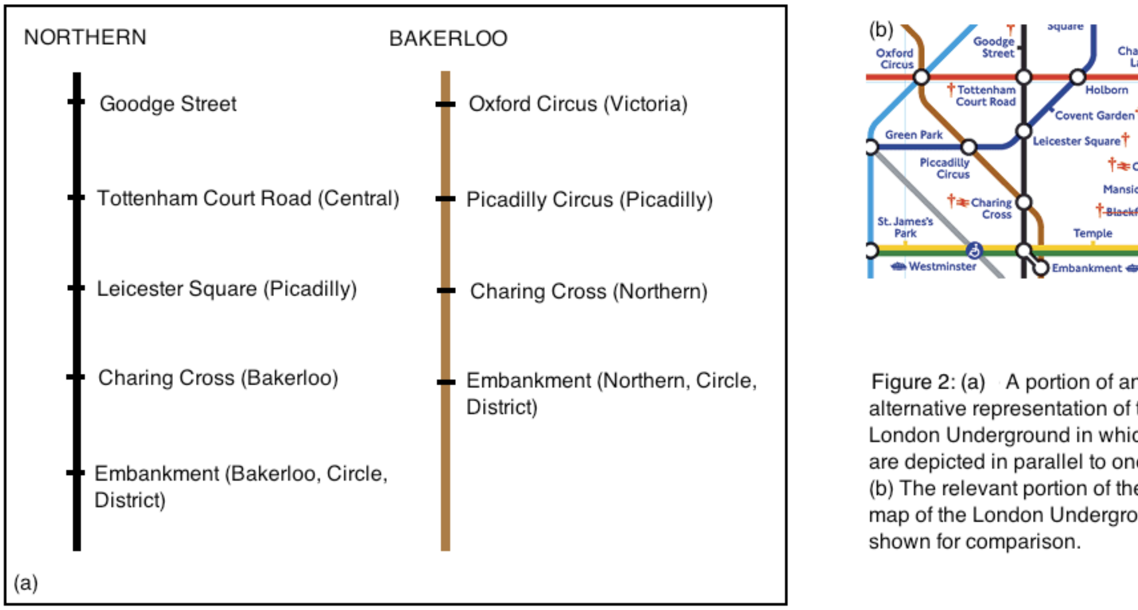

Figure 2: (a) A portion of an alternative representation of the London Underground in which lines are depicted in parallel to one another. (b) The relevant portion of the actual map of the London Underground, shown for comparison.

Such a representation of the London Underground makes the same bits of information about the relative positions of stations along each line and the locations of interchange stations just as syntactically salient as the map does, since they are equally perceptually available to the user. Just as a user has no problem detecting the white circles on the map, she also has no problem detecting the names of the other lines that can be accessed at an interchange station: she can easily see their names in parentheses next to the name of the station in question. However, it would be much more difficult to use such a representation for the purpose of determining the optimal route between stations than it is to use the map. While the individual pieces of information that are necessary to identify possible routes – positions of stations along each line and locations of interchange stations – are syntactically salient, the possible routes themselves, and thus information about which of these routes is best, are more difficult to discern. In order to find these routes, a user must perform additional inferential steps to decode, so to speak, the alternative representation. But this kind of extra work is precisely what a user seeks to minimize through the use of an epistemic representation.

Similarly, consider some alternative ways of representing molecular structure that also render its subcomponents – atomic positions and bond lengths and types – syntactically salient. For instance, a list of atomic coordinates together with a specification of the lengths and types of bonds between each pair of atoms makes all the relevant subcomponents syntactically salient; nevertheless, the overall structure of the molecule would, as with the alternative parallel-lines representation of the London Underground, require the user to perform a number of additional inferential steps to 'see' this structure. A two-dimensional image of the structure fares better, since it is possible to discern three-dimensional structure from such an image with fewer and simpler inferences. But the three-dimensional model is best.

