Map Semantics and the Geography of Meaning

Gabriel Greenberg, October 14, 2021

1 Maps and meaning

Maps are rich with potential for meaning. Their geographic contents are what enable simple acts of wayfinding, say from the grocery store to the gas station. But they can also reinforce the tectonic outlines of global ideologies, like the dominance of First World nations over the Third World.[1] The many meanings of a map can be visualized as a loosely organized sequence of concentric spheres. At the center is the most strict, conventional interpretation of a map as a representation of objects and features distributed in space. The spheres of meaning radiate out from that core in degrees of social dependence and political significance.[2]

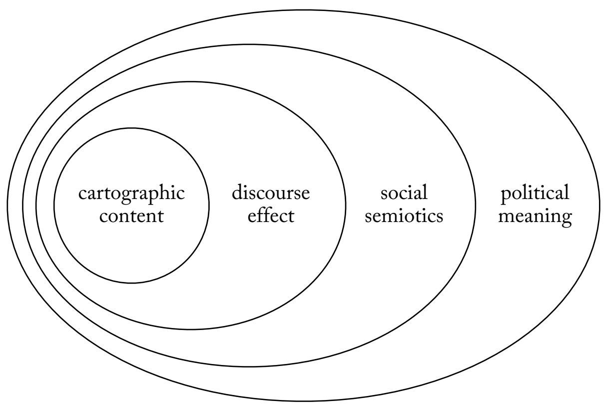

Figure 1: Spheres of meaning for maps.

The inner-most sphere of meaning is the literal meaning of the map, what I will call its **cartographic content**. It describes the way that objects and properties are distributed over some geography, and it is determined from the map itself by more and less natural semantic convention. The cartographic content of the map in Figure 2, for example, fixes the locations of concretia like roads, rivers, and population centers, but also relative abstracta like political boundaries, zip codes, and types of land usage. There are few limits on the kinds of properties that a map can encode at this level, so long as they are the sorts of thing that can be sensibly understood as having geographic

Acknowledgements. Thanks to Tendayi Achiume, Malihe Alikhani, Joshua Armstrong, Sheridan Bartlett, Elizabeth Camp, Sam Cumming, Catherine Hochman, John Kulvicki, Kevin Lande, and Matthew Stone for invaluable discussion leading to the ideas presented here.

[1] By "First World" nations I mean to refer to former colonial powers and by "Third World" nations I mean formerly colonized nations; see Achiume 2019, 1513-14.

[2] This spherical model of meaning isn't intended to be an exact theory, much less an original one. Compare Panofsky's more carefully worked out levels of iconology for pictorial art (Panofsky 1939, 5-8).

locations.

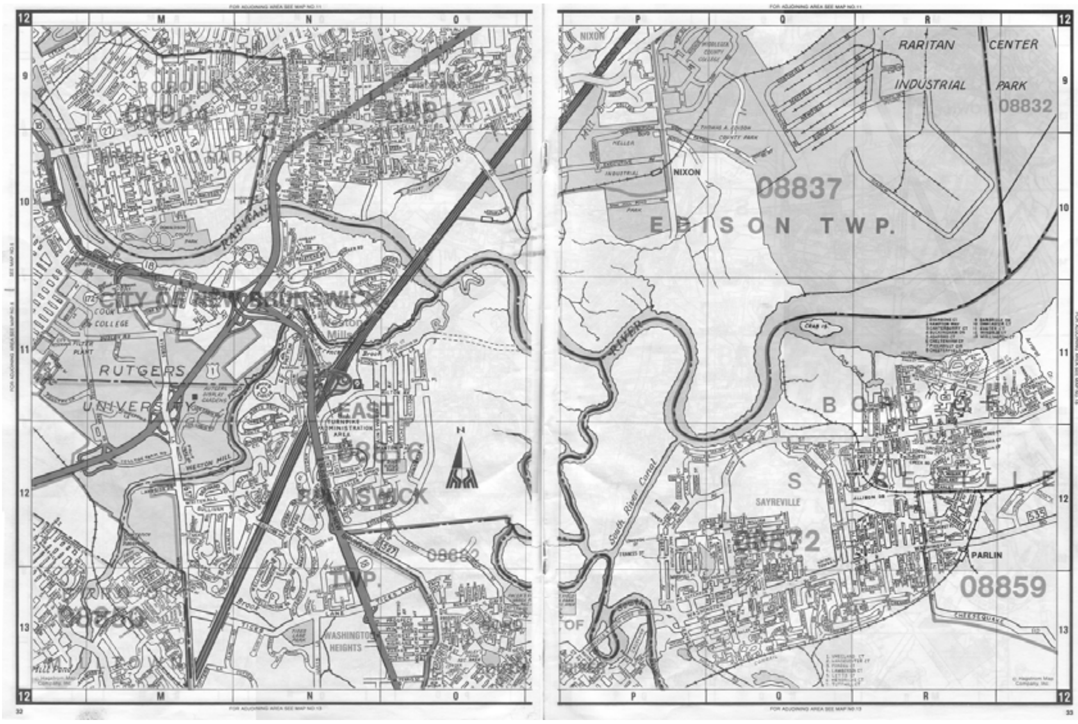

Figure 2: Excerpt from from *Middlesex County Atlas* (2002, pgs. 32-33).

Next, the **discourse effect** of a map, on a given occasion of use, corresponds to how the map's cartographic content is leveraged to contribute to the broader exchange of information in discourse. This, in effect, is the speech act level of representation for maps.[3] Here we can distinguish between *assertoric* uses of a map to describe how things are (like a road map) and *imperatival* uses to describe how things should be (like an urban planner's map for a city). Maps, of course, can also be counterfactual, cautionary, or fantastical. In most of these cases, the map maker will choose a *target* for the map: the time and place (actual or counterfactual) which the map functions to accurately describe. A map's discourse effect also includes the way a map contributes to the broader *multi-modal discourse* in which it may be embedded. Consider an actual road map: it includes not just a map narrowly defined, but also an index, legend, distance table, scale, compass rose, descriptive text, and a title; all of which combine to form a further, complex whole.

Whereas cartographic content and discourse effect have to do with the articulation and management of information, the **social semiotics** of a map describes the social forces which lead to a

[3]See the frameworks described by Murray and Starr 2018, 2021 and Lepore and Stone 2015, ch. 14 for the scope of discourse effect and the kinds of mechanisms that govern it.

map's production and display, and the characteristic effects it has on the psychologies and social-identities of its audiences. Some of the traditional interests of pragmatics may be found here.[4] We may ask, for example, what *purpose* the maps was created for, or what problem it was designed to solve: navigation, building, warfare, play, boundary making, colonization, or, more commonly, some mixture of uses.[5] Social semiotics also includes a map's *connotation*: the associations, assumptions, and expectations it reliably elicits from audiences in a given cultural setting.[6] For example, how does the design of the map signal its authorativeness (or tentativeness), its political affiliation, or its basis in scientific evidence? Here the manner of display is also relevant: is the map printed on one page of many in an atlas, or is it ostentatiously hung on the wall of a study, or is it part a company's logo? What is the displayer trying to say about themselves?[7]

At the outermost sphere is a map's **political meaning**: the way it shapes and is shaped by ideology and the societal distribution of power, often independently of the conscious awareness of those who create and use the map. Because maps are so intimately tied to land, and land to sovereignty and nationality, maps are an inevitable source of rich political meaning.[8]

Consider the way a map can make geography "available" to those with access to the map, automatically conferring a kind of epistemic power to the map's possessors. This might play out at a local level: having map of a nearby neighborhood allows you to explore, shop, or dine there, even though you may feel yourself an outsider. At a larger scale, cities that are featured in the map of a popular guidebook may make those cities welcoming to tourists. Of course, maps of the globe are the ultimate site of colonial ambition, rendering "undiscovered" lands available for conquest in the minds of the map holders.

> As much as guns and warships, maps have been the weapons of imperialism. Insofar as maps were used in colonial promotion, and lands claimed on paper before they were effectively occupied, maps anticipated empire. (Harley 1988, 132)

Likewise, the composition of a map can have irresistible effects on the ways that map consumers understand the distribution of status and power in the world. Most modern maps follow the convention of putting North on the top of the map. Applied to globe maps, this has the effect of placing most First World nations on top, and most Third World nations below. The stark visual re-centering that occurs when the map is inverted reveals the subliminal effects of the standard orientation. The inverted world map has the same cartographic content as the original, but at the

[4]The present taxonomy does not follow the classical division of communicative labor into semantics and pragmatics. Here I am guided by the arguments of Lepore and Stone (2015) to the effect that the traditional domain of pragmatics can be factored without remainder into discourse effects, on one hand, and *ad hoc* psychological and social engagement, on the other.

[5]See e.g. Wood 1992, ch. 1.

[6]See extensive analyses of map connotation in e.g. ch. 5 and MacEachren 2004, 330–51.

[7]Compare Taylor, forthcoming on representations of social-identity in the art-world.

[8]See Wood (1992, ch. 2) and Harley (1988) for vivid studies of political meaning in maps.

level of political meaning, it clearly shifts the global south into the center of attention.[9]

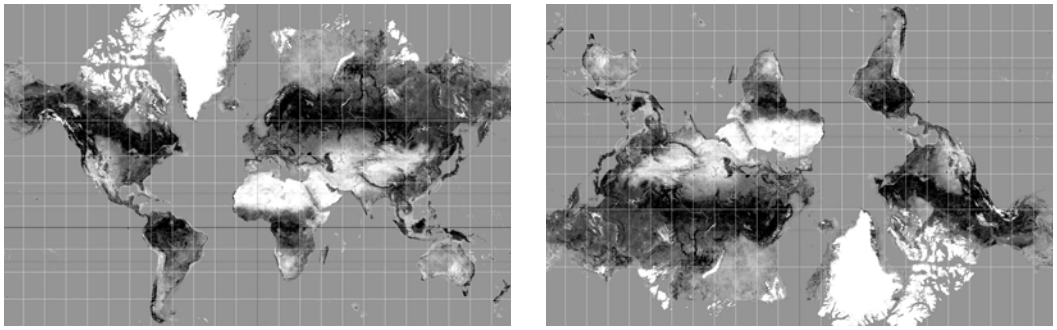

Figure 3: North-up and south-up Mercator projections. (Image from Wikipedia; the projection is cropped for ease of presentation.)

Or consider the late-20th-century debate over the rise of the Gall-Peters projection.[10] The classical Mercator projection, still widely used at the time, preserves local shape, and has functional benefits for nautical navigation. It also has the effect of inflating the size of landmasses further from the equator, so that the areas of First World nations in North America and Europe, as well as Russia and Greenland are all expanded relative to the map's total size, while Third World nations in Africa, South Asia, Central America, and South America, along with Australia, are all compressed. The Gall-Peters projection, widely criticized for its supposed aesthetic defects, represents all areas proportionally.[11] The inevitable effects of the two maps on audiences' assessment of power and importance are obvious.

