MAPS, PICTURES, AND PREDICATION

JOHN KULVICKI

Dartmouth College

Failing to indicate the presence of something in a map is tantamount to indicating its absence. Blue indicates water, and a lack of blue suggests a lack of water. No lines for highways on part of a map, which can otherwise indicate highways, indicates a lack of highways in that area. Michael Rescorla (2009) calls this the absence intuition, and claims it shows that maps cannot employ predication as languages do. This paper offers a new account of maps that respects the absence intuition without abandoning predication. Maps, pictures, and diagrams differ from language not in whether they involve predication, but in how they organize predicates. Maps introduce predicates holistically, in groups, as degrees of freedom to which any location on a map must commit. This proposal uncovers norms for mapmaking, leads to the first new semantics for maps since Roberto Casati and Achille Varzi (1999), and offers a new perspective on how maps relate to pictures. Maps and pictures are alike not just in the way they represent space, but also in that they both introduce predicates holistically. This proposal relates in interesting ways to John Haugeland's (1991) attempt to understand representational kinds in terms of features of their contents.

1. Introduction

Maps are picture-language hybrids. Their linguistic aspects are found not just in their labels—Topeka, Rest Stop—but in the way small numbers of features pair, fairly arbitrarily, with contents. Blue for water. Green for land. Black and white could serve the same purposes, just as mountain chevrons could represent swamps. The features covering maps are like lexemes of a language that play a predicative role. The map says, of a particular spot, that it is populated, mountainous, and crisscrossed by roads. Another map could say the same thing in another language, as it were, using different markers for each of those qualities.

Maps are pictorial in their use of space. The blue-green pattern pairs with a layout of land and sea by mimicking it spatially. A variety of shapes can represent

Contact: John Kulvicki <kulvicki@dartmouth.edu>

places at extreme latitudes, for example, but not anything goes. Maps' systematic recruitment of space is also found in pictures, where variety is constrained by rules of projection. These constraints are rough—sketched maps are as ordinary as sketched portraits—but they are important and they have no analog in language.

This plausible outline of maps' dual nature is incomplete because predication in maps departs, somehow, from predication in language. In maps, failing to indicate the presence of something—a town, a river, a road, or an island—is tantamount to indicating its absence. Linguistically committing to a town being somewhere does not commit one to towns being anywhere else. Michael Rescorla (2009) calls this the *absence intuition*. It is stressed by most philosophers who have written about maps, including Roberto Casati and Achille Varzi (1999: ch. 11), who have worked out the most detailed semantics for maps so far.

Rescorla (2009) argues that the absence intuition shows maps cannot be modeled using a Frege-Tarski approach to predication. He offers a different model, according to which absence of colors and other marks carries representational weight. This fits well with Casati and Varzi's map semantics, and Rescorla suggests it makes maps akin to pictures and many diagrams. Since a Frege-Tarski account works well for languages, he suggests he has found an interesting point at which representational kinds divide.

The following account of maps respects both the absence intuition and a Frege-Tarski model of predication. Maps, pictures, and diagrams differ from language not in whether they make use of predication, but in the way they organize predicates. In maps, predicates are introduced holistically, in groups, as degrees of freedom to which any location on a map must commit. This proposal constitutes a different way of understanding the coarse-grained distinction between maps, pictures, and diagrams on the one hand, and linguistic representations on the other. It leads to a new semantics for maps, inspired by Casati and Varzi's work. And it uncovers important norms for mapmaking. Violation of those norms sheds light on the absence intuition, which is not as universally compelling as some suggest.

Section 2 clears up some terminological issues, clarifies the absence intuition, and presents Rescorla's proposal for how to deal with it. Section 3 develops the idea that maps introduce predicates only as degrees of freedom, and shows how this is compatible with a Frege-Tarski model of predication. With the proposal in hand, Section 4 offers a theory of maps, including a semantics that departs from Casati and Varzi's. Section 5 defends the theory by considering bad maps. They violate important mapmaking norms, and provide an interesting test case for the two views of map semantics on offer.

The final three sections work out salient relationships between maps, pictures, and language. Section 6 shows that pictures satisfy the absence intuition. This leads, in Section 7, to the claim that maps are picture-like in that they introduce predicates holistically, but language-like in the way they introduce multiple degrees

of freedom. Maps' pictorial aspects go well beyond their use of space. This discussion relates in interesting ways to John Haugeland's (1991) attempt to understand representational kinds. Section 8 then identifies annotation as a distinct phenomenon that involves the full use of language grafted onto a map. Annotations do not respect the absence intuition and they need not be introduced holistically.

2. Locatable Features and the Absence Intuition

Maps tell us where stuff is, in locations we might inhabit. That, at least, captures exemplary instances of maps: city maps, atlases, and sea charts. Less central cases reveal population density, income inequality, oil and mineral deposits, rainfall, porosity, temperature, and anything else that might characterize some region of space. Just how far this goes is an open question. Can we map the ideological terrain? For now, however, it suffices to focus on uncontroversial examples.

The *syntactic features* of a map are those that matter for its representing what it does. They are helpfully distinguished from its *incidental* features, which do not matter in that way, and its *semantic* features (Kulvicki 2014: ch. 5). Semantic features are just the places a map represents and the qualities it represents those places as having. In some maps, for example, blue indicates the sea. Any blue patch indicates a body of water. Being a blue patch is a syntactic feature of the map, and being a body of water is a semantic feature, in that it is a feature the map represents. Admittedly, the phrase, semantic features, is a little awkward—the map is not a body of water—but it is terribly convenient and not apt to mislead.

In some maps, any shade of blue indicates water, so differences of shade are, representationally speaking, incidental not syntactic. In other maps, shades of blue indicate the depth of water, so the specific shades are syntactic, responsible for picking out a range of semantic features. Any particular map is printed on some kind of paper, carved in stone, or displayed on a screen. Those media are typically incidental as well. The spatial layout of a map, as noted above, is syntactically significant. Where things are on a map has consequences for where the map represents things as being.

Some syntactic features of maps can be placed at locations on it. Blue patches, mountain chevrons, and dots indicating towns, are all *locatable features*. Being a certain location on a map is not a locatable feature, even though it is syntactic, because that feature cannot be moved around, or placed at multiple points on a map. It is obviously possible to change around a map's locatable features: recolor sections of it, replace some markers with others, redraw the line indicating a coastline. These changes typically alter a map's semantic identity. They affect what it represents. It is also possible to change non-locatable features: cut a hole out of it, snip off its corners, extend a portion of it in some direction. These changes, too, affect its semantic identity. The altered map has a different content.

Maps' locatable features respect what Rescorla (2009: 182) calls the *absence intuition*. Though not exactly how Rescorla puts things, the following captures the intuition of interest:

If a locatable feature is absent from a portion of a map, the map suggests the absence of the quality that feature represents from the region that the map portion represents.

Coloring an area blue indicates sea somewhere, while not coloring it blue suggests the absence of sea. A mark indicates a town with more than 100,000 residents, and the absence of such a mark indicates the absence of such a place.