4.2 Semantic salience of phenomenon of interest

These alternative representations of the London Underground and molecular structure lack a feature that Kulvicki (2010) calls *semantic salience*. Semantic salience refers to the ease with which syntactically salient features of the representation convey their content to the representation's user. In order for a part of a representation to be semantically salient, it must be straightforward for a user to determine what that part represents: "there must be a plan of correlation between features of the representation and features of the data that is easy to grasp" (Kulvicki, 2010, p. 301). For example, suppose you want to represent a temperature gradient in which the temperature decreases continuously as you move away from a certain point. Representing the highest temperature with white, the lowest with black, and the remainder on a grayscale with lighter shades corresponding to higher temperatures makes this gradient semantically salient: it is easy to interpret the light-to-dark grayscale as corresponding to a high-to-low temperature scale. One could also arbitrarily assign shades of gray to temperatures, but the resultant image would not be semantically salient: it would be difficult to see that the temperature decreases continuously away from the hot spot, since one would have to refer back to the key that assigns temperatures to colours repeatedly before being able to make this determination.

The extent to which something is a successful epistemic representation is determined not only by the *syntactic salience* of the *subcomponents* of the vehicle representing parts of the phenomenon of interest, but also whether the *phenomenon of interest* is *semantically salient*. In the case of the London Underground, information about the phenomenon of interest – possible routes between stations – is semantically salient in the map, but not in the parallel-lines representation. The difference between these representations lies in the ways in which relative positions of the railway lines vis-à-vis interchange stations are represented. In the map, it is easy to see that the intersection of lines with white circles represents places where the lines literally cross, allowing users to travel between them. That is, the locations of interchange stations are semantically salient and this is necessary for making possible routes between stations semantically salient as well.

The parallel-lines representation, on the other hand, forces its user to perform extra steps in her reasoning to obtain the same information; the locations of the interchange stations – and thereby also possible routes between stations – are not semantically salient in this representation. Labelling interchange stations with the name(s) of the intersecting line(s) makes information about possible routes between stations more difficult to grasp than in the case of the map, in which interchange stations are marked with a white circle along two or more intersecting or adjacent lines. And it is precisely this information that is required to determine possible routes between stations when stations do not lie along the same line.

In the molecular modelling examples, the structures of the molecules – the phenomena of interest – are semantically salient in the physical models. It is straightforward to interpret the physical structure of the model as the structure of the molecule. This is not so for some alternative representations that also make the subcomponents of the structures syntactically salient. For instance, a list of the coordinates of each atom together with a list of the types of bonds that join them makes all the relevant subcomponents of the structure syntactically salient, but the overall structure of the molecule is more difficult to determine. As with the alternative parallel-lines representation of the London Underground, a user would have to perform a number of additional inferential steps to 'see' this structure. A two-dimensional image of the structure fares slightly better, but since the three-dimensional structure is the phenomenon of interest, the models make this structure more semantically salient than their two-dimensional counterparts.

Thus, we can now see the role that 'mirroring' may play in successful epistemic representation: it may contribute to the semantic salience of the phenomenon of interest. For instance, the London Underground map more closely mirrors the railway network than the parallel lines representation, since the lines in the network do, in fact, cross one another. Similarly, while molecular structure could be represented in a number of different ways, not all of them make the structure as easy to grasp as molecular models do. Thus, even though mirroring relations are not, strictly speaking, necessary for successful epistemic representation, they may play a role nonetheless in securing semantic salience.

5. Successful visual epistemic representation

I have outlined two features that contribute to successful visual epistemic representation: that the vehicle of representation contain sufficiently accurate information about the phenomenon of interest for the user's purpose, and that it effectively convey this information to the representation user. Both of these features may be present to varying degrees. Accordingly, vehicles of epistemic representation may be more or less successful, depending on the extent to which these features are realized.