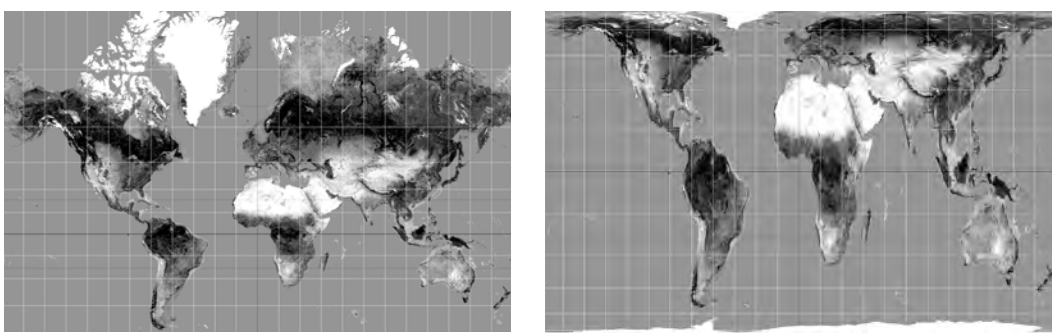

Figure 4: The Mercator projection (left) preserves local shape. The Gall-Peters projection (right) preserves area. (Images from Wikipedia.)

[9] Danforth 2014 and Williams 2016. Meier et al. 2011 provides psychological evidence for value associations with higher and lower spatial positions on a map.

[10] Crampton 1994.

[11] Hall 1992, 380 and Wood 1992, 56-61.

As map consumers, our interests range over all forms of meaning, from spatial display to political potency. Still, as the model of concentric spheres suggests, all these meanings depend at least to some extent on the inner-most sphere of cartographic content. Understanding how maps express this core is the task of this essay.

Part of what makes maps special is the diversity of expressive resources they marshal to fix their core cartographic contents. They are "picture-language hybrids" (Kulvicki 2015, 149), combining pictorial elements to construct space, and a host of symbolic and linguistic elements to enrich this space with features, landmarks, and terrain. In this respect, the study of map semantics joins the broader investigation within cognitive science into iconic and multi-modal representations, one that has already seen breakthroughs in our understanding of diagrams, gesture, sign language, and more.[12]

A central theme of this essay is that we ought to think of maps as a special kind of picture.[13] Like any other picture, maps consist of marked 2D planes that represent 3D environments. And tellingly, though this is sometimes overlooked,[14] maps always represent their subject matter from a point of view. This is usually a spread out "god's eye view," but it is a point of view nonetheless, showing a particular spatial and temporal part of the world from an overhead position. Just as pictorial perspective is the result of geometrical projection, maps too are governed by a range of projective methods. In what follows, I'll argue that projection from a viewpoint is the central principle of map semantics, just as it is for pictures.

Notwithstanding these analogies, maps differ from ordinary pictures by relying more on explicit symbolic signs, and less on implicit visual cues. Thus they standardly include a large share of symbolic elements, including a library of essentially arbitrary color and line-codes (the sort of thing specified in a legend), along with textual labels, and graphical symbols. And whereas pictures typically rely on a range of depth and shape cues to convey volumetric space, maps typically do not engage depth perception, 3D shape perception, or object recognition.[15] Wollheim put it this way:

> To [extract information from maps] we do not rely on a natural perceptual capacity, such as I hold seeing-in to be. We rely on a skill we learn. It is called, significantly, "map-reading": "map-*reading*." (Wollheim 1987, 60-61)

Crudely speaking, the interpretation of ordinary pictures relies more heavily on normal perceptual

[12]See, for example, work on diagrams (Shin 1994; Shimojima 2015), on speech with iconic gesture (Lascarides and Stone 2009a, 2009b), on images with linguistic captions (Alikhani and Stone 2019, 2018b), on iconic use of classifier constructions in sign language (Davidson 2015, 491-98), and iconic variables and predicates in sign language (Schlenker, Lamberton, and Santoro 2013, 103-20; Schlenker 2018). See Schlenker 2019 for an overview of recent work.

[13]This is a widespread, but not universally shared view. See Lopes 1996, 94 or Kulvicki 2015, 149 for sympathetic statements. See Dilworth 2002, 191-92 for discussion.

[14]See, e.g. Casati and Varzi (1999, 187): "Maps display no shortenings, no perspective effects: a map can be conceived as close as it gets to a view from nowhere about the area it depicts."

[15]187.

processes, whereas that of maps relies more heavily on symbolic conventions. As a result, an explicit semantics of maps is much more nearly within grasp, while a detailed understanding of the visual system is still a monumental work in progress.

In the remainder of this essay, I first discuss a variety of facts about the expression cartographic content that should inform any semantics for maps, in Section 2. Then, in Section 3, I outline a semantic theory which aims to explain these facts and illuminate the general mechanisms of map representation. Section 4 situates maps semantics within the broader geography of meaning.

2 Space, lines, and labels

Cartographic content is constructed by bringing together two aspects of a map's meaning: the space it expresses, on one hand, and the objects and properties, denoted by markings and labels, that it distributes throughout this space, on the other. In this section I highlight a handful of simple but fundamental facts of map representation, organized around the key representational elements of cartography: space, markers (including lines, colors, and textures), and tags (including labels and graphical symbols). Lessons are drawn for the eventual articulation of a semantics for maps.

2.1 Space

A **metric space** is one where every point in that space stands in a determinate distance to every other point. The familiar 3D physical space that surrounds us is metric in this sense. Such a space is constituted by an underlying network of metric relations which connect every point to every other point. Of course, a map itself defines a 2D metric space, while the world it represents is a 3D metric space.

How should we understand the relationship between map-space and world-space? To approach this question, I restrict attention to maps which express metric content, unlike quasi-topological subway maps, and to maps that are planar projections, as opposed to the spherical projections displayed by globe maps. Henceforth, this is what I'll mean by "map."

To make things concrete, let us introduce a set of axes in the represented world, where the XY plane runs perpendicular to the direction of gravity (or parallel to the ground, were the ground flat), and the Z dimension runs parallel to the direction of gravity (or perpendicular to the flat ground). We can now state two basic facts about the spatial content of maps.

Fact 1. Maps attribute complete metric relations in the XY-dimensions.

Fact 2. Maps (generally) do not attribute metric relations in the Z-dimension.

Fact 1 says that a map will determine the X and Y position, akin to longitude and latitude, of every location depicted, relative to every other location depicted. This is why, given a map and a landmark on the map, you can determine the direction and distance "as the crow flies" of every other landmark. The restriction to the XY-dimension means that maps attribute complete metric relations to the landmarks they depict— but only when these are projected onto the XY-plane.

Fact 2 says that maps do not attribute metric relations in the Z-dimension (unless they include special markings, such as contour lines). Maps, in other words, do not attribute relations of depth. This is not the same as attributing zero depth; maps do *not* require for their accuracy that the geographies they represent be flat planes. Rather, they are silent on matters of depth; their spatial content is indeterminate in the Z-dimension.

An important corollary of Fact 2 is that maps do *not* express fully metric relations between all the locations or landmarks they depict. Put another way, you cannot, just by reading a map, tell the distance between every pair of landmarks depicted. This fact is familiar to hikers. Suppose you are located at landmark B, and are trying to decide, from the map in Figure 5, whether it will be a faster walk to A or to C. Superficially, it looks on the map like the distance to B is shorter than the distance to A. But such looks can be deceiving. In fact, there is a sharp downward slope that separates A from B, revealed in Figure 6, while B and C lie along the same plane. As a consequence the AB distance is greater than the BC distance (not to mention a more arduous hike), even though, in Figure 5, the distance *on the page* between "A" and "B" is shorter than that between "B" and "C".

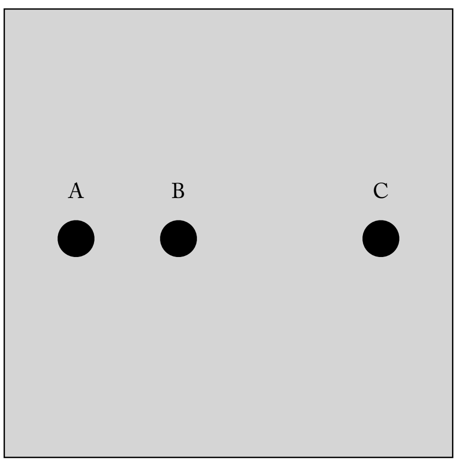

Figure 5: A map of the terrain.

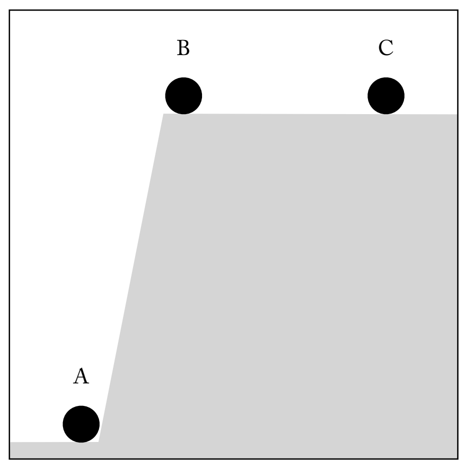

Figure 6: An elevation drawing of the terrain.

These elementary observations tend to be overlooked in contemporary semantic accounts of maps. As a consequence, we still have not got a fully satisfactory treatment of what is arguably the most basic feature of cartographic meaning: the expression of spatial relations.[16]

A number of authors adopt a quasi-linguistic view of maps, in which the locations of the map are treated as referring terms, and their denotations are locations in the world.[17] Markings on

[16]See Camp 2018, 36-41 for an extended discussion of this theme.

[17]Kulvicki (2020, 120) discusses and endorses this approach; Casati and Varzi 1999, 191 is a prominent example. See

maps denote properties, and these properties are attributed to the world-locations denoted by the map-locations occupied by the markings. At the extreme, views of this kind hold that the spatial content of a map is exhausted by attributing properties to locations in this way. The limitations of such a view have been argued forcefully by Camp (2018, 38), who concludes that: "a map makes the spatial relations among denoted object/properties directly available in a way that a list of symbol-pairs plus [an assignment function] does not." In other words, a purely referential approach to cartographic space fails to explain Fact 1.

Pratt (1993, 79-80), for example, begins his exploration of map semantics with an analogical extension of predicate calculus. Just as *names* in a sentence denote *individuals* in the world, *co-ordinates* in a map denote *locations* in the world. Locations, in turn, are taken to be geographic individuals. Since Pratt models map coordinates as pairs of numbers, he preserves a measure of metric relations *in the 2D space of the map*; unfortunately, the crucial attribution of metric relations in the 3D space of the world drops away. Pratt simply stipulates an assignment from map coordinates to locations; not only is the assignment unresponsive to the spatial properties of the map coordinates, it returns no spatial attributes to be associated with the locations that are its values. This is a consequence, Camp (2018, 27) argues, of treating maps too much like an exotic form of symbolic representation, and not taking seriously the distinctively iconic way in which maps encode space.