Most philosophers who have written about maps share this intuition, though they disagree about how to explain it. Casati and Varzi (1999: ch. 11) and Rescorla (2009) build it into the semantics of maps while Elliot Sober (1976: 132–133), Elisabeth Camp (2007: 163) and Ben Blumson (2012: §9) suggest the phenomenon is pragmatic. That is why it is important to be vague in unpacking the intuition: when a feature is absent the map “suggests” the absence of some quality.¹

There are costs on both sides. Placing the phenomenon in the pragmatics undersells just how ordinary it is. Every map seems to work this way, at least in large measure, which makes it a stretch to insist that the absence intuition results exclusively from the pragmatics of map interpretation. On the other hand, insisting absence is semantic can require awkward interpretations of locatable features.

For example, nautical charts indicate the presence of structures, like houses, along the shore.² The natural thing to say is that the house symbol stands for houses. Read this way, however, the symbols do not respect the absence intuition. Absence of a house marker does not indicate the absence of a house because the markers are only used for structures visible from the sea, some conspicuous, some not, that might be relevant for navigation (US Department of Commerce and US Department of Defense 2013: 6, 26). These charts also indicate submerged obstacles that might threaten a ship. Not all threats are known, however, so some charts indicate general areas in which there might be threats. Even when potential obstacles are known, they distinguish between those whose depths are known and those whose depths are not (2013: 52–56). These charts only respect the absence intuition if the qualifications—visible from the sea, navigationally relevant, of unknown depth—are read into the semantics of the marks. The symbol does not stand for a house, but for a navigationally relevant one visible from the sea. That is enough to push many toward a pragmatic understanding of the absence intuition.

1. Rescorla points out that others, including computer scientists, also endorse the claim. Ian Pratt (1993: 81) is clear about it, and see Schlichtmann (1979), Hayes (1985), and Lemmon and Pratt (1999).

2. Dominic Lopes suggested this example.

This paper makes a simple suggestion about how maps organize their locatable features, which casts the absence intuition as a semantic phenomenon. Because the suggestion is quite productive, this paper constitutes an indirect argument for thinking about absence in semantic terms.

Rescorla's semantic account of the absence intuition means that maps cannot be understood on a model of predication, sourced in Frege and Tarski, which works quite well for language (Rescorla 2009: 176–177). “Jim is tall” is true just in case the predicate “is tall” is true of Jim. In that sense, predicates are akin to functions that take objects like Jim to truth values. One might think that placing a locatable feature on a map is like predicating a quality—whatever quality the locatable feature represents—of some location—the one represented by that place on the map. The map says, of the areas represented by each blue region, that they are sea, and the map is true if the locations it represents as sea really are that way. This only captures part of the map's truth conditions, however. A map is inaccurate unless it represents all sea-covered regions as such, because absence of blue suggests the absence of sea. “This contrast between cartographic and linguistic representation poses a challenge to anyone who holds that attaching a marker to map coordinates is the same mode of semantic composition as attaching a predicate to a singular term.” (Rescorla 2009: 182; cf. Millikan 2004: 93)

Placing a locatable feature on a map affects its truth conditions. It does not affect truth conditions, however, in a way that fits with a standard model of predication. The locatable feature says, of the area denoted by that part of the map, that it is sea. But the map also commits to the absence of sea in places where the locatable feature is absent. The goal of the next two sections is to give a model for map semantics that respects both the absence intuition and a standard account of predication.

3. Predication and Degrees of Freedom

Consider a simple system that uses green for land and blue for sea. Representations in this system are 2-dimensional arrays with just two locatable features: blue and green. Being blue is incompatible with being green, so regions of these maps are blue if and only if they are not green. Similarly, being sea is incompatible with being land. Since patterns of land and sea are all these maps represent, they represent a region as sea if and only if they do not represent it as being land.

This system respects the absence intuition. Every part of a representation is filled with a color, and no part can be filled with both. The two colors represent incompatible qualities. The absence of blue in a region suggests the absence of sea, and the absence of green the absence of land. This system also fits a standard model of predication. Presence of a locatable feature at some place amounts to

predicating that feature of the location that place represents. The absence intuition survives because anywhere a map fails to say a location is land, it manages to say that it is sea. So, it is not so much the absence of blue that represents a lack of sea, but the presence of green, which represents land. Rescorla considers representations like this (2009: 185–187; see also Pratt 1993: 85–86), but rightly notes that they do not obviously generalize to most common examples of maps.

Imagine, for example, introducing a texture that indicates mountains. Any region can be a color and a texture at once: there are mountains on the sea floor, after all, and this system can represent that. Nothing about simple incompatibilities between syntactic or semantic qualities renders the absence intuition true here. Why is it that the absence of mountain marks indicates an absence of mountains? No quality incompatible with mountain marks fills the remaining spaces. Put differently, one need not remove anything from the plain green map in order to add the indicators for mountains: nothing is in the way. Cases like this constitute the main reason Rescorla insists that a standard model of predication is inapplicable to maps. Green stands for land, and adding texture indicates mountains. Some feature or other needs to be placed at each location, but not every feature needs to show up everywhere. Locatable features represent qualities where they appear, and the absence of those features indicates the absence of the corresponding qualities.

Perhaps introducing mountain textures is more complicated than Rescorla appreciates, however. Instead of three determinate syntactic features—blue, green, and mountain texture—there are four: plain blue, plain green, textured blue, and textured green. These features stand for, respectively: sea with no mountains, sea with mountains, mountain-free land, and mountainous land. The introduction of mountain marks renders texture syntactically relevant. Specifically, two ways of being textured matter: being plain and being marked. Before mountain marks mattered, texture was incidental, but now both mountain marks and smoothness are syntactic.

Once the mountain marks are relevant, something like a dimension of variation or degree of freedom starts to matter syntactically. What Rescorla calls the absence of a feature is better understood as the presence of an incompatible texture, as it were, smoothness. Before texture got added to the map system, the pattern within a color patch was syntactically and hence semantically irrelevant. Now it matters, for every patch.

The enhanced map, so understood, fits a standard model of predication. Plain green says of a region that it is relatively flat land. Textured blue indicates ocean over mountainous seabed. The map also satisfies the absence intuition. Absence of mountain marks somewhere indicates a lack of mountains, just as lack of blue indicates the absence of sea. In both cases, however, it is the presence of competing and incompatible qualities that accounts for the absence intuition. A smooth texture indicates the presence of relatively level land, and no part of the map can be

marked with both a smooth and a rough texture. Similarly, no part of the mapped terrain can be relatively flat and also mountainous.

In the system so described, it still makes sense to say that greenness represents land, blueness sea, and texture mountains. But greenness, for example, is always an abstraction from the most determinate syntactic features of any given part of the map. A map region is either green and smooth or green and textured. In either case, the green parts of a map represent land. Smooth green is not agnostic about mountains, however, because smoothness represents flatness.

This strategy generalizes well. Locatable features can be added in a way that respects the absence intuition, while keeping in line with a standard model of predication.