This is especially true because sometimes one feature of successful epistemic representation comes at the cost of another. Consider again the map of the London Underground. Recall that users of this map are interested in possible routes between stations, with the ultimate aim of determining which of these is most efficient. Given this aim, one might argue that this map could be made to better serve its users if it depicted the stations and railway lines to scale, thereby representing possible routes between stations more accurately. After all, this would give users the information they currently are able to extract from the standard map, as well as other potentially relevant information. For example, it would be possible to better estimate travel times from such a map[14] or to determine when it might be worthwhile to walk between stations, rather than taking the

14 Of course, in order to do this, one would have to take into consideration things like congestion at each station, which would determine how long trains stop there.

Underground. This additional information would make the more accurate for the purpose of determining the most efficient route between stations.

However, such an increase in accuracy would come at the expense of the syntactic salience resulting from the stations being positioned evenly along the lines and the lines crossing one another only at ninety or forty-five degree angles. For in this case, the areas of the map corresponding to parts of the Underground in which stations are close together would be densely populated by station marks and labels, making it difficult to distinguish between the stations. And the ability to easily locate different stations on the map is essential for the map's performing its primary epistemic function: to allow users to most efficiently determine the best routes between stations.

Tradeoffs between the accuracy of the information contained in the vehicle and the extent to which that information is conveyed to the user exist not as a matter of law, but as a matter of fact. In principle, there is nothing precluding a vehicle to both contain information that is maximally accurate for the user's purpose and to effectively convey that information. In practice, however, it is difficult to fully realize both of these features at the same time, except perhaps in the simplest cases. The crucial point is that the success of visual epistemic representation depends on both features and, accordingly, how well they are balanced with respect to the representation user and the purpose for which she wishes to gain information about the phenomenon of interest.

6. Conclusion

I have argued that the success of a visual epistemic representation depends on the extent to which it contains sufficiently accurate information about the phenomenon of interest for the user's purpose and makes this information readily accessible to her. In so doing, I have introduced some ways in which we may compare representations to one another, determining the degree to which they are successful in a particular context, defined in terms of a user's background, training, cognitive capacity and interests. Thus, this paper develops accounts focused on the user- and purpose-specificity of representation (mentioned in the beginning of Section 2) by specifying *how* the user and purpose affect the choice of representational vehicle as a tool for gaining information. The purpose determines how accurate the information contained in the vehicle of representation need be, while the background and cognitive capacity of the representation user determines what is required for that information to be effectively conveyed to her.

Rather than giving a set of necessary and/or sufficient conditions for epistemic representation, this account instead identifies features that *contribute* to its *successful* use as a tool for gaining information. It thus both departs from the standard methodology of analyzing what is required for a scientific model to be a representation of its target system *at all*. An approach that acknowledges that actual representations vary with respect to the extent to which they give their users epistemic access to a phenomenon of interest promises to more accurately reflect scientific practice, where a variety of representational vehicles are used for myriad purposes.

The account developed in this paper focuses on visual representation. A question for future work is whether it might be fruitfully extended to other modes of epistemic representation, such as mathematical models, and if so, how? Like visual representations, the success of mathematical models also depends in part on their containing information that is sufficiently accurate for the user's purpose. The notions of semantic and syntactic salience, however, are likely to have limited applicability to mathematical models, since visualization of their targets may be impossible or misleading. For instance, the position of a particle may be represented by a probability distribution, i.e. as a 'superposition' of different positions. Yet it is impossible to visualize what it would mean for the particle to be in such a superposition. Nevertheless, mathematical models make the information they contain readily available to users who have the requisite training. Whether the notions of semantic and syntactic salience could be fruitfully brought to bear in the realm of mathematical models, and if so, how important they would be for understanding these kinds of models as tools for conveying information, are questions worthy of further attention.