The lesson for a semantics of maps is that the mere denotation of locations by points on the map is not sufficient to establish the spatial aspect of cartographic content. The spatial relations of points on the map to one another must be iconically rendered into attributions of spatial relations in the geography. We therefore seeks a geometric principle that might relate the 2D space of the map to the 3D space of the represented geography. As Rescorla (2009, 197-98) concludes "A map's geometric structure is not just another element to be listed alongside its markers and coordinates . . . Rather, the markers and coordinates stand in geometric relations, and those relations bear representational import."[18]

Scholars facing this juncture have consistently proposed that the relevant relation is one of isomorphism with respect to metric relations (supplemented, perhaps, by a scalar transformation). "A map represents how various objects and properties are distributed in physical space. It does so by purporting to replicate relevant geometric features of the spatial region it represents" (Rescorla 2009, 178). In other words, many have found it natural to assume a semantic rule along the following lines.[19] Here I use the notation $[\![X]\!]$ to refer to the denotation or content of X.

Camp (2018, 28-31) for a critical discussion.

[18]Likewise, 39: "spatial relations in the map represent spatial relations in the world." See Camp 2007, 161 for more on spatial relations in maps.

[19]Rescorla (2009, 178) offers an explicit formulation much like the one in the text, and points to Russell (1923, 89) as an antecedent. Casati and Varzi (1999, 195) make topological isomorphism of the denotation function a prerequisite to semantic evaluation, rather than an explicit constraint on truth. (They acknowledge the need for metric constraints of the same kind on p. 196.) See Camp (2018, 39-40) for objections to Casati and Varzi's way of building isomorphism into the

(1) **Isomorphism Semantics**
A map M is accurate at a world w only if the distance between any two points p and p' in M is equal to the distance between $[\![p]\!]$ and $[\![p']\!]$ in w, times a scalar constant k.

But this apparently sensible proposal seems to be incompatible with Fact 2 and its corollary. For as we saw above, due to the vicissitudes of depth, distance on the 2D plane of the map is not generally proportional to distance in the 3D geography represented; isomorphism only holds when the geography itself is perfectly flat.[20] Thus the isomorphism semantics would entail that the map in Figure 5 is *in*accurate, because the distance relations between points on the map are not isomorphic to distance relations between the denoted locations in the world. And yet, Figure 5 is *perfectly* accurate (assuming the geography depicted in Figure 6): it doesn't comment about distance relations in general, only distance relations as projected to the XY plane. That it would be ineffective as a map for hiking— but effective as map of voting districts— is independent of its accuracy conditions.

Another way to put the point is that the isomorphism semantics neglects the role of viewpoint. It presupposes that the relevant spatial relation between map and world can be defined exclusively in terms of their intrinsic geometries. Yet every map displays its geography relative a viewpoint. At the very least, maps depict geography from *above*, and usually in a specific *orientation*. The isomorphism semantics has no role for viewpoint. As a result, it misses the fact that, relative to such a viewpoint, relations of distance in the Z dimension can be lost in projection, as in Figure 5, whereas they would be explicit if the viewpoint were oriented to the side, as in an elevation drawing like Figure 6.

What is correct about the principle of metric isomorphism, of course, is that distance on the 2D plane of the map is isomorphic and proportional to distances in the 3D geography *as projected onto a 2D plane*. But then the principle of isomorphism is no longer doing any real work; we might as well just say that the 2D map is a projection of the 3D geography. Thus the intuition behind the isomorphism approach quickly leads us to adopt a projection semantics for maps. Perhaps this should be no surprise, as projection is one of the central themes of modern cartography, and because projection semantics was originally developed for oridinary pictures, where issues of depth are paramount. In any case, the general moral is that it is projection, not isomorphism, which should provide the semantic backbone of cartographic representation. I pursue this idea in Section 3.

2.2 Markers

The space of a map is filled with lines, colors, and textures, collectively known as **markers**. Every marker has a **location** on the map and belongs to a **marker type**. Marker types might include,

truth-conditions. Camp herself endorses something like an isomorphism semantics on p. 41.

[20] Leong 1994, 138 resolves the problem by explicitly idealizing all geographies as flat projections.

for example, *brown color fill* or *green color fill*, or in the case of lines, *dotted lines* or *solid lines*. The meanings of the various marker types are typically supplied in a legend or map key. Though there is considerable leeway in the kinds of meanings that can be assigned, they are usually properties or categories. *Brown* might mean *land*, for example, while *green* might mean *forest*.

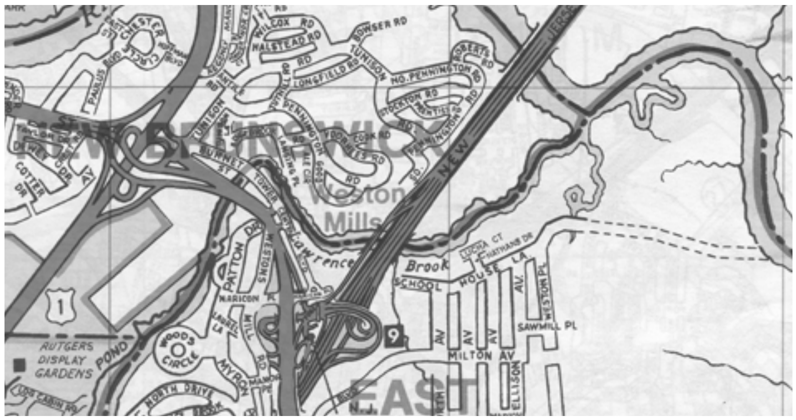

Figure 7: A variety of marker types. Excerpt from from *Middlesex County Atlas* (2002), pgs. 32-33.

As a demonstration of the diversity of possible marker types, one can find in the close-up of the Middlesex County map in Figure 7, at least eight different kinds of markers:

(2)

(i) 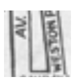 ... *small road*

(ii)  ... *unmaintained small road (?)*

(iii)  ... *highway*

(iv) ... *stream*

(v)  ... *zip code border (?)*

(vi) 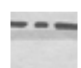 ... *town border*

(vii)  ... *body of water*

(viii) ... *land*

We can think of the syntax of a map marker as factored into the two components of (a) marker type and (b) relative spatial location on the map. Nearly everyone writing on the semantics of maps has recognized that these two aspects of marker syntax express two kinds of content: (a) the marker type expresses a property; and (b) the location of a marker on the map helps pick out a spatially situated object in the world to which that property is attributed. (Here the notion of "picking out" is intentionally vague; this could be reference or it could be mere quantificational description.) The kinds of object that a marker might pick out varies widely, potentially including landmarks, buildings, cities, rivers, patches of land, or swaths of geography.

No only do the two aspects of marker syntax express different kinds of content, they do so in different ways. Marker types *symbolically* express their associated properties. That is to say, the semantic connection between marker types and properties ultimately depends on arbitrary juxtapositions of sign types with meanings. Yes, these associations are often motivated— *blue* is chosen for bodies of water, *brown* for land— but the true semantic value is always fixed by the list-like conjunction of marker types with properties, rather than any quasi-pictorial interpretation.

By contrast, the relative locations of markers on the map *iconically* express relative locations in the world, and these in turn are used to pick out the designated object. In other words, marker locations bear an intrinsic and natural relation to locations in the represented world. The connection between marker locations and worldly location is *not* one of arbitrary juxtaposition; it is not symbolic.[21]

We can put these observations together in the form of two further facts about map representation that a semantic theory must account for.

Fact 3. Markers express properties by symbolic association of marker types with properties.

Fact 4. Markers pick out objects for attribution by iconic association of the relative location of the marker on the map with relative locations of objects in the world.

Facts 3 and 4 work together: the properties of Fact 3 are attributed the located objects of Fact 4. Much has been made of the contrasts between maps and predicative language.[22] In these two facts, we see the depth of both the similarities and the differences. On one hand, unlike language, the very same syntactic element— the marker— both expresses a property and picks out an object, in part by exploiting iconic properties of the map. On the other hand, two different semantic mechanisms are responsible for contributing two bits of content— the property and the object. Thus the deeper structure of subject-predicate attribution is still preserved. In this way maps contrast with holistic theories of picture meaning, like that of Hopkins (1998, ch. 4) for example, where property and object are represented without differentiation as part of a single situation, by a single semantic mechanism.[23]

2.3 Tags

Tags are symbols or symbolic expressions, including words and linguistic phrases, which are associated with specific subregions of a map.[24] Paradigm examples are names indicating the location and identity of landmarks, but a variety of other sub-clausal phrases can be linguistic tags, including numerals, nouns, adjectives, as well as definite and indefinite descriptions. (For now I'll

[21] The concepts of iconic and symbolic representation invoked here are developed in more detail in Greenberg 2021b.

[22] See e.g. Camp 2007, Rescorla 2009, and Camp 2018.

[23] See 29 for remarks sympathetic to such a holistic view of map semantics.

[24] Compare Leong 1994, 64-71, and Kulvicki 2015, 169-71 on labels and annotations in maps. The discussion here follows Greenberg 2019, which provides a general analysis of tags in maps and other visual media.

set aside indexical sentences, like "you are here," which can also play a tagging-like role.) Non-linguistic tags include a variety of specialized graphical symbols, like ⌂, ✝, or ⛺, which may be listed in a legend, or conventional for a type of cartographic discourse.

In the excerpt from the Middlesex County map below, for example, one finds names for towns and roads, a name for an entire railroad line, as well as numerals for local zip codes, and a numeral for a local route:

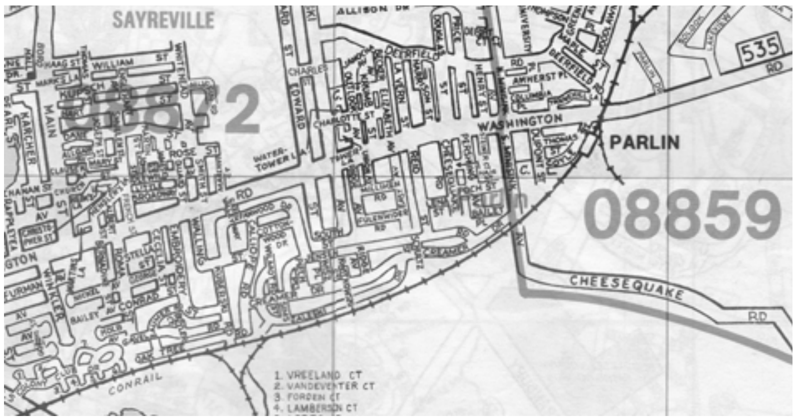

Figure 8: Tagging as the annotation of a map with linguistic symbols. Excerpt from from *Middlesex County Atlas* (2002), pgs. 32-33.

Much like markers, symbolic tags play the role of contributing identifying and descriptive information about particular objects for which the map provides spatial information. An adequate semantics of tagging must address three central facts about the use of tags in maps.

Fact 5. The relative location of a tag on the map (partly) determines the relative location of the object in the world that it describes.

Fact 6. There are a variety of ways, including placement at a variety of map locations, that a tag can determine the relative location of the object in the world that it describes.

Fact 7. Tags vary widely in their semantic function, including predication, naming, and more.

To illustrate these points, I'll refer to maps of the following landscape:

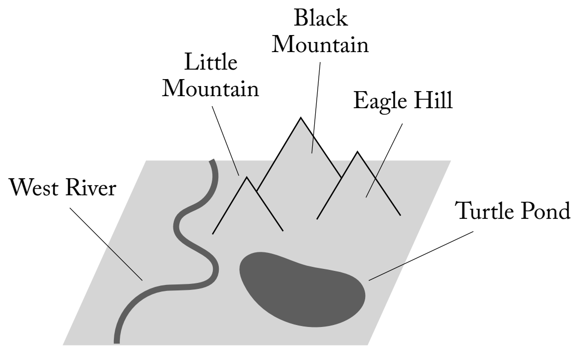

Figure 9: The landscape.

