Some Pictures Are Worth $2^{\aleph_0}$ Sentences

Philip Kitcher and Achille Varzi
Columbia University

(published in *Philosophy*, 75:3 (2000), 377–381)

According to the cliché, a picture is worth a thousand words. But this is a canard, for it vastly underestimates the expressive power of many pictures and diagrams. Even a simple map, such as the bare outline of Manhattan Island accompanied by a pointer marking North (Fig.1), is worth a vast infinity of sentences—including a vast infinity of useful true sentences. Here's why.

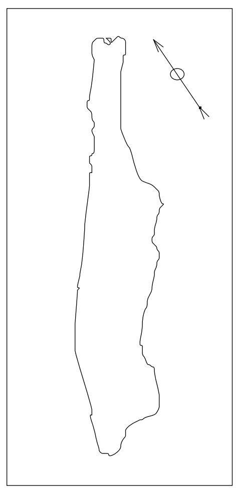

Figure 1

Let's first be clear about what the map is. The printed image of the Manhattan outline on the page before us is digital, consisting Seurat-style of a finite number of dots. The map itself, however, is a geometrical object, with the continuity properties of the real line. It contains continuum many points, that is, at least $2^{\aleph_0}$.

Second, this map (like other maps) is associated with *reading conventions* which all competent map users grasp. These conventions divide into two sorts: one sort of conventions link the syntactic elements of the map to entities in nature; the other sort tell us how to connect the properties of the visual display with properties of those entities. Thus, the map has a certain syntax and the reading conventions fix its semantics.

By way of illustration, consider a map of the world in which political territories are indicated by colored regions. Such regions are syntactic elements of the map, they stand for entities in the world (convention #1). Their colors are properties of the visual display, and the relevant reading convention is that differently colored regions correspond to territories with different political administrations (convention #2). This is, literally, a denotation assignment of the sort that is involved whenever syntactic elements are interpreted. The second convention also decrees that other properties of the visual display are projected to the corresponding territories. For example, if two regions overlap, so do the corresponding territories; if one region is next to another, then the corresponding territories are adjacent; if one region is bigger than another, so are the corresponding territories. Not all properties of regions, however, must be so projected. On old-fashioned maps, the British Commonwealth (or, in earlier days, the British Empire) was colored pink; but only an incompetent map reader would draw the conclusion that certain parts of the world, including Nigeria and the Indian sub-continent, were uniformly pink.

As a second example, consider the familiar map of the London Underground. The reading conventions for this map link various dots to stations (Charing Cross, Oxford Circus, and so forth) and various lines on the map to railway lines. These are conventions of the first sort. In addition, a competent map user knows that if two dots on the map are connected by a line, then the stations denoted are connected by a railway line, and the number of intermediate stations is exactly the number of intermediate dots. These are reading conventions of the second sort. But, as with the pink Commonwealth, there are certain things we shouldn't infer: the fact that the dot standing for Notting Hill Gate is left of one standing for Marble Arch doesn't entitle us to conclude that the former station is exactly due West of the latter. (Here, there's an important contrast with the standard map of the Paris Métro.) In short, the map preserves the topology of the London Underground but not other geometrical properties.

Back to the outline map of Manhattan. What are the reading conventions for it? First, we have to decide on the syntactic elements. To this end, let a *division* be any partition of the geometrical line into connected intervals. Then we take as syntactic elements the elements of any division. Since the set of connected intervals of the real line has cardinality $2^{\aleph_0}$, it follows that there are that many syntactic elements.

Next we must say what these syntactic elements denote (reading convention #1). The general answer is obvious: each element corresponds to a particular piece of Manhattan’s shoreline, a boundary in the physical world, if you like. Which boundary goes with which element? That’s a tricky matter. There is a vagueness problem and there is a granularity problem. Maybe there’s a consensus among cartographers about just which bit of nature counts as the shoreline, and which bit of shoreline corresponds to which element. If so, then there’s unique denotation. But maybe there are alternative ways to pick out the shoreline, or to decide just where an interval of the outline on the map starts on it. If so, then we’ll do best to think in terms of indeterminate reference: there is no unique way of interpreting the map. These problems are not peculiar to maps, however: they arise naturally also with regard to the semantics of ordinary linguistic expressions, such as ‘the shoreline of Manhattan’ or ‘the area where Martha lives’. We need not concern ourselves with such general problems here, so we’ll simplify by supposing uniqueness.[1]

Finally, we must specify how to project the geometrical properties of the display onto the world (reading convention #2). In the case we’ve chosen, this is more straightforward than in our illustrative examples. If a syntactic element has a particular curvature property, then the corresponding bit of shoreline must also have that curvature property, assuming the map is accurate. If the line joining the midpoints of two syntactic elements is at an angle θ to the North-South axis, then the relative direction of two corresponding points on the shoreline has to be θ from due North: so at least the map tells us. There is much more to it than just topology.

