From Lyons 2022 "Three grades of iconicity in perception"

Various sorts of hybrids and combinations are possible. Maps, of course, typically have symbolic elements in addition to their strong iconicity. They can also have weakly iconic and quasi-iconic elements, as when using color to represent air pressure or wind speed. Weak iconicity requires vehicle isomorphism, which requires internal vehicle structure; so when colors quasi-iconically represent wind speeds, then *patterns* of color will weakly iconically represent *patterns* of wind speeds. The relation between particular colors and particular wind speeds is scheme-level and a matter of co-similarity. Iconic representations can also combine discrete and continuous elements. A budgetary pie chart, for example, usually represents budget categories discretely and symbolically and represents percentages continuously and (weakly) iconically.

The diagram claims that all continuous representations are at least quasi-iconic.[14] I don't have a decisive argument for this, but I think it's probably true. All the instances I can think of are pretty clearly so. Additionally, I think that the noise inherent in continuous representations—the ease of mistaking one representation for a similar one—would be devastating if similar representations weren't constrained to have similar contents; thus, as a practical matter at least, continuity requires quasi-iconicity.

Again, this isn't about the *word* "iconic," and I wouldn't mind much if someone wanted to use it only for something much narrower, perhaps just those strongly iconic representations that represent in virtue of their spatial properties. What really matters are the distinctions illustrated by the Euler diagram. There are important differences between systems that represent by sharing properties and those that represent by having analogous properties. There are important differences between systems whose individual vehicles are similar to their representanda, and systems where the similarities only occur at the system level. And yet there are important similarities among all of these; co-similarity distinguishes all of these from symbolic representation and makes sense of the intuitive claim that iconic representation is somehow less "arbitrary" than symbolic representation.

4 Iconic *mental* representations

So far, I've deliberately focused on tangible, physical, non-mental representations, on paper maps and photographs and the like. My hope, however, is to connect this up with mental representation in an effort to better understand the latter. The most interesting question in this neighborhood, I think, is whether the mind employs any of the essentially spatial strongly iconic representations discussed above. I'll focus on maps, but I think the same considerations all apply to pictures as well. The idea that there are literally maps in the head now seems problematic, in a way that the idea that there are literally sentences in the head is not. Sentences are individuated only by their semantic and syntactic properties, so the thought that minds or brains could literally have sentences in them doesn't pose any fundamental trouble. It's just

[14] I.e., everything that's "analog" in the sense I would prefer (i.e., not digital) is "analog" in the sense some other authors prefer (i.e., not symbolic). Though not vice versa.

the thought that mental or brain states (or events, etc.) could have the same syntactic and semantic properties as some sentences. But strongly iconic representations represent by sharing properties with their objects, which requires them to *have* the relevant properties—in this case, spatial properties.

Surely there's no obstacle to the mind or brain exhibiting the isomorphisms required for weak and quasi-iconicity. Maybe that's iconic enough, without having to worry about strong iconicity? Can't we just say that some mental representations are at least quasi-iconic, and leave it there? I want to defend a fairly articulated view about the relation between iconicity and perception: although I think it is implausible that all perception is iconic in even the weakest sense (quasi-iconicity), I think it is nevertheless plausible that some perceptual representations may be iconic in the strongest sense. It is thus worthwhile to explore the possibility of strongly iconic mental (perceptual) representation, rather than just dropping the issue and settling for quasi-iconicity.

There's nothing problematic about the idea that mental representations might have temporal properties, which it might share with representanda and thus represent temporal properties in the environment in a strongly iconic way. But I think an even stronger, less obvious claim is defensible; I think the idea of strongly iconic *spatial* representations in the brain is not hopeless. Making sense of spatial vehicles in the brain is a difficult project that I can only touch on here, but it's a worthwhile project, as I hope to show below.

This difficult project, however, should be distinguished from a number of easy solutions that present themselves, all of which, I think, are dead ends.

4.1 Some dead ends

We want to try to make sense of the idea that mental representations could literally have spatial properties. One obvious possibility is to invoke sense-data. Sense-data are hypothesized mental items that literally have the very properties that external objects perceptually appear to have. For the table to look white, or rectangular, is for it to present a sense-datum that really is white, or rectangular. I won't say much about sense-data here, except to say that anyone who believes in them has a straightforward route to strong iconicity.

Another easy way out of the problem would be to simply appeal to phenomenal spatial properties. Suppose there's no special problem for the idea of phenomenal rectangularity and the like. Then maybe for any spatial property *F*, a mental representation "with" *F* (as opposed to merely *of F*) could just be a state that's phenomenally *F*? One problem with this move is that the question of whether (any) perceptual representations are iconic was supposed to be an empirical, scientific question, not one for introspection. The question wasn't supposed to be about how things seem to us, but about why they seem that way. Furthermore, a good many of the mental representations thought to perhaps be iconic are not (phenomenally) conscious, so phenomenal spatiality wouldn't help us with these questions anyway.