Fact 5 has to do with the semantic significance of tag placement. The placement of tags within a map contributes to the content of the map by indicating which depicted objects are associated with the contents of the tags, with consequences for accuracy conditions. For example, relative to the landscape above, (3) is accurate. But swapping the position of “Eagle Hill” and “Turtle Pond” results in (4) which is not accurate. In effect, tags locate linguistically expressed properties within map space, just as markers locate the properties they denote within map space.

(3)

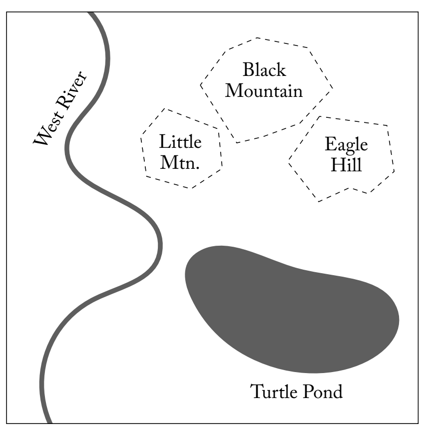

(4)

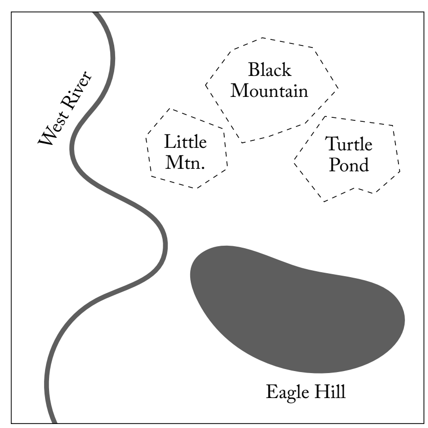

Figure 10: Relative to the landscape, map (3) is accurate, but (4) is not.

Yet the ultimate semantic contribution of tags can be achieved through a variety of expressive means; this is the point of Fact 6. In Figure 11, for example, the tag “Turtle Pond” is associated with the image by placement *proximal* to the part of the map it tags. “West River” is both proximal and *aligned* with the depiction of the river. “Black Mountain” is associated with a part of the image by spatial *inclusion* within it. And “Little Mountain” and “Eagle Hill" are associated with the map through *line linking*.[25] (In Figure 8 one can also see *numeral linking* at the bottom of image.)

Not only that, but the exact location of each tag is irrelevant to the content expressed. The content of Figure 12 is the same as that of Figure 11 even though most of the tags have been moved. Both the variety and flexibility with which tags can be associated with map regions distinguishes them from markers, which always comment on precisely the region of the map where they are located.[26]

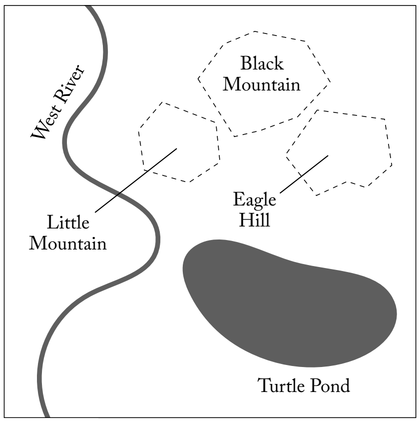

Figure 11: Tagging via line linking, inclusion, alignment, and proximity.

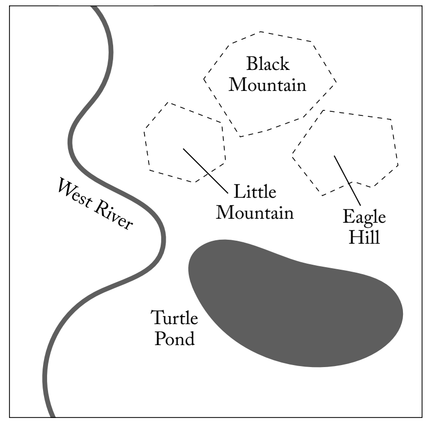

Figure 12: Local changes of tag position do not affect content.

The tension between Facts 5 and 6 can be resolved recognizing that the various of ways of tagging a region on a map all correspond to a single underlying relation of **linking** at the syntactic level. Inclusion, proximity, alignment, and line-linking are all *signals* of the same syntactic relation.[27] On this model, the location of a tag on the printed page is not itself part of syntax. Tags themselves have no location; they are associated with regions of the map plane by abstract syntactic links.

Finally, Fact 7 highlights the variety of ways tags can work semantically. As Figure 13 demonstrates, nouns and names can both be tags, but because nouns and names denote objects in different semantic categories, the semantic significance of tagging must itself be allowed to vary. A landmark tagged with the name "Black Mountain" expresses the content that the landmark *is identical with* ⟦Black Mountain⟧. An area tagged with the noun phrase "marsh area" expresses the content that the area *has the property* ⟦marsh area⟧. More extreme variations are common. The numeral "48°" in Figure 13 expresses the content that the tagged area *has atmospheric temperature*

[25] Alikhani and Stone (2018a, pg. 3555) call this *indication*.

[26] Leong (1994, 65) puts this point in terms of identity conditions of the map: local changes of tag location result in the same map.

[27] Compare Alikhani and Stone 2018b, pg. 2. See Greenberg 2019 for discussion of how the different signals of linking should be resolved.

⟦48°⟧. In Figure 8, the tag "08859" relies on the relation *has the zipcode*.

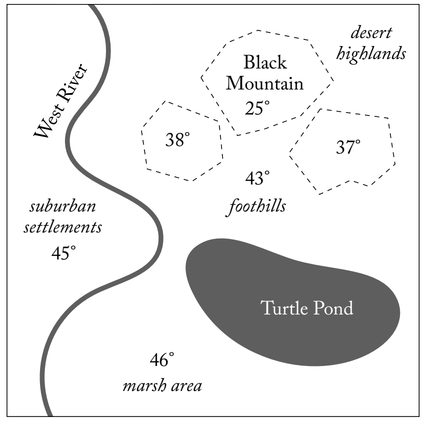

Figure 13: Tags as names, as noun phrases, and as numerals for local temperatures.

As these examples show, there is no single way to interpret a tag. Each tag is connected to the represented object by an implicit **tagging relation**. While the most common tagging relations are relations of *naming* and *predication*, a map maker can introduce any number of others, like *zip code* or *temperature* as convenience and recoverability allow. Since the operative tagging relation for a given tag is generally not explicitly signaled, they are determined instead by context and local conventions.[28] Such relations appear to be a species of the more general phenomena of *multi-modal coherence relations*. These are structural links in a discourse that function to bind together independent elements from different modalities, such as speech with iconic gestures, pictures with captions, or arrows with objects in a diagram.[29] Recognizing that maps rely on the same kind of system of flexible multi-modal relations allows for a well-motivated and empirical adequate account of the semantics of tagging.

3 A semantics for maps

A semantic theory for maps aims to explain how maps systematically express their cartographic contents. Braddon-Mitchell and Jackson (2006, 180-84), Camp (2007, 2018) and Blumson (2012), among others, have argued that map representation is in fact systematic in the way that

[28]See Armstrong 2016 for an account of local convention. Despite flexibility, there are constraints on how tagging relations may be expressed. Identity and predication appear to be defaults. Other relations are inferred when these defaults are incoherent or otherwise ruled out in context. A further constraint is typographic consistency: tags with the same typographic features are expected to encode the same tagging relation. Note the use of plain text versus italics in Figure 13.

[29]On coherence relations for speech with gesture, see Lascarides and Stone 2009a, 2009b; for pictures with captions, see Alikhani and Stone 2018b, 2019; for arrows in diagrams, see Alikhani and Stone 2018a.

invites semantic analysis. And a number of carefully worked out semantic theories for maps have already been developed, for example by Pratt (1993), Leong (1994), Casati and Varzi (1999, ch. 11), and Rescorla (2009). In what follows, I hope to build upon these advances, while heeding the lessons about space, markers, and tags from the preceding discussion.

Maps themselves are complex representational objects that function as the nexus point for a variety of interlocking interpretive rules and conventions. In this section I outline a semantic theory for maps that combines the key elements of pictorial projection, marker interpretation, and tagging, and sets them within the common framework of possible world semantics. I again restrict my attention to the simplest of metric maps in planar projection, setting aside quasi-topological (subway) maps, globe projections, and the many, many sophisticated forms of marking and tagging that populate real-world cartography.

Still, there are innumerably many planar, metrical map systems which vary in their specific treatment of space, markers, tags, and more. There is no sense in attempting a semantics for *all* such map systems, or trying to formalize the key features that might unite them. Instead, I will develop a semantics for a particular, artificially simplified or "formal" map system, what I will call **map system** S. My hope is that system S has enough in common with the natural map systems that a semantics for S may provide a template for other map semantics going forward.

3.1 Accuracy, content, and target

Maps are **accurate** or **inaccurate** relative to a map system, just as sentences are true or false in a language.[30] Whether a map is accurate relative to a system depends on the content the map expresses in that system and on how the world is. It also depends on what time and place within the world is supposed to be depicted by the map, and what position and orientation it is supposed to be depicted from. Collectively, these latter elements make up what we may call a **viewpoint**. A map is not simply accurate or inaccurate at a (possible) world, but at a world and a viewpoint.

We can think of the core, **cartographic content** of a map, relative to a system, as the set of conditions under which it is accurate. In traditional linguistic theory, the content of a sentence is understood as the set of possible worlds relative to which that sentence true. Since maps are accurate relative to a world and a viewpoint, we'll treat the cartographic content of a map as the set of world-viewpoint pairs at which the map is accurate. This is the inner-most sphere of meaning for maps.[31]

The content of a map as a whole can be factored into two aspects: **singular content**, which involves reference to particular objects; and **attributive content**, which consists of properties and relations that are attributed to objects, which are themselves either quantified or included in the singular content. Most of the cartographic content we care about is attributive: spatial relation-

[30] See, e.g. Leong 1994, 153-54 and Rescorla 2009, 178.
[31] See Blumson 2009, and Greenberg 2021a, §1 for parallel accounts of pictorial content.

ships, types of geography, types of road, location of borders. Explicitly represented singular content comes primarily from linguistic tags: names of cities, roads, and landmarks.

There is the additional question, discussed in Section 2.1, of whether maps also express indexical singular content— *that* piece of land, *that* road— even when they are not named or otherwise explicitly singled out. I am sympathetic with the idea that maps, and pictures generally, are vehicles of indexical reference in this way.[32] Yet as I argued above, following Camp (2018), even if there is such singular content, this fact has little bearing on the central question of how maps convey their spatial content. For the purposes of this analysis, I set aside the question of indexical singular content.

