Chapter 5
Iconic Representation: Maps, Pictures, and Perception

Tyler Burge

Maps and realist pictures comprise prominent sub-classes of *iconic* representations. The most basic, most important sub-class is perception. Other types are drawings, photographs, musical notations, diagrams, bar graphs, abacuses, hieroglyphs, and color chits. I will say something about what it is to be an iconic representation and why a prominent way of thinking about iconic representation is misconceived. Although I am primarily interested in what it is to be iconic, and in the iconic nature of perception, what I have to say will, I hope, illuminate the iconic nature of maps and pictures.[1] Both rely on iconic aspects of visual perception.

A primary theme of this article is that, like all representation, iconic representation gets its representational structure from the nature of the representational functions and competencies that underlie its use. In fact, representational structure marks aspects of representational functions and competencies. Iconicity is an aspect of representational format. Although it affects how a subject matter is represented, it is not an aspect of *basic* representational structure or function. The basic representational structure and functions of iconic representation are also present within the structure and functions of non-iconic language and non-iconic thought.

The key intuitive idea underlying the notion of iconic representation is that it is marked by natural correspondences between units of representation and entities in

I am indebted to Ned Block for comments on an earlier version.

[1]For a fine discussion of differences between maps and language, see Elizabeth Camp, 'Thinking with Maps' *Philosophical Perspectives* 21 (2007), 145–182.

T. Burge (✉)
Department of Philosophy, UCLA, Los Angeles, CA, USA
e-mail: burge@ucla.edu

© Springer International Publishing AG 2018

S. Wuppuluri and F. A. Doria (eds.), *The Map and the Territory*,
The Frontiers Collection, https://doi.org/10.1007/978-3-319-72478-2_5

the represented subject matter. Here is a somewhat fuller characterization of iconic representation.[2] A representational content, or representation, R, is *iconic* by virtue of meeting the following three conditions:

(1) There is a natural, systematic 1-1-into mapping from one or more types of structural representational units, or constituents, of R, or from a repertoire that includes R, to corresponding types of entities in the subject matter that R functions to represent.[3]
(2) The mapping in (1) preserves correlations between some *relations* among structural representational units of R, or within the repertoire in which it is embedded, and natural relations among entities in the represented subject matter.
(3) R represents the relevant entities, and relations among them, in the represented subject matter, partly by way of the mapping cited in (1) and (2).

Condition (1) allows a whole representation to count as a constituent. Condition (2) allows identity as a limit case of a relation.[4] These points accommodate *very*

[2]Representations are not the only sorts of things that can be iconic. In more extensive, forthcoming work on iconicity, I explicate iconicity for *information registration*. A state X informationally registers state Y if and only if (a) instances of states X and Y statistically co-vary in a significant way, (b) instances of X tend to be caused by instances of Y, and (c) X's meeting conditions (a) and (b) is functional. Information registration is not representation. In my terminology, truth is propositional veridicality; accuracy is non-propositional veridicality. Representation requires having either accuracy conditions or truth conditions as part of the nature of the state that represents. Initial registration of the retinal image in visual systems does not have, and is not taken in science to have, accuracy or truth conditions. A bacterium informationally registers light. Although the occasional scientist attributes seeing to bacteria and even trees, no bacterium's states are explained in the statements of laws of any science as having accuracy conditions. Information registrations, however, can and commonly do meet conditions for being iconic. Registrations of the retinal image have a structure and function that map iconically to spatial aspects of the retinal image, and degrees of light intensity. In such cases, the function of the natural mapping is entirely biological, not representational. Non-representational, non-perceptual sensory states commonly bear iconic relations to sensed aspects of the environment. For more on the distinction between representation and non-representational information registration, see my *Origins of Objectivity* (Oxford: Oxford University Press, 2010), chapter 8.

[3]That the mapping is functional implies that it could fail to match structural elements in the subject matter. So there can be non-veridical and purely fictional iconic representational mappings. Fictional pictorial mappings are parasitic on real mappings. Non-veridical mappings are parasitic on veridical ones.