New locatable features are either compatible with the others or not. Adding incompatible features is simple as long as they satisfy the following *incompatibility constraint*:

incompatible locatable features represent incompatible qualities.

Land and sea are incompatible, for example, as are flat and mountainous terrain. Any system of maps that does not conform to this constraint has trouble representing situations in which the represented qualities overlap. Section 5, on bad maps, considers such cases. For now, assume that new incompatible locatable features represent incompatible qualities. They are unproblematic because, as we have seen, they simply fill out a space of incompatibles. The presence of one of these qualities excludes the others, and because of that the absence intuition is respected. Absence of blue requires presence of green. If we add red to represent marshland, the situation is more complicated but not fundamentally different.

The difficult part is adding more and more mutually compatible features, like markers for towns, roads, and geological formations. Adding a new feature, compatible with all the rest, amounts to adding a dimension of variation, or degree of freedom. Within the degree of freedom, all locatable features are mutually incompatible, and between degrees of freedom, features are generally compatible. Mountain marks pair with smooth texture as mutually incompatible, but syntactically significant, aspects of a map. Once mountain texture is on the menu, it is easy to add more textures for different kinds of land: alps, piedmont, hills, bumps, etc. Each of those textures is incompatible with the others, and what each represents is incompatible with what the others represent. Untextured, smooth areas are the zero value along this degree of freedom. Being smooth carries representational weight just as the marks do.

Different road markers are mutually incompatible, and they indicate different, incompatible kinds of roads. Even in the simple case, in which there is only one kind of mark for roads, this degree of freedom has a zero value, which corresponds

to no roads being in the area. Markers for towns work this way as well. Different marks for towns may not occupy the same place, but a town marker can otherwise be placed anywhere. Once town markers are introduced, that degree of freedom has a syntactically significant zero value, in which no marker is present. An area of a map unmarked by town or road markers looks just like an area of the simple blue-green map, of course. But in the simple map that area does not manifest a zero value along the road- or town-marker degree of freedom.

Rescorla suggests that maps are distinctive because of the way they handle absences of syntactic qualities. Failure to indicate something is *eo ipso* the indication of its absence. The current strategy, by contrast, insists that the absence of a syntactic quality should be understood in terms of the presence of an incompatible one, along a degree of freedom. Casting maps in this manner preserves the key aspects of the standard model of predication. To know what a map says of a given area, just list the values it takes along all its degrees of freedom. Those values predicate, of the region represented, that it has the qualities in question. The map is true just in case each region has the qualities it is represented as having. Maps are quite different from other kinds of representation, however, in the way that predicates are introduced. The cost of introducing a new predicate in a map is either refining an existing dimension of variation or enfranchising a new one.

4. A Theory of Maps

A *mapping system* is a way of filling arrays with locatable features. The focus here is on locatable features and how they figure in predication. For present purposes, then, arrays are sets of locations, typically in two dimensions, that represent regions of some other space. Typically, neighborhood relationships among an array's locations are matched by those among the locations represented by that part of the array. Sometimes, stronger constraints are in place, but for now this general picture suffices (see Casati & Varzi 1999: 191).

The locatable features are grouped into *incompatibility classes*: classes of mutually incompatible members. For example, if colors are used, they all go into a class, because no map region can be two colors at once. Different textures for hilliness wind up in their own class, just as road markers and town markers each get their own. Locatable features in different incompatibility classes are typically compatible with one another. If they are not, we run the risk of having a bad map.

Incompatibility classes reflect the mapping system's *degrees of freedom*. A mapping system is a way of *filling* arrays in that every location is assigned a value from each degree of freedom. A value is either a member of the corresponding incompatibility class, or, sometimes, a null value.

It is tempting to identify degrees of freedom with incompatibility classes, but

doing so depends on being quite catholic about qualities. The null value of the mountain marks incompatibility class is not smoothness, if we understand smoothness to involve the lack of all texture. Instead, the null value is failing to have any of the mountain marks. Things get even more complicated when the incompatibility classes are stranger, like the marks that indicate varying kinds of towns. The most convenient way to think about this is that readily identifiable locatable features constitute an incompatibility class. Maps assign a value for each degree of freedom, and that value can be zero. Whether those zero values are properties matters less than the fact that having a zero value for that degree of freedom is syntactically significant. Having no mountain texture represents smooth terrain. Having no city mark represents relatively uninhabited land.

Incompatibility classes need not include qualities that relate to one another along physically interesting dimensions. It can be a matter of convention, not physics, that the qualities in a class are incompatible in the first place. Shades of color are useful locatable features because they are mutually incompatible and saliently ordered with respect to one another. A temperature map is easy to read if it uses shades of warm color ordered from dark to bright to indicate increasing temperature. But there need be no interesting sense in which small solid dots, annuli, squares, stars, and triangles are values along a dimension. In fact, it might even be physically possible for a couple of town markers to occupy the same location. Norms governing mapmaking can prohibit their co-occurrence, even if their physical features do not. The null value for the degree of freedom corresponding to such an incompatibility class is similarly gerrymandered: not having any of the town mark features. Compatibility is often as conventional as incompatibility. One cannot add visible texture to a flat surface without changing aspects of local color, but it makes perfect sense to say that mountainous texture and the colors for land and sea cohabit.

Thought of in this way, maps respect the absence intuition. Lack of a feature indicating mountains indicates lack of mountains, lack of a road or town marker indicates a lack of roads, or towns, and so on for all of the degrees of freedom. Moreover, satisfying the absence intuition this way is compatible with the Frege-Tarski model of predication. Maps are all like the simple blue-green example with which we started. In that case, absence of land was indicated by the presence of a competing color, which stands for something incompatible with land. In general, lack of a feature amounts to presence of a competing value along some degree of freedom, even if that value is zero. If we take a region of a map to pick out a particular region of space, the values for each degree of freedom assign features to that region. The map expresses something like a function from objects—represented regions of space—to truth values.

It should not seem terribly strange that degrees of freedom can have syntactically significant zero values. Even representations quite distant from maps recruit

zero values in syntactically significant ways. A red light indicates the dishwasher is running and a green light that it is finished. Alternatively, a red light indicates run time and no light stands for the final state.

Rescorla considers and quickly rejects a version of this strategy. He imagines that each locatable feature is paired with a “hidden marker” that indicates absence of whatever the feature represents. Mountain marks are paired with a hidden no-mountains marker, for example. He wonders “why we should clutter our theory with so many hidden markers just to preserve [a Frege-Tarski model].” In addition to being ugly, Rescorla suggests the move is unmotivated. “If applying a marker to map coordinates is the same semantic operation as applying a predicate to a singular term, why does it fall under such special restrictions in the cartographic domain?” (2009: 187)

The present proposal is systematic and uncluttered. Locatable features sort into incompatibility classes, and some degrees of freedom need a zero value. This renders predication in maps holistic in an interesting way, distinct from language, but nevertheless preserves a model of predication that is language-like. Predicates are not added piecemeal, but *en bloc*. Added motivation for this proposal comes from the way it helps us understand the relationship between maps and pictures, but that will have to wait until Sections 6 and 7.