Still, there is a clear sense in which a given map is, say, a map *of Vermont*, or a map *of Vermont in 1800*, or a map *of your neighborhood following the proposed construction of the highway*. These are descriptions of what we may call the **target** of a map. That target of a map is a particular world-viewpoint pair which the map is supposed to accurately represent. Often the map's target is partially specified in a title or caption. The target of a map is not strictly part of its cartographic content, but belongs to the sphere of meaning having to do with speech acts and discourse effects. The content of a map is what the map *says*; its target is the situation in space and time that a given use of the map is supposed say something *about*.

When a map is intended to be accurate in the actual world, the target of the map consists of the actual world and a viewpoint located within the space of the actual world. Maps of past, present, and future geographies correspondingly have targets whose viewpoints correspond to different times in the actual world; counterfactual and generic maps may have targets sampled from other possible worlds.[33] When maps do target the actual world, they acquire a kind of particularity. They are directed at a specific time and place. This goes a long way towards addressing the intuition that maps, by nature, are representations of particular locations.[34]

3.2 Viewpoints and viewpoint constraints

A map is accurate at a given world only relative to a **viewpoint**: an oriented position in space and time. A viewpoint, in this sense, is not literally a point, but a complex structure which gives maps spatial and temporal perspective, and can be defined, relative to a world, by a series of interlocking elements, illustrated in Figure 14.[35]

(5) **Elements of viewpoint**

(i) the position and shape of a projection source, relative to a world w;

(ii) the position and shape of a cartographic plane, relative to the projection source;

[32] See Lopes 1996, ch. 5 and Greenberg 2018 for defenses of this kind of view for ordinary pictures, and Kulvicki 2020, 120-23 for maps.

[33] See Cummins 1996, Greenberg 2018, and Pavese 2020 for development the theory of target.

[34] Kulvicki 2020, 119-22 defends this intuition.

[35] See Ross 1997, ch. 5, Blumson 2009, and Greenberg 2021a, §1 for parallel discussion of viewpoint in pictures.

(iii) a bottom-to-top vector on the cartographic plane;
(iv) a back-to-front vector on the cartographic plane:
(v) a single temporal location for all these elements.

Item (i), the projection source, establishes the position and orientation of the viewpoint. For the planar projection maps considered here, the projection source itself is a plane, while for linear perspective projections, it would be a point.[36] Item (ii) gives the relative position of the cartographic plane, the surface onto which features in the world are ultimately projected to constitute a map. When considering a particular map, the cartographic plane will always be proportional to the map itself. Elements (iii) and (iv) establish the orientation of the cartographic plane, and (v) identifies a temporal index for all.

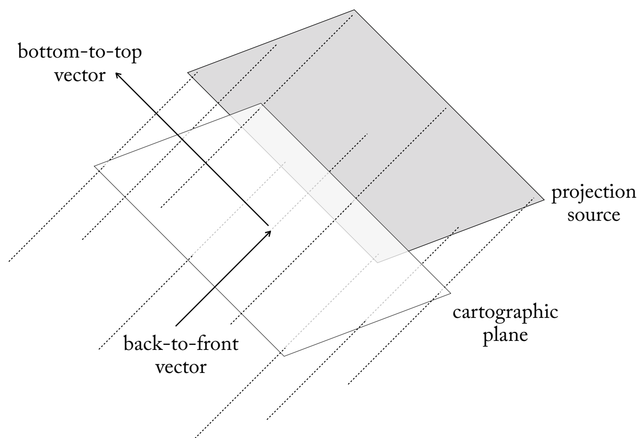

Figure 14: Elements of viewpoint (shown from below).

Generally, maps are to be understood as depicting the world from above, as defined by the direction of gravity. There are also norms of orientation: customarily the top of the page is aligned with North, but other orientations can be achieved with a compass rose. We can understand these norms as providing conventionalized restrictions on the relationship between the world depicted and the position and orientation of the viewpoint from which it is depicted. Such **viewpoint constraints** amount to interpretive assumptions that help determine the content of a map, but do not depend on the map's specific composition.[37]

[36]Globe projections require more elaborate notions of viewpoint. For example, the projection source might be located at the center of the globe, in front of, rather than behind the cartographic plane; and projection itself might take place in several discrete steps.

[37]Cumming, Greenberg, and Kelly (2017) and Cumming et al. (2021) identify a range of viewpoint constraints at work in film; they appear to be endemic to visual media. Dilworth (2002) discusses similar constraints for pictures (pgs. 185-86) and maps (pgs. 198-99), though Dilworth believes that, in general, map content depends on the physical orientation of the map, not its intrinsic orientation. See Leong (1994, 154) for a version of the North Constraint.

(6) **Cartographic Gravity Constraint**
$G\text{-}Con(w, v)$ iff the back-to-front vector of v is anti-directional with the direction of Earth's gravity at v in w.[38]

(7) **North Constraint**
$N\text{-}Con(w, v)$ iff the bottom-to-top vector of v is co-directional with the south-to-north direction of Earth in w, when projected to v.

(8) **Compass Rose Constraint**
$C\text{-}Con(w, v)$ iff the S-to-N vector of the compass rose in M, embedded at v, is co-directional with the south-to-north direction of Earth in w, when projected to v.[39]

I call (6) the *Cartographic* Gravity Constraint to distinguish it from the Gravity Constraint of ordinary pictures, where the *top-to-bottom* vector of the picture is normally aligned with the direction of gravity. The cartographic constraint holds for nearly all terrestrial maps, but alternative conventions that prioritize other vectors will hold, for example, for maps of space stations. The North Constraint and the Compass Rose Constraint can come into conflict. In modern map-making, as I discussed in Section 1, the North Constraint tends to hold by default, but it is over-ridden by the presence of a compass rose.[40] Camp (2018, 28-29) introduces the idea of an "unanchored" map, such as a plan for a possible estate with no settled building site, or a template for the layout of a chain of drug stores; in such cases, it seems that no orientational constraints apply.

What viewpoint constraints apply to a given map can be signaled by an explicit sign (like a compass rose), by general or local convention, or by ad hoc indicators of the map maker's intention. For convenience, I'll treat the operative viewpoint conventions as part of the current map system, allowing that they will vary from system to system. The whole set of available viewpoint constraints may be defined by the broader network of interpretive conventions from which particular map systems emerge.

Although viewpoint constraints apply independently of the composition of the map itself, they are an important part of the apparatus of interpretation. They play an essential role in narrowing and fixing cartographic content.

3.3 Projection semantics

In this section I outline an approach to cartographic semantics that enlists geometrical projection, rather than metric isomorphism, as its organizing principle.

[38] When taking a projection of the ground, the map as read is facing *up*, away from the ground. Thus the back-to-front vector of the cartographic plane faces *away* from (is anti-directional with) the direction gravity.

[39] The Compass Rose constraint requires a notional embedding of the map, including the compass rose itself, at the viewpoint; more on this idea in the next section.

[40] There are local exceptions to this, like New York City street and subway maps, in which top-to-bottom is conventionally aligned with the lower-Manhattan-to-upper-Manhattan vector. Other default orientations were common in the pre-colonial era (Danforth 2014).

Geometrical projection describes a family of methods for mapping a 3D to scene onto a 2D surface, via a system of straight lines that are structured by a viewpoint. Familiar forms of projection from the study of depiction include perspective projection— where the lines of projection converge on a point— and parallel projection— where the lines of projection run orthogonal to a plane. More exotic projections are involved in large-scale cartography, due to the challenges of mapping the sphere of the globe onto the 2D plane. In any case, we may define a projection function *proj* which takes possible worlds and viewpoints within those worlds as arguments and returns 2D projections as values.

The idea behind a **projection semantics** is that, for a picture to be accurate at a world and a viewpoint, it must be a projection of that world, from that viewpoint.[41] To work out the content of a picture is determine the kinds of worlds and viewpoints from which it could have been projected. Thus the content of a picture can be thought of as all of the world-viewpoint pairs from which the picture can be projected. Schematically, where P is a picture and $[\![P]\!]$ is its content, such a semantics takes the following form:

(9) **Projection Semantics**

a. P is accurate at w, v iff $proj(w, v) = P$

b. $[\![P]\!] = \{\langle w, v\rangle \mid proj(w, v) = P\}$

To adapt a projection semantics for maps, we will have to get into the definition of projection a bit. There are two basic steps. First, a viewpoint is positioned in space relative to the world. This induces a set of projection lines, which extend from the viewpoint, through the cartographic plane, into the world. This effectively sets up a correspondence between points in the world to points on the cartographic plane. In the method of projection illustrated in Figure 15, the projection lines are all orthogonal to a planar projection source and orthogonal to the cartographic plane. I assume a parallel projection like this for maps throughout this essay. In perspective projection, by contrast, the lines of projection converge on a point.

Second, features in the world are mapped to feature on the cartographic plane, following the lines of a projective correspondence. If, for example, a point in the world belongs to a body of water, then the corresponding point on the cartographic plane might be assigned the feature *blue*, and forests in the world might be paired with *green*. This projection of world-features to map-features is carried out for every line of projection. Conceived of this way, projection is a method for taking a region of the world, ranging over all of the basic features there, and transmitting them to the map. Of course, any number of different schemes for associating world-features with map-features are possible.

[41] For developments of projection semantics for pictures see Greenberg 2011, 2013; Abusch 2015; Greenberg 2021a. Earlier analyses of depiction in terms of projection include Hopkins 1998, ch. 3, Kulvicki 2006, ch. 3, and Hyman 2006, 5.

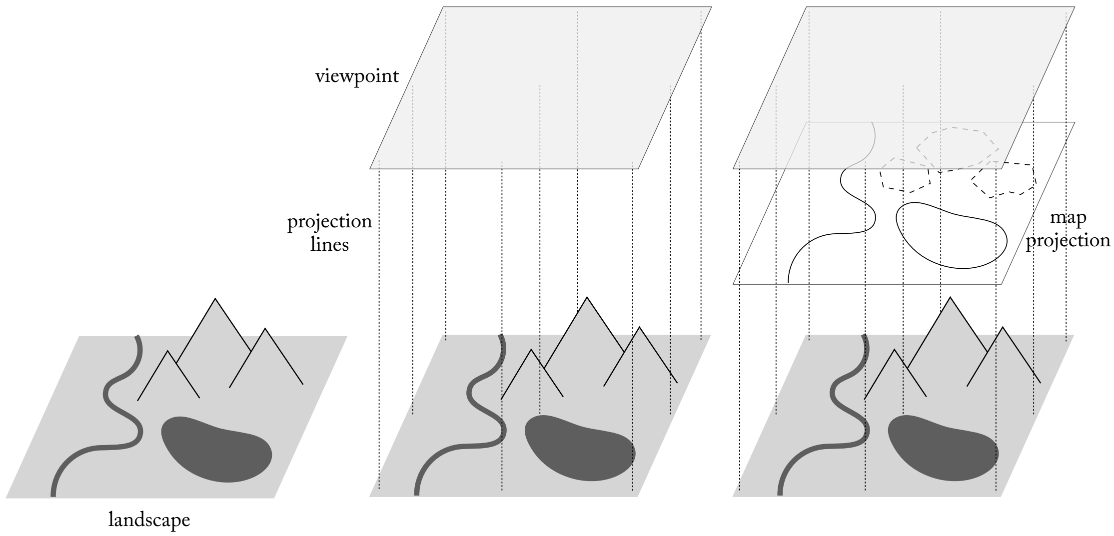

Figure 15: Construction of a map projection.