It is possible to allow minor divergences from strict 1-1 mappings. Perhaps for convenience two representational elements could be mapped to a single represented item, as when two circles occur on a subway map for stops at the same station on different subway lines.

I am not fully convinced by such examples. Commonly, different circles represent different positions within the same station. When they do not, it is commonly possible to regard the two different circles as the same representational element, repeated for convenience–or as occupying different maps (one for each subway line). I owe the example to Ned Block. Although I do not insist on strict 1-1 mappings, I take them to be paradigmatic.

[4]I state the first two conditions separately, although condition (2) could be taken to be implicit in what is meant by a natural, systematic mapping in condition (1).

simple iconic representation. For example, a color chit lacks a relational structure, ordinarily understood. A color chit might represent iconically through its color's being the same as the color that is represented.

Take a slightly less simple example of iconic representation. Suppose that the following map represents the light-rail line between the Western Avenue stop and the USC stop:

The dots and the positions of the names iconically represent the stations and their relative positions, and the lines iconically represent the relevant portions of the light-rail line. There are non-iconic elements in this iconic representation. The names for the stops are non-iconic.

In accord with (1), there is a natural, systematic 1-1-into mapping from dots to stations, and from lines to tracks between stations. The natural mapping is spatial. Relative positions of dots and names and the relative compactness of dots are mapped naturally to the relative positions and relative compactness of the stations. The extended nature and linearity of the lines and the relative positions of the lines are mapped to the extended and linear natures of the rails and their relative positions.

In accord with (2), spatial relations among the dots and lines preserve some spatial relations among the stations and tracks. Thus the between-ness relation among the dots preserves the between-ness relation among the stations: the middle dot is between the outer dots, and the Vermont Ave. station is between the USC and Western Ave. stations. On the other hand, distance relations are not preserved under the mapping.

In accord with (3), the map represents relevant spatial relations partly via the mappings cited in (1) and (2).

So the map represents the light-rail line and its stations iconically.

The central notion in the explication (1)–(3) is that of a *natural* correspondence. I have no definition. Paradigmatically, natural relations, including 1-1 mappings, are of the sort that natural science (physics, chemistry, biology, geology, and so on) or mathematical science represents.[5] In the light-rail map example, spatial mappings are evinced as natural by the fact that natural sciences study spatial relations.

Metrical or topological relations in spatial arrangements, relations of intensity among light reflectances or among sounds, temporal relations, relations of greater or lesser size or speed, relations of natural parts to natural wholes, relations of sound pitch or degree of pressure are examples of natural physical relations. The idea of

[5]This list of sciences is paradigmatic, not definitional. I take the notion of naturalness to be intuitive. The key point is that the mappings are not *in themselves* representational or intentional. Natural mappings are close cousins of what Grice called natural meaning. See H.P. Grice, 'Meaning', *The Philosophical Review* 66 (1957), 377–388.

natural mathematical relations is less obvious. Simple operations on the natural numbers (doubling, adding two, dividing by two, factorization) are clear examples of natural operations. Inevitably, what counts as natural in mathematics depends on degree of expertise and amount of background knowledge.

Let me say something about what a natural mapping is *not*. Natural mappings are not *in themselves* representational mappings. (See notes 2 and 5.) Being iconic is a non-representational feature of representation or representational content. When *representation* is iconic, iconicity is an aspect of how a *representatum* is represented. The natural mappings that make representation iconic are prior and independent of whether they are capitalized upon in representation.