This general strategy suggests a semantics for maps that is a variant on the one offered by Casati and Varzi. They claim maps are composed of *atomic map stages* (1999: ch. 11). An atomic map stage is an array like the original, but which only contains a single locatable feature—blue, for example—exactly where it is on the map as a whole, and nothing elsewhere. Another stage contains green, another mountain marks, another large city markers, etc. (They limit their presentation to colors but the strategy for extending their discussion should be clear.)

An atomic map stage says that some quality is present at certain places, specifically those denoted by the parts of the map where the locatable feature is found. The stage is true just in case the quality indicated by the locatable feature is present where the stage says it is and absent everywhere else. That is, an atomic map stage is false if it locates a quality where it is not or fails to locate a quality where it is. The conjunction of all atomic map stages yields the truth conditions for the map as a whole. This is an elegant and effective strategy for thinking about the compositional structure of maps, but it is not quite compatible with what has been said here.

The lesson of this section has been that locatable features sort into incompatibility classes, corresponding to degrees of freedom. So, instead of isolating individual locatable features, a map stage is an array that replicates just the layout of the locatable features within a degree of freedom. There is no reason to call such a complex thing atomic, but that is in line with the thought that predication happens

holistically in maps. That stage is true just in case the qualities indicated are where they are said to be. Those degrees of freedom that include a zero value indicate the absence of any of the other features within that degree of freedom where they zero out. Those degrees of freedom that do not require a zero value are those that fill every space with some among a family of incompatible qualities, like fields of color. Each map stage is a filled array, but an array filled by only one degree of freedom. The conjunction of all map stages yields the truth conditions for the map as a whole.

The Casati-Varzi semantics agrees with the one offered here in almost all cases. They come apart in very specific contexts, which speak in favor of the present proposal. It is time to think about what happens when maps go bad.

5. Good Map, Bad Map

Section 3 bracketed cases in which maps do not satisfy the incompatibility constraint: incompatible locatable features represent incompatible qualities. Though maps need not satisfy this constraint (Rescorla 2009: 185), those that fail to do so run afoul of a simple norm, *Coverage*, to which maps aspire:

For any quality P that a mapping system can represent, a map of any region can represent the full extent of P in that region.³

Coverage does not prohibit false maps. Maps can easily locate qualities in places they are not found. The norm only insists that maps be able to represent the full extent of qualities they can represent. This condition entails the incompatibility constraint. If incompatible locatable features represent compatible, but distinct, qualities, then a region where those qualities overlap is such that at least one of those qualities cannot be represented across its full range.

The Casati-Varzi semantics coincides with the present one for all mapping systems that satisfy Coverage. In such systems, if two locatable features are incompatible, then they represent incompatible qualities. Degrees of freedom, then, represent sets of mutually incompatible features. In that context, one can break up a map stage as understood here into Casati-Varzi atomic stages without changing the truth conditions of the whole. An atomic stage says that some quality is found in the places represented by regions occupied by a locatable feature, and not elsewhere. The not-elsewhere part of their semantics is handled in the current view by

3. Thanks to an anonymous referee for convincing me that this simple formulation of Coverage does the work I need it to do.

the fact that every other location is filled with some value of the degree of freedom representing an incompatible quality.

Maps that fail to satisfy Coverage pull apart the semantics offered here from Casati and Varzi's, and they also undermine the absence intuition. Bad maps thus provide interesting test cases for the theories.

Imagine a topographical mapping system that represents conifer forests with fields of green and ranges of crowned kinglets with regions of red, as in Figure 1 below. The birds occupy conifer forests, but other areas as well, so these qualities are compatible, even though red and green are not. Also assume that the system has no third feature—striped red and green, for example—to stand for conifer forests inhabited by the birds. This system works well for areas in which there is no overlap between kinglets and conifers, but it runs into trouble where they cohabit because it violates Coverage. Tasked to represent overlapping qualities with incompatible locatable features, a mapmaker must choose between indefinitely many options, none of which represents the kinglets everywhere they are and the conifers everywhere they are.

Imagine that our mapmaker produces Figure 1. According to Casati and Varzi—and Rescorla, who endorses their semantics—no matter what she does the map will be false. By hypothesis, kinglets and conifers overlap in this region. So, the atomic map stage for at least one of the colors fails to place the feature where a quality resides, and thus wrongly suggests absence. Maps that do not satisfy Coverage are simply unable to represent many situations accurately. Given the semantics offered in the previous section, however, the map can be true, despite being bad. One stage of the map consists of all the incompatible colors and their zero value, and that stage is true just in case all of the represented properties are where the stage says they are. Assuming none of the colors or the zero value represents anything as being anywhere it is not, the map stage is true. The map is incomplete because the system to which it belongs does not satisfy Coverage, but it is not false for all of that.

Deciding which semantics to endorse depends on what happens to the absence intuition in cases where Coverage is violated. If it remains in full force, Casati and Varzi give us the right answer: the bad map cannot accurately represent the region in question. If bad maps undermine commitment to the absence intuition, however, the semantics offered here gives a nice explanation of why.

The mapmaker is in a bind. She wants to represent everything everywhere it is, but she is hamstrung by a poor mapping system. At best, she can avoid representing conifers and kinglets where they are not and be sure to represent at least one of these wherever they overlap. Imagine that she does this, and moreover that she does it for a well-informed audience. Consumers of her map know that red and green are incompatible, while kinglets and conifers are not. They do not know

Figure 1. Kinglets are compatible with conifers, but the red color in this map is incompatible with green, resulting in a bad map. The topographic map source is a detail from US Geological Survey (2012).

whether kinglets and conifers overlap in the region she has represented, but they know that no such map could tell them whether they do.

In such circumstances, the absence intuition loses its force. Standards for interpreting these representations are clear, but they do not respect absence. The mapmaker has reasonable intentions to be truthful, and her audience knows that she cannot represent the full range of both qualities if they overlap. So, absence of red, where green happens to be, is not taken to represent absence of kinglets. Absence of green, where red happens to be, is not taken to represent absence of conifers. Despite that, this representation is a map. It is just a bad map, and it engenders agnosticism, not disbelief. There might be kinglets in that conifer forest.

Charity runs deep. Everyone wants maps to be good, so one might try to reinterpret the colors in such a way that they satisfy Coverage. Perhaps red represents important areas of kinglet activity and green represents conifer stands without many kinglets. The mapmaker might even have allowed such reinterpretations to color her choices. But that just reveals a desire to have maps live up to certain norms. When they obviously do not, they do not respect absence. The map is incapable of representing all of the kinglets and conifers if they overlap, so it comes across as incomplete in a manner we prefer maps not to be.