To use projection as an interpretive principle, we reverse engineer this process. Following the two basic steps involved in projection, a projection semantics can be divided into two components. First the **projection semantics** proper tells us how we should interpret the spatial distribution of markers within the space of the map. Second, a **marker semantics** assigns meanings to each of the markers. Dividing the semantics into these two elements provides for a more compositional semantics, and more modular analysis. This, in effect, is the plan for what follows.

Before moving to state the projection semantics, let us first define a minimal syntax for maps without tags, what I'll call a **pure map**. For present purposes, a pure map is a 2D plane segment where distinguished regions are associated with marking types.[42]

(10) **Syntax of pure maps in S**

A structure $M = \langle P, d, R, mark, V \rangle$ is a **pure map** in S iff:

(i) P is a set of points;

(ii) d is a Euclidean metric over P which defines a 2D rectilinear space;

(iii) R is a set of significant regions of P;

(iv) *mark* is a function from R to sets of marking types;

(v) V is a pair of orientation vectors on P: bottom-to-top and left-to-right.

Here P is the set of points from which the map is composed, and d supplies its basic spatial structure. R isolates those regions of the map which are semantically significant; I'll assume

[42]See Leong 1994, 41 and Kulvicki 2015, 151 for the distinction between a map's syntactic structure and its physical form.

that the regions in question are normally contiguous, but possibly overlapping, and correspond to psychologically natural segmentations of the graphical space.[43] Next, *mark* associates each semantically significant region with a set of marker types.[44] Finally, V gives the intrinsic orientation of the map (regardless of how the physical vehicle is pinned to the wall).[45]

To think of the pure map as a projection of a worldly scene, one has to put objects in the world into geometrical relations with the map itself. This requires notionally embedding the map in the space of the world being projected, where the embedded map is scaled up to match the size of the geography being projected.[46] This is sensible if one keeps in mind that a "map" here is an abstract object constructed around a 2-dimensional space. I'll say that the *embedding of a map M at world w and viewpoint v*, written $M_{w,v}$, is identical to M except that the points that make up map (P) belong to the space of the world w. The internal structure of the viewpoint v determines the scale, position, and orientation of the embedded map.[47]

We can now state the projection semantics, where $markers_S$ is defined as the domain of $[\![\cdot]\!]_M$.

(11) **Projection semantics for pure maps in S**
M is accurate at w, v in S iff:
for every region r in $M_{w,v}$: there is an object o in w such that:
(i) for all projection lines ℓ from v: ℓ intersects $r \leftrightarrow \ell$ intersects o;
(ii) for all $\tau \in markers_S$: $\tau \in mark(r) \rightarrow [\![\tau]\!]_M(o)$ at w, v.

The semantics says that a map M is accurate at world w and viewpoint v if and only if, when M is embedded at w, v, every region in M corresponds to an object o in the world in the following way. (i) First, for any ray from v, it intersects with r when and only when it intersects with o. This sets up a spatial correspondence between regions in M and objects in w. (ii) Second, for any marker type available in S, if r has that marker type, then o has the property which is the denotation of that marker type.

The projection semantics has the same basic structure as the map semantics described by Leong (1994), Casati and Varzi (1999), or Rescorla (2009), since it divides labor between a geometrical constraint between map space and world space, on one hand, and interpretation of marker

[43] The semantic function of image sub-regions is anticipated in Abusch's analysis of visual co-reference (Abusch 2012, §3-5; Abusch 2015, §4). I'll assume here that the shapes of the regions in question are perfectly definite, though this is certainly an idealization in many cases.

[44] This definition of map syntax allows for many parameters of variation, but imposes few limits. A more complete and explanatory account would put substantive constraints on this structure. For example, Kulvicki (2015, 153-59) has proposed that marking types are divided into incompatibility classes, such that only one type per class can be associated with any given region. Leong (1994, 62) distinguishes basic categories of marker types— line, area, spot, and boundary— and requires (p. 138) that certain marker types can only be paired with certain categories (like dotted lines with lines, and colors with areas).

[45] Dilworth (2002, 3; 2003, 41-42) defends the idea that intrinsic orientation is a constitutive element of a picture.

[46] The embedded map, then, is much like the 1:1 scale map of Borges' (1998) fable, except the fact that it is abstract would protect it from the inclemencies of weather.

[47] The orientation vectors of M match the orientation vectors of v following the right-hand rule.

types, on the other. The virtue of the projection semantics is that it preserves correspondence of metrical relations in the XY dimension, without committing the content of map with respect to relations of depth in the Z dimension.

Ultimately, projection semantics is a natural fit for maps, not only because maps seem to extend the pictorial way of representing the landscape, but also because it makes sense of the facts about cartographic space reviewed in the last section. Of course, what I have described here is a very simple projection semantics for a very basic map system. There are infinitely many other potential methods of projection, and much of modern cartography has been devoted to the formal study of existing methods and the invention of new ones.

3.4 Marker Semantics

The projections semantics has to be supplemented with a semantics for marker types. For the sake of illustration let us suppose that system S has five marker types:

(12) **Marker types in *S***

(i)  ... white fill

(ii)  ... light grey fill

(iii)  ... dark grey fill

(iv) 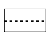 ... dotted line

(v) 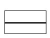 ... solid line

The marker semantics itself has the same structure as a linguistic lexicon. It take the form of a list-like pairing of marker types with properties. We can define the interpretation of marker types in S as follows, where $[\![\cdot]\!]_M$ is the specialized interpretation function, in S, for marker types:

(13) **Marker semantics for *S***

(i) $[\![$ 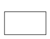 $]\!]_M = \lambda x.\lambda w.$ x is part of unforested land in w;

(ii) $[\![$ 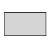 $]\!]_M = \lambda x.\lambda w.$ x is part of a forest in w;

(iii) $[\![$  $]\!]_M = \lambda x.\lambda w.$ x is part of a body of water in w;

(iv) $[\![$ 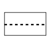 $]\!]_M = \lambda x.\lambda w.$ x is part of the border of a territory in w;

(v) $[\![$ 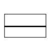 $]\!]_M = \lambda x.\lambda w.$ x is part of a road in w.

In fact, this statement of the marker semantics is only approximate, because it sidesteps the crucial role of viewpoint. For example,  doesn't just mean that there is unforested land in a given direction; after all, wherever there is unforested land, there may be forested land, in the same direction, on the opposite side of the earth. Instead, 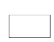 means, more nearly, x is part of *unoccluded* unforested land in w relative to v. And occlusion, in turn, isn't defined optically, but in terms of projection lines uninterrupted by land masses (since atmospheric phenomena might be optical occluders, but are not relevant). So a more complete version of (i) above might be:

(14) $\llbracket \square \rrbracket_M = \lambda x.\lambda w.\lambda v.$ x is part of a forest in w and there is a projection line from v that intersects x without intersecting any other land masses.

In what follows I gloss over the important complications introduced by the viewpoint relativity of marker denotations.

3.5 Tagging semantics

The next issue we must address is the interpretation of the symbolic tags which annotate a map.[48] As I discussed in Section 2.3, part of the challenge here is finding the right division of labor between syntax and semantics. I concluded there that tags are associated with regions of the cartographic plane by abstract structural **links**.[49] The links which connect tags to regions need not be explicitly marked on the printed page, but they are nevertheless signaled through a variety of means, as we saw in the initial discussion.

The final ingredient in the syntax of tagged images addresses the variety of tagging relations. I posit a set of **relation-symbols** which are again explicit in the syntax, but implicit on the printed page. Each link between a symbol and a region is associated with one such symbol. Formally, I treat the *tag* function as a mapping from picture regions to sets of pairs of tag symbols and relation-symbols.

We can capture the core syntactic idea formally by adding the following structure to our definition of a pure map, where R is the set of significant regions:

(15) *tag* is a (partial) function from R to pairs $\langle \rho, \phi \rangle$, where ρ is a relation-symbol, and ϕ is a tag.

To incorporate tagging into our semantics for pure maps, we'll add the following clause, where $\llbracket \cdot \rrbracket_R$ is the interpretation function for relation-symbols; $\llbracket \cdot \rrbracket_L$ is the interpretation function for linguistic tags; and $tags_S$ is the domain of $\llbracket \cdot \rrbracket_R \times$ the domain of $\llbracket \cdot \rrbracket_L$. Here r is a variable ranging over picture regions, o is a variable ranging over worldly objects.

(16) For all $\langle \rho, \phi \rangle \in tags_S : \langle \rho, \phi \rangle \in tag(r) \rightarrow \llbracket \rho \rrbracket_R(o, \llbracket \phi \rrbracket_L)$ at w, v.

Clause (16) states that, for every possible relation-tag pair, if a region r is linked with that pair, then the corresponding property is attributed to object o. The property in question is not simply $\llbracket \phi \rrbracket_L$, the content of the tag, but rather $\llbracket \rho \rrbracket_R(x, \llbracket \phi \rrbracket_L)$, the content of the tag as it is modulated by the content of the relevant tagging relation. I'll represent *identity* and *predication*, the two most common tagging relations, by the relation-symbols as "`id`" and "`pred`" respectively, and the more specialized relation *zipcode* as "`zip`":

[48] The syntax and semantics described in this section are essentially like those introduced by Leong (1994, 64-71, 154), except that Leong does not make provision for tagging relations. The key elements of the present analysis were originally presented in Greenberg 2019.

[49] Such links are a sub-segmental variant of the text-to-image links posited by Alikhani and Stone 2018b.

(17) **Relation-symbol semantics in S**

a. $[\![\texttt{id}]\!]_R = \lambda x.\lambda y.x = y$

b. $[\![\texttt{pred}]\!]_R = \lambda F.F$

c. $[\![\texttt{zip}]\!]_R = \lambda x.\lambda z.$ the zipcode for region x is z[50]

For example, if the tag is "Wantastiquet Mountain" and the relation is "`id`", then, by (16) the semantic contribution of the tag can be derived as follows:

(18) a. $[\![\texttt{id}]\!]_R(o, [\![\text{Wantastiquet Mountain}]\!]_L)$ at w, v

b. $(\lambda x.\lambda y.x = y)(o, m)$ at w, v

c. $o = m$ at w, v

If the tag is "08816" and the relation is "`zip`," the result is:

(19) a. $[\![\texttt{zip}]\!]_R(o, [\![08816]\!]_L)$ at w, v

b. ($\lambda x.\lambda z.$ the zipcode for region x is z) $(o, 08816)$ at w, v

c. the zipcode for o is 08816 at w, v

The semantics for tags can now be integrated into the projection semantics for pure maps, to yield a more nearly adequate semantics for a complete map system. This will have the desired effect of allowing the content of symbolic tags to enter into the content of the tagged image at specific, object-dependent locations in pictorial space.