Acknowledgments

For comments on earlier drafts of this paper and helpful discussions, I'd like to thank Gillian Barker, Jossi Berkovitz, Anjan Chakravartty, Alex Djedovic, Eugene Earnshaw Whyte, Roman Frigg, Fermín Fulda, Cory Lewis, Greg Lusk, Margaret Morrison, Isaac Record, Mauricio Suárez, Jim Weatherall, and Rasmus Grønfeldt Winther. I'm grateful also to audiences at the Models and Simulations 6 conference at the University of Notre Dame and the Center for Philosophy of Science at the University of Pittsburgh, where I presented this work. Finally, I thank two anonymous referees, whose constructive comments significantly improved this paper.

References

Adriaans, P. (2013). Information. In F. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy* (Fall 2013.).

Bailer-Jones, D. (2003). When scientific models represent. *International Studies in the Philosophy of Science*, *17*(1), 59–74. doi:10.1080/02698590305238

Barwise, J., & Seligman, J. (1997). *Information Flow: The Logic of Distributed Systems*. Cambridge: Cambridge University Press.

Batterman, R. W. (2009). Idealization and modeling. *Synthese*, 169(3), 427-446.

Bolinska, A. (2013). Epistemic representation, informativeness and the aim of faithful representation. *Synthese*, (190), 219–34. doi:10.1007/s11229-012-0143-6

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

Contessa, G. (2007). Scientific Representation, Interpretation, and Surrogative Reasoning. *Philosophy of Science*, *74*(1), 48–68. doi:10.1086/519478

Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge: Cambridge University Press.

Floridi, L. (2004). Outline of a Theory of Strongly Semantic Information. *Minds and Machines*, *14*(2), 197–221. doi:10.1023/B:MIND.0000021684.50925.c9

Floridi, L. (2005). Is Semantic Information Meaningful Data? *Philosophy and Phenomenological Research*, *70*(2), 351–370. doi:10.1111/j.1933-1592.2005.tb00531.x

Giere, R. N. (2004). How Models Are Used to Represent Reality. *Philosophy of Science*, *71*(5), 742–752. doi:10.1086/425063

Giere, R. N. (2009). An agent-based conception of models and scientific representation. *Synthese*, *172*(2), 269–281. doi:10.1007/s11229-009-9506-z

Graham, G. (1999). *The Internet: A Philosophical Inquiry*. London: Routledge.

Grice, H. P. (1957). Meaning. *The Philosophical Review*, *66*(3), 377–388.

Grice, H. P. (1989). *Studies in the Way of Words*. Cambridge, Mass.: Harvard University Press.

Hoover, K. (2012). Pragmatism, Perspectival Realism, and Econometrics. In A. Lehtinen, J. Kuorikoski, & P. Ylikoski (Eds.), *Economics for Real: Uskali Mäki and the Place of Truth in Economics* (pp. 223–240). London: Routledge.

Judson, H. F. (1996). *The Eighth Day of Creation: Makers of the Revolution in Biology* (expanded.). Plainview, New York: Cold Spring Laboratory Press.

Kulvicki, J. (2010). Knowing with Images: Medium and Message*. *Philosophy of Science*, *77*(2), 295–313. doi:10.1086/651321

Mäki, U. (2009). MISSing the World. Models as Isolations and Credible Surrogate Systems. *Erkenntnis*, *70*(1), 29–43. doi:10.1007/s10670-008-9135-9

Olby, R. (1974). *The Path to the Double Helix: The Discovery of DNA*. London: MacMillan.

Piccinini, G., & Scarantino, A. (2010). Computation vs. information processing: why their difference matters to cognitive science. *Studies In History and Philosophy of Science Part A*, *41*(3), 237–246. doi:10.1016/j.shpsa.2010.07.012

Scarantino, A., & Piccinini, G. (2010). Information without truth. *Metaphilosophy*, *41*(3), 313–330.

Suárez, M. (2004). An Inferential Conception of Scientific Representation. *Philosophy of Science*, *71*(5), 767–779. doi:10.1086/421415

Teller, P. (2001). Twilight of the perfect model model. *Erkenntnis*, 393–415.