In sum, the goal is to provide an account of maps that respects the absence intuition while being compatible with a Frege-Tarski model of predication. The proposal is that predicates are not added to maps piecemeal, but either as new degrees of freedom or extensions of already existing ones. Each region of a map commits to some value—sometimes zero—along each degree of freedom. Thinking about predicates this way suggests changing Casati and Varzi’s semantics for maps. The new semantics coincides with the old for all mapping systems but those that run afoul of Coverage. Maps that do not respect this simple norm are poor maps, but great test cases for the new semantics. These maps do not elicit a strong absence intuition, and thus are readily regarded as incomplete, rather than inaccurate. The new semantics provides precisely such truth conditions for bad maps, while the old one does not. An added bonus to this way of thinking about maps is that it shows how maps relate to other kinds of representation, like pictures, images, diagrams, and graphs. They all add predicates as degrees of freedom, as the next section shows.

6. Pictures and Absence

Something like the absence intuition applies to picture interpretation. A picture representing blue sky over the Green Mountains is blue where it represents the sky and green where it represents mountains. Absence of green indicates the absence of mountains, while absence of blue tracks absence of sky. Though plausible enough, two facts about pictorial content threaten this claim.

First, local colors represent aspects of surfaces and the way they are illuminated. Dark regions of pictures can depict dark well-lit surfaces or shadowed light ones. Absence of a light color at a certain point need not represent the absence of a light-colored surface there.

Second, pictures represent things as occluding others. Mountains do not preclude the presence of sky beyond, and so it is hasty to claim that the absence of blue indicates the absence of sky past the mountaintop. In fact, Ned Block (1983), and following him, Flint Schier (1986: ch. 4) and Dominic Lopes (1996: 118–119) suggest that pictures are distinctive representations in part because they make *explicit non-commitments*. Unlike linguistic representations, pictures can bar themselves from committing to something’s presence by committing to the presence of something else that obscures it. A picture of someone wearing a big hat explicitly does not commit to that person’s hair color, for example. Pictures are readily interpreted as incomplete in that the absence of a quality representing hair, in this case, does not in any way indicate hair’s absence.

Illumination and occlusion are examples of what Ernst Gombrich called the “ambiguities of the third dimension” (1961: ch. 8). They are features of things in

space, not two-dimensional arrays, so only a complex relationship between surface features and content allows pictures to represent things in the round. Semantic complications like this make the absence intuition less salient for pictures than it is for maps, but pictures respect it nonetheless. Seeing this requires thinking about pictorial semantics in light of pictorial syntax.

Syntactically, pictures are filled with a limited palette of locatable features. At each location, a picture manifests some color or other. They sort into an incompatibility class, in that the presence of a color somewhere precludes the presence of any of the other colors there. Alternatively, the colors sort into three incompatibility classes: hue, saturation, and brightness. Syntactically speaking, pictures seem like the simple blue-green map discussed in Section 3. They place one locatable feature at each location. Some pictures might work a bit differently, but for now the point is whether these central cases satisfy Coverage, or at least the incompatibility constraint. This, in turn, shows how they respect the absence intuition.

Pictures might seem to violate the incompatibility constraint because surface colors are mutually incompatible, but different surface colors can represent identically colored surfaces. That is, two incompatible colors of a picture surface—light red and black, for example—can represent identically colored things: one well-illuminated, one in deep shadow. But this is not a violation of incompatibility. Pictures' surface colors correspond to complexes of surface color and illumination. A shadowed red surface is incompatible with one well-lit. The contents associated with pictures' locatable features are, in that sense, a bit more specific than one might have thought.

One might also worry that pictures violate incompatibility because one and the same surface color can represent rather different things, for example, a red surface illuminated by white light or a white surface illuminated by red. Indeed, whether a local patch of red represents one kind of thing or another often depends heavily on the other patches of color that make up the picture. Ambiguities like these do not make any obvious trouble for incompatibility, however, because they do not suggest that incompatible features are tasked with representing compatible qualities. In this sense, the contents associated with pictures' locatable features are less specific than one might have thought.

Insofar as they respect the incompatibility constraint, pictures also respect the absence intuition. Absence of a quality like bright red at some location indicates the absence of any of the complexes of surface color and illumination represented by that quality. Yes, in pictures bright red might stand for red surfaces, white surfaces, and even other colors illuminated oddly, but those indeterminacies do not interfere with the absence intuition. Each locatable feature can stand for a distinctive range of colors and illuminations, so absence of a feature indicates absence of that range of possible contents. The semantic complexity of pictures works within constraints that respect absence.

Occlusion is the other threatening complication. Explicit non-commitment results from the fact that, in the round, one object can stand in front of another. Pictures commit explicitly to the facing surfaces of things and how they are illuminated, and they respect absence exclusively in terms of those explicit commitments. Absence of blue where there are mountains suggests absence of any of the things a shade of blue might represent, as far as the picture can explicitly commit to things. Similar worries might be raised for maps, of course—the antipodes of New England is sea—but neither maps nor pictures are held to these constraints with respect to explicit non-commitments.

John Haugeland (1991, anthologized 1998) found it theoretically valuable to distinguish skeletal from fleshed out contents of representations. Skeletal contents are quite indeterminate. As we saw earlier, local color in a photo can represent any number of worldly features. The skeletal content is what all of those potential scenes have in common. Haugeland thought that by stripping down content to its simplest form, he would find what makes kinds of representation distinct from one another.⁴ He suggests, for example, that all a photo represents skeletally “is certain variations of incident light with respect to direction” (1998: 186). His idea is not that pictures represent light, *per se*, so much as that they represent patterns of color across two dimensions. These patterns can be due to variously colored surfaces, lights, haze and combinations thereof. In short, Haugeland thought that the skeletal contents of pictures were just what we arrived at in showing the respects in which they satisfy the incompatibility constraint and absence intuition.

A picture’s fleshed out content is, for example, mountains and blue sky from a certain perspective. It counts as fleshing out a skeletal structure because it identifies a relatively more specific situation, consistent with the pattern of color, as the content. There are many ways of fleshing out any skeletal content, and norms govern how it is done. For example, a flat plane colored just like the picture being interpreted manifests a pattern of color consistent with the picture’s skeletal content, but it is rarely appropriate to interpret pictures as representing nothing more than colored planes (Kulvicki 2006: ch. 9).

For now, the dynamics of fleshing out are beside the point. Thinking about whether and how pictures satisfy the incompatibility constraint leads to skeletal content in a way no one has previously considered. Skeletal content also makes it easy to see how pictures respect the absence intuition. They have colors everywhere, and any given color excludes the rest. Each color makes a distinctive contribution to skeletal content, so the absence of a color from a place indicates the absence of the related aspect of content from that place. Because we usually think

4. Elsewhere, I suggest that the idea is impressively useful, though not quite in the way Haugeland envisioned (Kulvicki 2006, chs 3–4). The present discussion constitutes another reason the skeletal/fleshed-out distinction is valuable.

about pictures in terms of their fleshed out contents, not their skeletal contents, absence is a much less salient feature of them than it is of maps.

Finally, skeletal content makes it easy to see how pictures work, predicatively, like maps. Pictures are arrays filled with locatable features, which sort into an incompatibility class. These features are paired, skeletally, with contents. Unlike maps, pictures do not readily admit new incompatibility classes of locatable features. As we will see, this layering of incompatibility classes in maps and diagrams is part of what makes them language-like.