Natural mappings are also not established by convention. They exist independently. An iconic representation's being a representation can depend on convention. The use of dots and lines to represent stations and tracks, in the map example, is conventional. The natural mapping relations have been selected conventionally for representational use. But the natural mapping relations that the convention utilizes–the mapping between spatial relations among the dots and lines, and spatial relations among the stations and tracks–are not conventional. They do not result from agreement, or unconscious but non-compulsory coordinations.[6]

A consequence of condition (3) is that an individual that uses an iconic representation (in producing it or in receiving it) must be sensitive to and competent in using the natural mappings. The individual must respond to the mappings "naturally". Thus an individual representer must be sensitive–perhaps unconsciously–to the fact that relations among structural elements of the representation are analogs of relations among some structural elements of the represented subject matter. The natural mapping must not only be, in itself, a non-representational, objective relation. It must be *natural for* users of the mapping–at least natural enough to allow relatively easy use. What is natural for a user can vary with the user's species and learning history.

I will say more about iconicity and naturalness in other work. I turn here to some ways of thinking about iconicity that are seriously mistaken. The views that I criticize are centered in Jerry Fodor's claims about iconic representation.

Fodor maintains, 'it is having a canonical decomposition that distinguishes discursive representations from iconic ones'.[7] Fodor understands compositionality in language, which he takes to be non-iconic, discursive representation, as follows:

> A…representation in L is syntactically compositional iff [if and only if] its syntactic analysis is exhaustively determined by the grammar of L together with the syntactic analyses of its lexical primitives. A…representation is semantically compositional in L iff its semantic interpretation is exhaustively determined by its syntax together with the semantic interpretations of its lexical primitives.[8]

The characterization of decompositionality in language is unobjectionable.

---

[6]David Lewis, *Convention* (Cambridge: Harvard University Press, 1969).

[7]Jerry A. Fodor, *Lot 2: The Language of Thought Revisited*, *op. cit.*, 173.

[8]*Ibid*, 172.

A *canonical* decomposition is a privileged, correct decomposition. Fodor takes representation to be iconic in that it lacks a canonical syntactic or semantic decomposition. He infers from this view that

(a) iconic representations have no constitutive structure;
(b) constituents of iconic representations are homogeneous ('each constituent contributes in the same way');
(c) iconic representations lack logical forms;
(d) iconic representations lack a distinction between semantic constituents that contribute individuals and constituents that contribute properties;
(e) iconic representations lack truth [or accuracy] conditions;
(f) iconic representations lack ontological commitments;
(g) iconic representations do not impose principles of individuation on domains in which they are interpreted;
(h) iconic representation is not representation-as.[9]

Fodor takes pictures to be paradigms of iconic representation.[10] He mentions graphs. But other than certain psychological states, pictures constitute the only case of iconic representation that he discusses in any detail. Fodor's discussion of pictures is supposed to constrain how one thinks about iconicity in perception, and presumably maps.

Fodor rests most of his reasoning on what he calls 'the picture principle':
PP(1): If P is a picture of X, then parts of P are pictures of parts of X.
Fodor notes that according to PP(1), '*all* the parts of an icon are ipso facto constituents'.

He argues for the principle as follows:

> (ARG) Take a picture of a person, cut it into parts whichever way you like; still, each picture part pictures a person part. And the whole that you have if you reassemble all the picture's parts is a picture of the whole person that the parts of the picture are pictures of.[11]

As far as I can tell, (ARG) derives from Stephen Kosslyn. Kosslyn writes: 'Depictive representations convey meaning via their resemblance to an object, with parts of the representation corresponding to parts of the object. In this case, a "part"

[9] *Ibid*, 174–177. I think that Fodor intends the claim more broadly, to mean that iconic representations lack any structure relevant to being veridical.

[10] *Ibid,* 173.

[11] *Ibid*, 173. The principle and the argument for it are also stated in Fodor's 'The Revenge of the Given', in B. McLaughlin and J. Cohen, *Contemporary Debates in Philosophy of Mind* (Blackwell, Oxford, 2007), 108.

can be defined arbitrarily, cutting up the representation in any way; no matter how you cut it, the part will still correspond to a part of the object…'.[12]

PP(1) is false. The argument for it, based on (ARG), is unsound. The conclusions about the semantics of pictures that Fodor draws from the principle are mistaken. The claims (a)–(h) and the claim that iconic representations lack canonical decomposition are all false.