There is a lot of talk among neuroscientists about "cognitive maps" (Tolman, 1948; O'keefe & Nadel, 1978; Derdikman & Moser, 2011), especially in the hippocampus and related structures. What's meant by virtually all such talk, however, is that certain brain states encode navigation-enabling information about what is where. These internal representations have the *content* of maps and are "map-like" in the important sense that they carry information about metric relations among items in the world, and not just which direction home is. But this doesn't speak to our question here, which is about the *format* of these representations. Just as a subway rider could use a verbal list of the stations in order, rather than an actual *map* with this information, the mere fact that, say, place cells in the hippocampus carry information about conjunctions of features found in particular environmental locations, doesn't begin to argue that this information is encoded in a spatial, map-like format. But the format matters, as the subway map example illustrates: the very same information (content) is much easier to use when encoded in a map format, rather than a verbal one.

One last easy way out, which seems to have tempted nearly everyone who's written on the subject, is to note that the brain employs a large number of topographic maps. There are retinotopic maps throughout the early visual cortex, where adjacent brain regions code for adjacent retinal areas. There are tonotopic maps in the cochlea and auditory cortical regions, where high pitches are coded at one end and low pitches at the other, with in between pitches in between. And so on. These are highly significant phenomena, of course, but they don't have anything to do with the iconicity of mental representations. Of course the brain has spatial properties and some of these might be shared with distal objects; the question is whether *mental representations* might have spatial properties, and topographic maps do nothing to answer this question.

Take retinotopic maps. In a seminal study (Tootell et al., 1988), a monkey was given a radioactive dye while looking at a bullseye pattern; afterwards, the monkey's primary visual cortex was placed on a radiation-sensitive photographic film, where a (distorted) bullseye pattern was clearly visible. There were literally lines, curves, and wedges in the monkey's brain. (Or so let us concede.) Yet, there's no reason to think that at this very early stage of processing—V1—anything is being represented *as wedge-shaped*. In fact, there are reasons to think not, i.e., to think that lines, curves, wedges and the like are represented *as* lines, curves, etc. in different brain regions and later. The topographic map is often said to be a "representation" of the activity of retinal photoreceptors. Yet there's no reason to think that the brain *ever* forms a single unified mental representation of the retina. It's not needed, so long as a large number of highly coordinated representations of very local states of retinal activity exist. The contents at this level aren't things like "curve," "square," "wedge," but rather "edge at retinal location *x*," etc. Their topographic organization presumably facilitates the coordination of these local representations (as well as contrast enhancement by way of lateral inhibition, etc.), but the representations are still local.

Second, it is well known that, due to cortical magnification, the topographic maps are distorted: there are more cortical neurons with foveal receptive fields than with peripheral receptive fields. Consequently, a square presented to one side of the fixation point will produce a trapezoid in the cortex, the side nearer to the fixation point receiving a larger share of cortex. But if strongly iconic representations represent by having the very properties they attribute, this raises a puzzle: why shouldn't this represent the object as being trapezoidal, or projecting trapezoidally? There's nothing physically rectilinear in the cortex, so no strongly iconic representation as of rectilinearity.

The solution, of course, is what we knew anyway, that it's functional spatiality, rather than physical spatiality that's doing the representational work. If we could rearrange and jumble these cortical neurons while keeping their connectivity and timing the same, the topographic organization would disappear, but all the semantic and syntactic properties of the mental representations would stay the same, including the iconicity or not of these representations. Topographical realization is simply irrelevant to the question; we have to seek iconicity at the representational level, not the implementational level. All this means we can't avoid the task of coming to a better understanding of what functional spatiality—which I equate with spatiality at the level of the representation, rather than the realization—might mean.

4.2 Functional implementation

Let a *representational scheme* consist of (a) a set of representations and (b) a semantics for these representations. In compositional schemes, (a) is specified by listing primitive representational elements and providing a set of combination rules for generating well-formed representations out of these, and (b) is given by specifying contents for the primitives along with semantic composition rules to generate contents for the complex representations. Schemes will be individuated rather finely. Two schemes could differ even while having the very same "syntactic" elements (the same primitives and same possible combinations among them), if they had different semantics.

Much of cognitive psychology is concerned with discovering what kind of representational scheme a given system is using/implementing. This research, I think, is guided by the following considerations. To implement a representational scheme of a given type (e.g., Venn diagram, subway map, Roman numerals, etc.) the implementation needs to match the scheme in at least the following, functional, ways:

(i) It needs to have the same *expressive power* as the scheme. Standard Venn diagrams, for example, can express relations among three predicates, but no more; any system that can express more is not implementing a (standard) Venn diagram scheme. If pictures do have distinctive skeletal contents, then a system can't be implementing a pictorial scheme if its states don't have those contents.

(ii) It needs to exhibit the same *systematicities*. As Fodor and Pylyshyn (1988) famously noted, cognitive capacities tend to come in clusters. Among other things, this means that as certain representational capacities are gained or lost, they do so in groups. Gaining a new primitive automatically brings with it a host of new complexes; losing a combination rule means losing all the complex representations that rule made possible, etc. As Cummins (1996b) points out,