Not all pictures work the same way. Color photos, for example, have different skeletal and fleshed out contents than black and white photos. Some paintings and drawings make use of highly realistic techniques, while others are sketchy. An anonymous referee pointed out that if we mix such styles—a drawing that is detailed here, sketchy there, colored here, black and white there—there is a sense in which those hybrids do not respect the absence intuition. Absence of color does not indicate an absence of represented color in the black and white portion of the picture. Locally, of course, they do. Within a portion of the picture that respects a given set of standard of interpretation, absence is respected, though between them it is not.

7. Maps, Pictures, and Language

Haugeland's goal was to understand kinds of representation. He suggested that *icons*, which include pictures, diagrams, graphs, and maps, have different kinds of skeletal contents than *logical*, or linguistic representations. Iconic contents, which include the contents of maps and pictures,

might be conceived as variations of values along certain dimensions with respect to locations in other dimensions. Thus, variations of temperature with respect to time, or altitude with respect to latitude and longitude, would be the contents of familiar icons. The former dimensions are called dependent, the latter independent, because, for each point in the relevant independent region, a point is determined in the dependent space, but not necessarily vice versa. (Haugeland 1998: 192)

Different species of this genus make use of different dependent and independent dimensions. Some pictures display black and white alone, while others “have two independent and three dependent dimensions” (1998: 192). Graphs of temperature over time have one independent dimension, representing time, and another dependent dimension representing the temperature. Maps' contents can be very rich, stretching to tens of dependent dimensions, depending on the system one uses.

Haugeland says nothing about the syntactic features of representations, and the way they bear semantic weight. He even suggests that the important distinctions between representational kinds “are to be found in what is represented—the contents” (1998: 174) rather than anywhere else. Even if Haugeland is wrong on that score (for an argument, see Kulvicki 2006: ch. 6) there is something right about his characterization of iconic content.

This paper, along with Casati and Varzi and Rescorla, keeps syntactic qualities and compositional mechanisms in mind. How do syntactically and semantically significant parts of a representation combine to yield a meaningful whole? The accounts of maps described here can all agree with Haugeland’s characterization of iconic content (cf. Rescorla 2009: 177 n. 3) but they disagree about the compositional mechanisms that yield such contents. For Rescorla, and Casati and Varzi, locatable features can be introduced to maps one by one, but they do not act as predicates in a language. Their absence plays an interesting representational role. Because of that, they can say that map contents are unpacked in terms of dimensions. On the current account, locatable features can only be added as components of degrees of freedom. Every location on a map instantiates some value for each degree of freedom and locatable features can thus function as predicates.⁵

It should be clear by now that maps are picture-like not just insofar as they are spatial arrays that represent things in space. Their predicates also fit into classes, which respect the incompatibility constraint and the absence intuition. Pictures are distinct from most maps and diagrams in that they use relatively few sets of locatable features, paired in a fairly flatfooted manner with skeletal contents. Colors in pictures represent color and illumination. In fact, most of the philosophical action in understanding pictures has focused on the way their simple surface features pair perceptually with rich pictorial contents (Kulvicki 2014). For all of that, however, pictures are severely limited in the range of things they can represent. By and large, they represent potentially visible scenes. Even those who suggest that there are non-visual pictures—Lopes (1997), Hopkins (2000), Kulvicki (2006: ch. 5)—think that tactile or auditory pictures represent tactile or auditory qualities.

A simple map, which uses only two colors—for land and sea—and a single mark for roads, has more degrees of freedom than the typical color photograph. Maps are highly conventionalized in the way they make room for various degrees of freedom: compatibility is more a matter of convention than any physical fact. The challenge in producing maps is in part deciding which conventionalized (in) compatibilities can be readily understood. This is generally a non-issue for pictures.

5. My proposal also fits well with Haugeland’s characterization of logical contents, as contrasted with iconic ones: “the primitive elements of logical contents . . . are always identifiable separately and individually. That is, they can enter into atomic contents one by one, without depending on their concrete relations to one another, if any” (1998: 191). That is precisely what maps’ locatable features and their associated contents fail to do.

Figure 2. Detail from Takana and colleagues' (2014) map of Mars, showing the areas around the Viking and Pathfinder landing sites, along with some relevant legend entries.

The photo might make rich use of its one degree of freedom—all colors within it are syntactically important—but no other degrees of freedom operate. Maps are colored planes, much like pictures, but granted a more complex syntactic structure.

Added complexity, but typically enhanced precision, also characterizes map semantics. A degree of freedom can represent just about any range of qualities, and it is easy to know, via a legend, exactly what qualities each feature within it represents. It helps, very much, if the qualities represented by a degree of freedom are mutually incompatible, but they need not be. In addition, the range of qualities represented need not be a physically interesting dimension, let alone one readily perceived.

Geologic maps exemplify the complexity and precision of maps. Consider Kenneth Tanaka and colleagues' (Tanaka, et al. 2014) geologic map of Mars. It deploys at least four degrees of freedom: colors, textures, black line markings, and blue line markings. The colors indicate the age and character of land, the textures are a relief map indicating elevations above a certain threshold via shading, and the black lines indicate specific geologic features that are often not evident merely in the textures. Blue lines are for channels where water has flowed. Figure 2 is a detail from the map, along with some relevant legend entries.

Some of these represented qualities, like age and elevation, fit neatly into physically interesting dimensions. Wrinkle ridges, grabens, and outflow channels, however, are geologically interesting formations that do not fit into easily specified physical magnitudes. Locatable features within any group—line markings, colors, textures—are mutually incompatible and often related to one another along many dimensions. Because geology is not merely a study of surfaces, the map legend explaining the colors also indicates which kinds of land overlay, grade into, embay, and otherwise interact with the one in question. To give a sense of the complexity involved, the description of the orange color marked lHt for late Hesperian transition unit, reads in part:

Superposes units eHt, eHv, eHh, Htu, HNt, Nhu, mNh, mNhm, and eNh; gradational with unit Hve; interfingers with unit Ahv, overlain by units Hto, lHl, AHtu, Aa, mAl, lAv, and the younger part of ANa. (Tanaka, et al. 2014)

Each color indicates a kind of land incompatible with all of the others, and the map legend then indicates relationships between these kinds of land. It takes a practiced eye to tell from the map alone which if any of the areas mentioned in the description above are overlain by lHt. The map represents the surface makeup, and the legend indicates all of these relations.

Between degrees of freedom, all of the locatable features are mutually compatible, and there is even some interesting overlap in their contents. For example, the relief textures often clearly suggest an impact crater. Consider the formation just above the Pathfinder landing site, and others like it, especially in the map's yellow regions. The map also has a line symbol indicating crater rims, which can be seen in the far left center and upper right portions of this detail. Such overlap is harmless, and even helpful, but it does raise a question.