I start with some small critical points. Though in themselves small, they are connected to deeper issues. In the second sentence of PP(1), Fodor can be charitably taken to mean that each picture part is a picture of a person or person part.[13] This claim is clearly mistaken. Nearly all pictures of persons have parts that picture things that are not persons or person parts. Most pictures of people do not picture them naked. Parts of a picture that picture the buttons on the person's shirt are not pictures of parts of the person. If a part of a picture represents a highlight or shadow on the person's forehead, it does not represent a part of the person.

Further, nearly all pictures of a person picture a background for the person. PP(1) holds that for every picture of a person, the parts of the picture are pictures of parts of the person. But parts of the picture of a person that picture parts of the background do not picture parts of the person.

PP(1) is false for many further reasons. For example, indiscernible micro-parts of the picture–molecules either beneath or on the surface–do not depict anything. Parts of surfaces that result from losses of paint usually do not represent anything.

I lay these problems aside. Analogs of them will return, because they connect to fundamental difficulties. One might think that, so far, I have just shown how careless Fodor has been. One might think that his position can easily be repaired.

One partial repair would be to construct a principle for the whole *scene* that the picture depicts. The point of the repair is to show that every part of the picture pictures a part of the scene. Even highlights are parts of the scene, even though they are not parts of anything else pictured in the scene. Take a realist painting of three real giraffes. One might illustrate the principle by claiming that the top half of the

[12]Kosslyn's idea is expressed in his *Image and Mind* (Cambridge, Mass.: Harvard University Press, 1980), 33; and *Image and Brain* (Cambridge, Mass.: London, 1994), 5. Fodor does not credit Kosslyn. See also Kosslyn's 'Mental Representation' in *Tutorials in Learning and Memory*, J. Anderson and S. Kosslyn eds. (New York: W. H. Freeman and Company, 1984), 105–107. Kosslyn's syntactical and semantical ideas are vulnerable to the same points I make against Fodor's. For further expressions of the Kosslyn idea, see M. Tye, *The Imagery Debate* (Cambridge, Mass.: MIT Press, 1991), 44; D. Braddon-Mitchell and F. Jackson, *Philosophy of Mind and Cognition*: *An Introduction* 2nd edition (Oxford: Blackwell, 2012), where, 179ff., they claim, 'there is no natural way of dividing a map at its truth-assessable representational joints'. Although Braddon-Mitchell and Jackson do not mention Fodor or Kosslyn, they in effect echo the view, basing it on the claim that 'there is no natural *minimum* unit of truth-assessable representation in the case of maps'. They present this view as if it were obvious. I discuss minimality of size toward the end of this article.

[13]Taken literally, the second sentence in (ARG) implies that all parts of a picture depict the person. This view is clearly mistaken. A tiny picture part that depicts the left side of a mole on the person's cheek is not a picture of the person. I assume that here Fodor is simply being careless in his formulation.

picture depicts a part of the scene. This is claimed to be so, even though the top half cuts across the middle of the giraffes' bodies, cuts across trees and their branches, and depicts an amalgam of foreground body parts and background tree parts. One might add, aping ARG, that one can "cut up" the picture any way one likes; and the cut-up parts will picture parts of the scene. Any arbitrary cutting could be reassembled to produce the original picture. One might take this argument to support 'Picture Principle (2)':

> PP(2): Every part of a picture pictures [or represents] a part of the scene that the picture pictures [or represents].[14]

One might take PP(2) to show that pictures lack canonical decompositions and to show (a)–(h).

PP(2) and the argument for it are intuitive for some. But intuitiveness does not vindicate them, or support conclusions about pictures and iconic mental representations that Fodor infers from PP(1). The basic problem for both PP(1) and PP(2) is that the arbitrary representational units (whether primitive or combinations of primitives) that they allow correspond to no units grounded in use and understanding of pictures. Any serious semantics for pictures–like any serious semantics for any representation–must be grounded in representational usage and representational competence. I will try to make this problem vivid by developing it slowly.