3.6 Map semantics

We are finally in a position to compile the forgoing elements into a unified semantics for maps in the simplified map system S. Let us begin with syntax:

(20) **Syntax for map system S**

A structure $M = \langle P, d, R, mark, tag, V \rangle$ is a map in S iff:

(i) P is a set of points;
(ii) d is a Euclidean metric over P which defines a 2D rectilinear space;
(iii) R is a set of significant regions of P;
(iv) $mark$ is a function from R to sets of marker types;
(v) tag is a (partial) function from R to pairs $\langle \rho, \phi \rangle$ of relation-symbols and tags;
(vi) V is a pair of orientation vectors on P: bottom-to-top and left-to-right.

[50] I assume that zipcodes, as written, are interpreted, but they are not treated as numerals, but rather as numeral sequences; their contents are number sequences.

We are after a semantics for maps that is unified, but not uniform: as we've seen, maps bring together a diverse array of representational resources. The semantics for a system like S are defined in terms of a series of interpretive sub-systems, including a method of projection for the pure map, a natural language semantics for linguistic tags, and denotational assignments for marking types, for non-linguistic tags, and for tagging relations. I'll notate the various interpretation functions used in the final semantics as follows.

(21) **Interpretation functions**

- $[\![\cdot]\!]_S$ ... maps
- $[\![\cdot]\!]_M$ ... marker types
- $[\![\cdot]\!]_L$ ... linguistic tags
- $[\![\cdot]\!]_R$ ... tagging relations

Then we may think of the map system, S, as a structure containing all of these interpretation functions as components, in addition to a set of viewpoint constraints. (More elements would likely have to be added in a more fully formal treatment.)

(22) **Map system *S***
$S = \langle \mathcal{M}, [\![\cdot]\!]_S, [\![\cdot]\!]_M, [\![\cdot]\!]_L, [\![\cdot]\!]_R, \textit{vp-constraint}_S, ... \rangle$

We can now state the semantics of S as follows, where $\textit{markers}_S$ is defined as the domain of $[\![\cdot]\!]_M$, and $\textit{tags}_S$ is the domain of $[\![\cdot]\!]_R$ × the domain of $[\![\cdot]\!]_L$.[51]

(23) **Semantics for map system *S***
For any map M in S: $[\![M]\!]_S$ = the set of all $\langle w, v \rangle$ such that:
for every map region r in $M_{w,v}$: there is an object o in w:
(i) for all projection lines ℓ from v: ℓ intersects $r \leftrightarrow \ell$ intersects o;
(ii) for all $\tau \in \textit{markers}_S : \tau \in \textit{mark}(r) \rightarrow [\![\tau]\!]_M(o)$ at w, v;
(iii) for all $\langle \rho, \phi \rangle \in \textit{tags}_S : \langle \rho, \phi \rangle \in \textit{tag}(r) \rightarrow [\![\rho]\!]_R(o, [\![\phi]\!]_L)$ at w, v;
(iv) for all $\delta \in \textit{vp-constraints}_S$: $\delta(w, v)$.

(24) **Accuracy for map system S**
M is accurate at w, v in S iff $\langle w, v \rangle \in [\![M]\!]_S$.

Clause (i) of (23) is the key projection clause, establishing a spatial correspondence between regions of the map's cartographic plane and spatially located objects in the world. Clause (ii) requires that, for every region decorated with a given marker, its projective counterpart in the world instantiates the property which is the denotation of that marker type. Clause (iii) likewise requires

[51] The present analysis follows the precedent of Leong (1994, 154), whose semantics brings together a similar set of interpretive clauses.

that a region's projective counterpart instantiates the property which results from interpreting the associated tag and tagging relation. Finally, clause (iv) ensures that the map's content conforms with the system's operative viewpoint constraints.

The resulting statement of cartographic semantics is compact, but it has none of the minimalist elegance of a Tarskian theory of truth. So be it: interpretive elegance went out the window with the invention of maps. Maps are not austere registers of homogenous information. They are maximalist, doing as much with as they can within the confines of the basic spatial substrate. Future semantic theories for naturalistic map systems, I am sure, will require the addition of further clauses and constraints, and of even greater complexity and variety than those envisioned here.

3.7 Exhaustivity

There remains one issue in the semantics of maps which demands comment in part because it is the subject of such extensive discussion in the literature.[52] This the question of **exhaustivity**: whether maps should be interpreted as providing complete and exhaustive representations of their subject matters, or, whether they may be read as supplying merely partial or incomplete representations. Practically speaking, the question is: if a given marker type τ is absent at a given region on a map, does this imply that there is no object which is $[\![\tau]\!]$ in the corresponding region of the world? According to the so-called "absence intuition," which endorses exhaustivity, the answer is yes.[53] The alternative is that the absence of τ merely indicates the map's quiesence on the question of the whether $[\![\tau]\!]$ is realized in the corresponding region.

While some have thought that all maps, by definition, are exhaustive, the truth is almost certainly more complex than that.[54] Expectations of exhaustivity seem to vary by type of map, type of marker, subject matter, and even among regions on a single map. At one extreme, in a mass-produced topographical map intended for hikers or surveyors, absence of the depiction of a body of water almost certainly entails the absence of a corresponding body of water. At the other extreme, an informal map drawn on a napkin to given directions to the nearest gas station is free to include only landmarks the mapmaker feels are relevant, without risking the representation of the absence of others.

To a first approximation, the assumption of exhaustivity seems to vary with the register— the degree of formality and implied rigor— of the map itself. Exhaustive maps appear to be the special provenance of mass-production in the age of information, where it is commonly assumed that any feature which *can* be surveyed, will be. Modern map consumers expect print and digital maps to

[52] See e.g. Pratt 1993, 81-83; Leong 1994, 145-53; Casati and Varzi 1999, 192-94; Rescorla 2009, 181-96; Blumson 2012, 428-30; Kulvicki 2015; Bronner 2015; Camp 2018, 31-35; Kulvicki 2020, 123-26.

[53] Rescorla 2009, 181-83.

[54] Rescorla (2009, 181-96), following Casati and Varzi (1999, 192-94), defends the claim that exhaustivity is a definitional feature of maps. But other commentators have consistently drawn attention to counterexamples, e.g. Leong 1994, 149; Bronner 2015, 45; Kulvicki 2015, 152; Camp 2018, 33. Work by actual cartographers does not seem to support universal exhaustivity either; see.e.g. Robinson 2019; Wood 1992, 79-94.

possess a kind of encyclopedic authority.

There remains the further the question of whether exhaustivity should be viewed a semantic assumption, built into cartographic content, or whether it belongs to the realm of discourse effect, or to classical pragmatic implicature.[55] I won't take a stand on this issue here, or try to improve on the empirical generalization about when and where exhaustivity applies. Instead, I want to show how the availability of exhaustive and non-exhaustive interpretations of markers can be explained within the framework of projection semantics.

First, it should be noted that the projection semantics articulated here is non-exhaustive. The key formulation comes in the clause for the interpretation of markers:

(25) For all $\tau \in \mathit{markers}_S : \tau \in \mathit{mark}(r) \rightarrow [\![\tau]\!]_M(o)$ at w, v.

This clause requires that, for any marker type τ on any region r of the map, the object in the world which is the projective counterpart of r must have the property $[\![\tau]\!]$. This is non-exhaustive, because it allows that there may be objects which have the property $[\![\tau]\!]$ for which there is no corresponding region with τ.

Rescorla (2009, 181-96), along with Casati and Varzi (1999, 192-94), achieve exhaustivity by adding an additional clause to their semantics. The primary clause says that regions with a given marker τ must denote locations that have the property $[\![\tau]\!]$; the supplementary clause says that all regions without τ must denote locations that do *not* have the property $[\![\tau]\!]$. In the semantics presented here, the same effect is achieved by strengthening the marker clause above to a biconditional.[56]

(26) For all $\tau \in \mathit{markers}_S : \tau \in \mathit{mark}(r) \leftrightarrow [\![\tau]\!]_M(o)$ at w, v.

What is interesting, in the context of a projection semantics, is that exhaustivity, as realized by the strengthened marker clause, really signals a return to the original understanding of projection. According to this conception, a method of projection impartially surveys features of the world and transforms them into image features. If we assume that a map is accurate only if it is a projection of the world, then an accurate map is an exhaustive record of the projectable features of the world.

[55] Casati and Varzi 1999, 192-94, Rescorla 2009, 181-96, and [123-26]kulvicki2020modeling develop semantic accounts; [81-83]pratt1993map and [153-56]leong1994towards offer broadly discourse-level proposals; Blumson 2012, 428-30, Bronner 2015, Camp 2018, 31-35 advocate for pragmatic explanations.

[56] Kulvicki (2015, 153-59) offers an alternative analysis, also compatible with the semantics presented here, where exhaustivity is achieved through a refined theory of markers. Kulvicki divides marker types into "incompatibility classes:" one and only marker from each class can be associated with a given region. Thus a given region might be represented both as forest (green) and high altitude (textured lines), but not as both forest (green) and land (white). Where there is no visible marker for a given class, Kulvicki posits implicit "zero values." Zero-values are markers that are interpreted as denoting the property of instantiating none of the other properties denoted by the same incompatibility class.

Where 0 is a zero-value from incompatibility class Π:
$[\![0]\!]_M = \lambda x.\lambda w.\forall \chi \in \Pi : \chi \neq 0 \rightarrow \neg([\![\chi]\!]_M(x)$ at $w)$.

This is an unusual denotation, to be sure, but an interesting analysis of the problem.

Exhaustivity just amounts to the assumption that the map really is the result of applying projection to the world.

Here we might distinguish between **forward-projection** and **reverse-projection**. *Forward-projection* is an algorithm that impartially surveys features of a world, relative to a viewpoint, and returns a constraint on images— what an image must include to be a projection of that world. *Reverse-projection* is algorithm that impartially surveys features of an image, and returns a constraint on a world and viewpoint— what the world must include to have projected that image. Both rely on the same projective geometry, but differ with respect to how they treat features of the world and the map. Forward-projection implies an exhaustive coding of world features; reverse-projection implies an exhaustive de-coding of map features.

In its loosest formulation, projections semantics says that an image is accurate at a world and viewpoint if and only if it is a projection of that world and viewpoint. We now see that the projection semantics I've developed here relies on reverse-projection, not forward-projection. It uses space projectively to establish the alignment of map regions and objects in the world. But markers on the map are surveyed impartially, and imposed back on the content, not visa versa. This was precisely to avoid the implication of exhaustivity.

Exhaustivity arises when it is assumed that maps are also forward projections of their subject matters: they are the result of impartially recording all relevant world features onto the cartographic plane. Such an assumption might be modulated for certain kinds of features, or for certain regions of a map; and a variety of semantic or discourse-level mechanisms might effect such a modulation. What this suggestion illuminates is the nature of exhaustivity itself. It is not the result of an arbitrary supplemental constraint on map interpretation, nor a generalized pragmatic implicature, nor a principle of informational conservativity. Instead, it follows from the nature of map interpretation; it is a consequence of a strengthened understanding the projective correspondence of map and world.

4 The geography of meaning

I began this essay by isolating the narrowest sphere of literal, cartographic meaning as my subject matter. It is here, arguably, that maps are most clearly distinguished from other forms of representation. I went on to examine the roles that space, viewpoint, markers, and tags all play in the determination of a map's total cartographic content. Finally, I attempted to provide, at least in outline, a projection-based semantics of maps that integrates and explains these factors.