Are the crater-like textures craters or something else? Do the crater line symbols fail to respect the absence intuition? The pamphlet accompanying the map (Tanaka, et al. 2014: 6) explains that the mark is used only for craters whose rims exceed 100 kilometers in diameter. This map thus respects absence, and has two compatible ways to mark the presence of some craters. Marking all impacts with

lines would make for a rather cluttered map of a heavily cratered world. Nautical charts, as we saw in Section 2, often employ symbols for houses that seem to violate the absence intuition because not all houses are represented. But such maps specify that the symbols represent navigationally important houses: those visible from a boat. Respect for absence is built deeply into mapmaking practice.

The point of this extended example is to illustrate the ways in which maps can represent multiple classes of qualities, many of which are far from perceptible, and organized in complex ways with respect to one another. The kinds of qualities one can correlate with an incompatibility class are open-ended, much as the thing or quality represented by a given lexeme is. Likewise, the number of qualities that can work together, in a conventionally mutually compatible fashion, is open-ended. In these senses, maps are like language, while in their use of space and deployment of predicates only within degrees of freedom, they are like pictures. Maps are thus interesting hybrid forms of representation.

The distinction between maps and pictures is quite salient when the Mars map is set against a color photograph, but it is important to note that there are many interesting cases in between the two, from simple maps to syntactically complicated pictures. An anonymous referee helpfully pointed out that some pictures make heavy use of outlines to indicate object boundaries. Being an outline has no significance for the surface colors of the objects depicted, even if the outline itself has a color. So, outline marks are their own locatable features, distinct from those that represent surface color. Also, standard techniques for drawing comics have symbols for representing motion and speech that work independently of their other syntactic features. These pictures exploit the same syntactic tricks that maps do. Whether one calls all such cases pictures, full stop, and whether the central cases discussed above are truly exemplary of the kind, depends on one's theory of depiction. Similarly, there is no need to call comics maps, in large measure because they are used for such different purposes.⁶ The point of this section is to unpack a genetic link between pictures and maps that goes well beyond their use of space.

8. Annotations

It is tempting to refer to any recognizable bit of language on a map as a label, but this confuses two quite different uses of linguistic forms. On the one hand, words can be locatable features. Instead of a man/woman icon, wc marks restrooms, and 'picnic' replaces a sketched table. The Mars map uses letter codes along with each color, as another way of indicating the age and character of the land. In the lower

6. Dylan Horrocks's *Hicksville* (2010: ch. 4) contains an interesting exchange on the relationship between maps and comics. Thanks to Aaron Meskin for the pointer.

middle of Figure 2, Hto stands for Hesperian transition outflow regions, and toward the middle, INh stands for late Noachian highlands. Each color suffices to indicate these features, but the letter combinations are helpful keys for interpreting the challenging range of colors. The map would indicate all of the same features if either color or the labels were removed—they are completely redundant—but it would be an unhelpful map if it lacked either. The color helps users see the shapes of similarly constituted regions as well as the locations of similar, but spatially disjoint areas. The text makes for ready decoding. Understood this way, words or letter strings are simply locatable features belonging to one or more incompatibility classes.

Words also function as bits of language affixed to a map. A title suggests that the map shows Mars, not Cleveland or Venus. Craters and long-dormant volcanoes get names. One can even imagine adding long descriptions: “the landing site of the first Mars expedition, which left Earth on. . . .” These are all *annotations*. Though their locations on maps are often semantically significant, annotations are not locatable features. They do not belong to incompatibility classes or respect the absence intuition. Exactly how they do work is more controversial, but the theory of maps presented earlier offers a way to think about this practice.

All the theories discussed so far treat locations on maps as referring elements. This spot on the map refers to some region of Mars. Locatable features, on the current view, work predicatively, so the map says, of the region referred to by a given location, that it is cratered late Noachian highlands. Annotations work predicatively too. They are attached to spatial parts of the map, and say, of the regions to which those parts refer, that they are thus-and-so. This is true for full-blown descriptions, descriptive elements—rough, dangerous, icy—and proper names.

The tourist map marks a location with a description—an interesting example of ophiolite!—or just: cool ophiolite! If the map is correct, the user is convinced that she will find interesting ophiolitic rocks at that point. There might be other interesting ophiolites around, but lack of this description at each and every one of them is not taken as a problem for the map. A geologic map might make use of a degree of freedom that includes a mark for ophiolites, and in that case failure to place the mark everywhere it is appropriate amounts to a semantic foul. For annotations, absence of a description of some place on a map does not indicate absence of any worthwhile descriptions of it. It also does not indicate that other areas are not the way the described area is. There might be interesting ophiolites elsewhere. Definite descriptions—the site of the Donner encampment—respect absence, but that is not because they work as locatable features, so much as that they only refer if uniquely satisfied. Uses of descriptions are accurate when they accurately describe the place referred to by the location on the map to which they are attached.

Proper names work much like descriptions and descriptive elements in this context. Attach ‘Elm St’ to a road marker and one says, of the region of the world

to which that part of the map refers, that it is Elm Street. ‘Cleveland’ says, of the referent of a map region near the name, that it is Cleveland. Since reference is secured by the location, it is like saying “This is Cleveland” or “This place is called Cleveland.” Other places might also be called Cleveland.⁷ Two Kansas Citys are represented by many maps, but there is no foul in labelling only one of them. One may not conclude from the fact that someplace is named Kansas City that other places are, or are not, so named. Names can even annotate the map as a whole, as when the Mars map is labelled Mars. Referential uses of names are common, and maps are sources for many of them. Cleveland is north of New York, for example. That is not to say the map uses names referentially.

Sometimes it can be difficult to decide whether an aspect of a map is an annotation or a locatable feature. The previous section claims that the Mars map has *at least* four degrees of freedom, because it is an open question what to do with the Pathfinder and Viking mission marks. The crosses seem like locatable features, even if the labels are annotations. Treating them as such commits the map to completeness with respect to Mars landing sites. The red crosses are, in that case, part of a fairly slim degree of freedom. Since there are few landing sites—seven altogether—one might just regard the red crosses, along with the names, as annotations.

This way of thinking about annotation fits well with the semantics for maps offered earlier. The referring elements are locations, and both the locatable features and the annotations act predicatively. There are, doubtless, other ways of doing things, but so far not much at all has been said about annotation in maps and any theory of maps had better say something. Perhaps the most controversial part of the proposal is that proper names are used predicatively when they annotate. But while there is much controversy over predicative accounts of names—see, e.g., Bach (2002) and Fara (2015)—there is little over whether names can be used predicatively.

Annotation cuts across all representational practices. One can annotate diagrams, graphs, and, of course, text. Locations of annotations often matter, even when text is annotated, because the location of the annotation establishes the part of the representation, or the aspect of its content, upon which it comments.

9. Conclusion

Maps borrow much more from pictures than the way they represent space. Like pictures, they acquire predicates as degrees of freedom, not individually. This accounts for the absence intuition in a way compatible with an ordinary, Frege-

7. It sounds like a great name for a philosopher’s city in which all distinctions are made, and matter.

Tarski model of predication. And it cuts representational kinds in two. Some take on predicates individually, while some take them on as dimensions of variation. Maps, diagrams, graphs, and the like fall on one side of this divide, while, broadly speaking, linguistic representations fall on the other. This claim is not, as Hauge-land suggested, purely one about content, but one that helps us understand the compositional structure of these representations. That is, the semantics are reflected in the syntax.