It should now be clear why the map is worth $2^{\aleph_0}$ sentences. Our point is not that a picture is worth so much because it is not a linguistic entity but rather because a map says a lot of things at once. For there are at least $2^{\aleph_0}$ syntactic elements, so obviously there are at least $2^{\aleph_0}$ sentences. These include $2^{\aleph_0}$ logically true sentences of the form $x = x$, which do not correspond to any specific content of the map (albeit they are, in a sense, part of what the map says). But there are also $2^{\aleph_0}$ true sentences which are peculiar to this map, and whose truth is not a matter of logic. For there are at least $2^{\aleph_0}$ geometric properties and relations that we can project from the geometrical configuration with respect to each of them (each pair of them, each triple of them, and so forth). Each such projec-

[1] On the link between linguistic vagueness and geographic ontology see e.g. P. A. Burrough and A. U. Frank (eds.), *Geographic Objects with Indeterminate Boundaries*, London: Taylor & Francis, 1996.

tion yields an interpreted sentence, and virtually all of these sentences are false. (Sad to say, even given the most careful surveying, we can't expect the cartographer to have the angles exactly right.) But some of these sentences—in fact as many as $2^{\aleph_0}$—are true; and among these are the ones that are most useful to us. Let's see why.

Take any two relatively small syntactic elements, say the closed intervals [*a*, *b*] and [*c*, *d*] (Fig. 2). Consider all the lines joining any point in one to any point in the other. These will form a band. Let the extremal angles to the North-South axis be $\theta - \delta$ and $\theta + \delta$ respectively (Fig. 3). Consider now all statements of the form "P_1 is within φ of θ East of North with respect to P_2", where the P_i are the referents of the syntactic elements chosen. These statements are a pedantic way of putting what we expect the map to tell us, to wit that one place on the shoreline is, within a certain limit of error, in a particular compass direction from another specified place. If the syntactic elements are small, then δ will be small, and, if the map meets ordinary standards of accuracy, all lines within the band will fall within an interval of width φ around the direction marked out by θ. So the map will yield lots of true *error-explicit direction statements*. How many? $2^{\aleph_0}$ of course. For there are at least that many pairs of small intervals along the coastline. So the map includes $2^{\aleph_0}$ true sentences of this sort, among which is the (large finite) number we may be interested in.

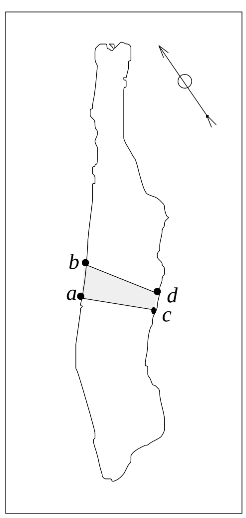

Figure 2

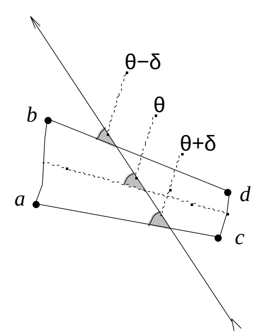

Figure 3

In conclusion we face two obvious objections. First: All this is too vague. We are treating maps as syntactic objects that receive a semantic interpretation *via* suitable reading conventions. But even a simple map requires a lot of spelling out when it comes to identifying the relevant syntactic elements and the relevant properties of the visual display. Indeed, we agree that providing a formal semantics for real life maps (not to mention other pictures) is a difficult task. But we are confronted with similar problems when it comes to the semantics of ordinary language. In that case, one standard solution proceeds from the assumption that each sentence of ordinary language corresponds to one or more expression of a suitable formal language—expressions that exhibit a *logical form* for the sentence and for which a rigorous semantics can be given. Likewise, one way to get clear about the semantics of maps is to proceed from the assumption that each ordinary map corresponds to one or more idealized maps for which the reading conventions are well defined.[2] Our reference to the map of Manhattan is to be understood in this way, which is why we can steer clear of such problems such as granularity and vagueness. It is the logical form of that map that involves $2^{\aleph_0}$ true sentences.

Second objection: All this is cheating. By taking the map as an idealized geometrical object we multiplied statements beyond necessity. Instead we should have recognized either the digital character of the display or the limits of human visual acuity. Now we don't normally think of maps or other pictures as thoroughly finite objects. If we did, we'd have to worry about whether different tokens of the same type really did contain exactly the same number of points. If the commercial printer has a higher resolution than ours are we unable to print out the same map? We think not. But we are tolerant of different conceptions. Allow then that the number of divisions isn't $2^{\aleph_0}$ but some vast finite number. For each pair of small elements we can still derive from the map true error-explicit direction statements. How many? Well, the error parameter φ can still take continuum many values. So there are at least $2^{\aleph_0}$ true contingent sentences in the map. Not all the ones that we suggested earlier, to be sure, but still just as many. Enough to show that the old cliché has sold pictures short.[3]

---

[2] This is the idea behind the semantics for formal maps outlined in R. Casati and A. C. Varzi, *Parts and Places*, Cambridge (MA): MIT Press, 1999, ch. 11.

[3] We would like to thank Laura Perini for many helpful conversations about the semantics of maps and pictures. Perini's own views are rather different from those expressed here, and her innocence of our errors will be clear once her own research on the topic is published.