The value of such an investigation is not just to understand maps in themselves, but to understand how they relate to the broader geography of meaning. Typologies of representation which focus exclusively on the format of representational vehicles, or on the kinds of contents they express, miss out on what is arguably the central feature of any system of representation: the semantics by which it associates signs with contents. The deeper contours of the semiotic terrain are not constituted by *what* is represented, but by *how* it is represented.

In comparing maps to language, scholars have noted that both are systematic and productive, but there has been debate about the extent to which maps employ predication, and further debate about the extent to which they are propositional.[57] These are all substantive philosophical questions. But I think we can cut to the chase, to some extent, by directly comparing semantic theories for each type of representation.

With respect to their most basic constituents, semantics for map systems and languages have clear affinities. The semantics of map markers have essentially the same structure as linguistic lexicons. And tags, often enough, just *are* bits of language. Where the two kinds of representational truly pull apart is in the way that complex structures composed from basic elements are interpreted.

Map semantics are based on natural spatial transformations of the cartographic plane. Projection from a viewpoint, applied uniformly to every point on the map, gives structure to the attribution of markers and tags. There are no such natural or uniform transformations in the semantics of (non-iconic) language. The meanings of complex linguistic expressions are derived from the recursive application of composition rules to tree-like structures with lexical items at the nodes. The composition rules themselves assign structure types to logical operations, like function application, without regard to those structures' spatial or dimensional properties. This, I believe, is the kind of contrast Camp has in mind when describing her requirements for an adequate semantics of maps: "rather than employing a combinatorial operation that is digital, universal, and asymmetrical, we need one that is holistic, spatial, and symmetrical" (Camp 2018, 41).

With a semantics of maps in hand we can also begin to get a sense of where maps stand in relation to the surrounding terrain: not just of the distance between maps and language, but also of the more proximal representational forms. For example, previous isomorphism analyses suggested that maps belong to the same family as other isomorphism-based diagrams, like Venn and Euler diagrams. The present account, by contrast, suggests that maps are much more nearly like pictures in the way they structure space. What they have in common with diagrams is the more general feature of basing their semantics on a natural relation or transformation.

The other thing that maps have in common with diagrams is the way they effortlessly blend symbolic and iconic interpretation. In this respect, they join the wider world of multi-modal discourse, where gestures, facial expressions, pictures, and diagrams are brought together with speech, annotations, labels, titles, and symbols. This, of course, is the natural setting of human communication. Maps exemplify this combinatorial instinct by bringing spatial transformations, symbolic markers, and multi-modal discourse relations together into a unified cartographic arena. In the geography of meaning, maps are a crossroads.

[57] See e.g. Camp 2007; Rescorla 2009; Blumson 2012; Grzankowski 2015; Camp 2018.

References

Abusch, Dorit. 2012. "Applying discourse semantics and pragmatics to co-reference in picture sequences." *Proceedings of Sinn und Bedeutung 17.*

———. 2015. "Possible worlds semantics for pictures." In *Blackwell Companion to Semantics,* edited by Lisa Mathewson, Cécile Meier, Hotze Pullman, and Thomas Ede Zimmermann. Forthcoming. Wiley, New York.

Achiume, E Tendayi. 2019. "Migration as decolonization." *Stan. L. Rev.* 71:1509.

Alikhani, Malihe, and Matthew Stone. 2018a. "Arrows are the Verbs of Diagrams." In *Proceedings of the 27th International Conference on Computational Linguistics,* 3552–3563.

———. 2018b. "Exploring coherence in visual explanations." In *2018 IEEE Conference on Multimedia Information Processing and Retrieval (MIPR),* 272–277. IEEE.

———. 2019. ""Caption" as a Coherence Relation: Evidence and Implications." In *Proceedings of the 2nd Workshop on Shortcomings in Vision and Language,* 58–67.

Armstrong, Josh. 2016. "The problem of lexical innovation." *Linguistics and Philosophy* 39 (2): 87–118.

Blumson, Ben. 2009. "Pictures, Perspective and Possibility." *Philosophical Studies* 149 (2): 135–151.

———. 2012. "Mental Maps." *Philosophy and Phenomenological Research* 85 (2): 413–434.

Borges, Jorge Luis. 1998. "On the exactitude of science." In *Collected Fictions,* edited by Andrew Hurley, vol. 325. New York: Penguin.

Braddon-Mitchell, David, and Frank Jackson. 2006. "Philosophy of mind and cognition: An introduction."

Bronner, Ben. 2015. "Maps and absent symbols." *Australasian Journal of Philosophy* 93 (1): 43–59.

Camp, Elisabeth. 2007. "Thinking with Maps." *Philosophical Perspectives* 21 (1): 145–182.

———. 2018. "Why Maps Are Not Propositional." In *Non-Propositional Intentionality,* edited by A. Grzankowski and M. Montague. Oxford.

Casati, Roberto, and Achille C. Varzi. 1999. *Parts and Places: The Structures of Spatial Representation.* MIT Press.

Crampton, Jeremy. 1994. "Cartography's defining moment: The Peters projection controversy, 1974–1990." *Cartographica: The International Journal for Geographic Information and Geovisualization* 31 (4): 16–32.

Cumming, Sam, Samuel, Gabriel Greenberg, Elsi Kaiser, and Rory Kelly. 2021. "Showing Seeing in Film." *Ergo* forthcoming.

Cumming, Samuel, Gabriel Greenberg, and Rory Kelly. 2017. "Conventions of viewpoint coherence in film." *Philosopher's Imprint* 17 (1).

Cummins, R. 1996. *Representations, Targets, and Attitudes.* Cambridge, MA: MIT Press.

Danforth, Nick. 2014. "How the north ended up on top of the map." *Al Jazeera America.* `http://america.aljazeera.com/opinions/2014/2/maps-cartographycolonialismnortheurocentricglobe.html`.

Davidson, Kathryn. 2015. "Quotation, demonstration, and iconicity." *Linguistics and Philosophy* 38 (6): 477–520.

Dilworth, John. 2002. "Varieties of visual representation." *Canadian journal of philosophy* 32 (2): 183–205.

———. 2003. "Pictorial orientation matters." *The British Journal of Aesthetics* 43 (1): 39–56.

Greenberg, Gabriel. 2011. "The Semiotic Spectrum." PhD diss., Rutgers University.

———. 2013. "Beyond Resemblance." *Philosophical Review* 122 (2): 215–287.

Greenberg, Gabriel. 2018. "Content and Target in Pictorial Representation." *Ergo* 5 (23).
———. 2019. "Tagging: Semantics at the Iconic/Symbolic Interface." In *Proceedings of the 22nd Amsterdam Colloquium*, 11–20.
———. 2021a. "Semantics of Pictorial Space." *Review of Philosophy and Psychology.*
———. 2021b. "The Iconic-Symbolic Spectrum." Manuscript.
Grzankowski, Alex. 2015. "Pictures have propositional content." *Review of Philosophy and Psychology* 6 (1): 151–163.
Hall, Stephen S. 1992. *Mapping the next millennium: the discovery of new geographies.* Random House Incorporated.
Harley, John Brian. 1988. "Maps, Knowledge, and Power." In *The Iconography of Landscape: Essays on the Symbolic Representation, Design and Use of Past Environments,* edited by Denis Cosgrove and Stephen Daniels. Cambridge: Cambridge University Press.
Hopkins, Robert. 1998. *Picture, Image and Experience: A Philosophical Inquiry.* Cambridge University Press.
Hyman, John. 2006. *The Objective Eye: Color, Form, and Reality in the Theory of Art.* University of Chicago Press.
Kulvicki, John. 2006. *On Images: Their Structure and Content.* Oxford: Clarendon.
———. 2015. "Maps, pictures, and predication." *Ergo, an Open Access Journal of Philosophy* 2.
———. 2020. *Modeling the Meanings of Pictures.* Oxford University Press, USA.
Lascarides, Alex, and Matthew Stone. 2009a. "A Formal Semantic Analysis of Gesture." *Journal of Semantics* 26 (4): 393–449.
———. 2009b. "Discourse coherence and gesture interpretation." *Gesture* 9 (2): 147–180.
Leong, Mun-Kew. 1994. "Towards a semantics for a visual information system." PhD diss., Stanford University.
Lepore, Ernie, and Matthew Stone. 2015. *Imagination and Convention: Distinguishing Grammar & Inference in Language.* Oxford University Press.
Lopes, Dominic McIver. 1996. *Understanding Pictures.* Oxford: Oxford University Press.
MacEachren, Alan M. 2004. *How maps work: representation, visualization, and design.* Guilford Press.
Meier, Brian P, Arlen C Moller, Julie J Chen, and Miles Riemer-Peltz. 2011. "Spatial metaphor and real estate: North–south location biases housing preference." *Social Psychological and Personality Science* 2 (5): 547–553.
Murray, Sarah E, and William B Starr. 2021. "The structure of communicative acts." *Linguistics and Philosophy* 44 (2): 425–474.
———. 2018. "Force and Conversational States." In *New Work on Speech Acts,* edited by Daniel Fogal, Daniel W. Harris, and Matt Moss. Oxford University Press.
Panofsky, Erwin. 1939. *Studies in iconology: humanistic themes in the art of the Renaissance.* Oxford University Press.
Pavese, Carlotta. 2020. "Practical representation." In *The Routledge Handbook of Philosophy of Skill And Expertise,* 226–244. Routledge.
Pratt, Ian. 1993. "Map Semantics." *Lecture Notes in Computer Science:* 77–77.
Rescorla, Michael. 2009. "Predication and cartographic representation." *Synthese* 169 (1): 175–200.
Robinson, Anthony C. 2019. "Representing the presence of absence in cartography." *Annals of the American Association of Geographers* 109 (1): 286–300.
Ross, Jeff. 1997. *The Semantics of Media.* Kluwer Academic Publishers.
Russell, Bertrand. 1923. "Vagueness." *The Australasian Journal of Psychology and Philosophy* 1 (2): 84–92.

Schlenker, Philippe. 2018. “Visible meaning: Sign language and the foundations of semantics.” *Theoretical Linguistics* 44 (3-4): 123–208.

———. 2019. “What is Super Semantics?” *Philosophical Perspectives* 32 (1): 365–453.

Schlenker, Philippe, Jonathan Lamberton, and Mirko Santoro. 2013. “Iconic variables.” *Linguistics and Philosophy* 36 (2): 91–149.

Shimojima, Atsushi. 2015. *Semantic Properties of Diagrams and their Cognitive Potentials.* Center for the Study of Language / Information.

Shin, Sun-Joo. 1994. *The Logical Status of Diagrams.* Cambridge University Press.

Taylor, Paul. Forthcoming. “Representation, Identity, and Ethics in Art.” In *The Oxford Handbook of Ethics and Art,* edited by James Harold. Oxford University Press.

Williams, Caroline. 2016. “Maps have ‘north’ at the top, but it could’ve been different.” *BBC online.* `https://www.bbc.com/future/article/20160614-maps-have-north-at-the-top-but-it-couldve-been-different`.

Wollheim, Richard. 1987. *Painting as an Art.* Thames / Hudson.

Wood, Denis. 1992. *The Power of Maps.* Guilford Press.