Pictures come across as a bit primitive seen with maps in mind. Central examples admit a single degree of freedom, which is readily perceived and understood. Moreover, this degree of freedom pairs in a rather simple way with the qualities pictures represent, at least if we focus on their bare bones contents. Colors pair with combinations of surface color and illumination. Pictorial contents are fleshed out to detailed scenes thanks to the close relationship between pictures and our perceptual systems. Perception is the source of pictures' rich, fleshed-out contents, and a source of philosophical consternation, but also a profound limitation. Pictures represent perceptible scenes, albeit in impressive detail.

Maps profoundly sharpen, and then expand, the expressive power of pictures. They sharpen it, in that members of incompatibility classes can pair precisely and comprehensibly with represented qualities. A color patch on a picture stands for any of an open-ended set of situations, while the green on a map means land, or late Hesperian lowlands. This dot means a city between ten and one hundred thousand inhabitants; that squiggle a tectonic contraction. Maps expand pictures' expressive power, first, by detaching color patches semantically from readily perceived states of affairs. Now, colors can stand for any features one likes. Second, they allow liberal layering. Many degrees of freedom can, at once, collaborate to communicate rich contents. Both the feature-quality pairings and their mutual (in) compatibilities are highly conventional, and in these respects maps borrow a page from language. Camp suggests that maps "can easily represent the locations of and relations among very many objects and properties in an explicit yet cognitively transparent way . . ." (2007: 162). Hopefully, the foregoing has gone some way to unpacking just how maps manage to do this.

Acknowledgements

Versions of this paper were presented at the *Shaping the Trading Zone* conference at Leeds, organized by Steven French and Otávio Bueno, and the London Aesthetics Forum workshop, *The Aesthetic Aims of Science*, organized by Stacie Friend and Mauricio Suárez. Murat Aydede and Dom Lopes organized an afternoon workshop on the paper in Vancouver, and my colleagues at Dartmouth showed up in force for a workshop just as our fall term was starting. My colleague David Plun-

kett read a draft even though he was on leave, and Liz Camp and Roberto Casati sent excellent comments. Ken Tanaka was good enough to answer some questions about his impressive Mars map. The referees for *Ergo* were prompt, charitable, careful, and helpful, and Robert Mason provided good copyediting advice. I am grateful for all of the assistance, and the paper is much better because of it.

References

- Bach, Kent (2002). Giorgione was So-Called Because of His Name. *Philosophical Perspectives* 16(1), 73–103. <http://dx.doi.org/10.1111/1468-0068.36.s16.4>
- Block, Ned (1983). The Photographic Fallacy in the Debate about Mental Imagery. *Nous* 17(4), 651–661. <http://dx.doi.org/10.2307/2215087>
- Blumson, Ben (2012). Mental Maps. *Philosophy and Phenomenological Research* 85(2), 413–434. <http://dx.doi.org/10.1111/j.1933-1592.2011.00499.x>
- Camp, Elisabeth (2007). Thinking with Maps. *Philosophical Perspectives* 21(1), 145–182. <http://dx.doi.org/10.1111/j.1520-8583.2007.00124.x>
- Casati, Roberto, and Achille Varzi (1999). *Parts and Places*. MIT Press.
- Fara, Delia Graff (2015). Names are Predicates. *Philosophical Review* 124(1), 59–117. <http://dx.doi.org/10.1215/00318108-2812660>
- Gombrich, Ernst (1961). *Art and Illusion*. Princeton, Phaidon.
- Haugeland, John (1998). *Having Thought*. Harvard University Press.
- Haugeland, John (1991). Representational Genera. In William Ramsey, Stephen Stich, and David Rumelhart (Eds.), *Philosophy and Connectionist Theory* (61–90). Lawrence Erlbaum. (Reprinted in Haugeland 1998, 171–206.)
- Hayes, Patrick (1985). Some Problems and Non-Problems in Representation Theory. In Ronald Brachman and Hector Levesque (Eds.), *Readings in Knowledge Representation* (3–22). Morgan Kaufmann.
- Hopkins, Robert (2000). Touching Pictures. *British Journal of Aesthetics* 40(1), 149–167. <http://dx.doi.org/10.1093/bjaesthetics/40.1.149>
- Horrocks, Dylan (2010). *Hicksville*, 2nd ed. Drawn and Quarterly Press.
- Kulvicki, John (2006). *On Images: Their Structure and Content*. Clarendon Press. <http://dx.doi.org/10.1093/019929075X.001.0001>
- Kulvicki, John (2014). *Images*. New Problems in Philosophy series. Routledge.
- Lemmon, Oliver and Ian Pratt (1999). Putting Channels on the Map: A Channel-Theoretic Semantics of Maps? In Lawrence Moss, Jonathan Ginzberg, and Maarten de Rijke (Eds.), *Logic, Language, and Computation* (Vol. 2, 143–164). CSLI.
- Lopes, Dominic (1996). *Understanding Pictures*. Clarendon Press.
- Lopes, Dominic (1997). Art Media and the Sense Modalities: Tactile Pictures. *Philosophical Quarterly* 47(189), 425–440. <http://dx.doi.org/10.1111/1467-9213.00069>
- Millikan, Ruth (2004). *Varieties of Meaning*. MIT Press.
- Pratt, Ian. 1993. Map Semantics. In Andrew Frank and Irene Campari (Eds.), *Spatial Information Theory* (77–91). Springer.
- Rescorla, Michael (2009). Predication and Cartographic Representation. *Synthese* 169(1), 175–200. <http://dx.doi.org/10.1007/s11229-008-9343-5>

- Schier, Flint (1986). *Deeper into Pictures*. Cambridge University Press. <http://dx.doi.org/10.1017/CBO9780511735585>
- Schlichtman, Hansgeorg (1979). Codes in Map Communication. *Canadian Cartographer* 16, 81–97. <http://dx.doi.org/10.3138/T166-2241-G815-70X1>
- Sober, Eliot (1976). Mental Representations. *Synthese* 33(1), 101–148. <http://dx.doi.org/10.1007/BF00484711>
- Tanaka, Kenneth, James Skinner, James Dohm, Rossman Irwin, Eric Kolb, Corey Fortezzo, Thomas Platz, Gregory Michael, and Trent Hare (2014). Geologic Map of Mars: US Geological Survey Scientific Investigations Map 3292 [Map], pamphlet 43 pp. <http://dx.doi.org/10.3133/sim3292>
- US Department of Commerce and US Department of Defense (2013). US Chart 1: Symbols, Abbreviations and Terms Used on Paper and Electronic Navigational Charts (12th edition). Retrieved from <http://www.nauticalcharts.noaa.gov/mcd/chartno1.htm>, accessed September 28, 2014.
- USGS (2012). Claremont North Quadrangle, New Hampshire, Sullivan County. US Geological Survey Map [Map].