Let us be more specific about what a *part* of a picture is. Let us focus, as PP(1) and PP(2) should have, on parts that are on the surface and intuitively relevant to understanding the picture.

There is a notion of a Goodmanian part that includes any aggregate of scattered parts of the picture as making up a part.[15] For example, the part that depicts the upper half of the left-most giraffe's left ear and the part that depicts the right-most third of the highest leaf on the right-most background tree are not contiguous. One Goodmanian part of the picture consists of these non-contiguous picture parts.

Fodor does not rule out Goodmanian parts. His phrase 'cut it into parts whichever way you like' and his talk of 'reassembling' the parts do not stipulate that the assemblies, short of the whole reassembled picture, must be among erstwhile contiguous parts. Whether or not Fodor intended to include Goodmanian picture parts, let us pursue these matters a step further.

PP(1) and PP(2) retain *some* intuitive force on the Goodmanian understanding. The illustrated scattered "part" of the picture can be taken to depict a scattered part

---

[14]E. J. Green and J. Quilty-Dunn, 'What is an Object File?', forthcoming *British Journal of the Philosophy of Science*. Their principle PP(2) is clearly inspired by Fodor, although they do not present it as a repair of Fodor's mistakes. They argue for PP(2) in ways nearly identical to the way in which Fodor argues for PP(1)–again using Kosslyn's example of cutting up a picture in arbitrary ways.

[15]N. Goodman and H. Leonard, 'The Calculus of Individuals and Its Uses', *Journal of Symbolic Logic*, 5 (1940), 545–55; N. Goodman, *The Structure of Appearance* (Cambridge, Mass: Harvard University Press, 1951; 2nd ed. Indianapolis: Bobbs-Merrill, 1966).

of the scene. Why should one not allow Goodmanian parts to count as parts, in understanding these principles?

It is not credible that so broad a notion of depiction is relevant to a serious semantics. It is not credible that a serious semantics takes arbitrary scattered parts of the picture as representational units. The picture's semantics hinges on its use–and on the psychological competencies, processes, and types of understanding that figure in its production and appreciation. Its use and the associated psychological competencies reside in the perceptual segmentation of pictures, the intentions of the painter, and the conventions of interpretation for realist paintings. Nothing in the use, production, or appreciation of the picture corresponds to, or is explained in terms of, any such unnatural representational units. Usage, production, and appreciation treat the part of the picture that pictures a whole giraffe as a unit. They do not treat as a representational unit the scattered Goodmanian picture part that I cited.

The part-whole relation for pictures that is relevant to representational units, or representational constituents, for a semantics for pictures is constrained. It does not follow off-the-cuff intuitions about the parts of pictures.

Suppose that PP(2) is understood to *exclude* Goodmanian scattered parts. Every part of a part is to be taken to be contiguous to some other part of the part. The same problems remain.

Take a part of the picture whose left side corresponds to a small sliver of a giraffe's right flank, and whose right side corresponds to a melange of a part of a tree trunk, parts of a couple of branches, parts of leaves, and parts of patches of sky behind the foliage. Take the left and right sides of the part to be contiguous. One might find it intuitive that this picture part represents a part of the scene. It is the part consisting of that sliver of the giraffe, that part of the tree trunk, the mix of branch and leaf parts, and the visible parts of sky behind the foliage. One might grant that that is a part of a scene. One might grant that that part gets represented in a rough intuitive sense. But a serious semantics of the picture should not and does not follow such intuitions.

The semantics of the picture hinges on use of the picture and relevant psychological competencies, processes, and understanding. Nothing in the use, production, appreciation, or understanding of the picture corresponds to such "units", or suggests that such units are otherwise consequences of these factors. Perceptual segmentation, intentions of the painter, and conventions of interpretation for realist paintings simply do not cut the painting up in that way. That part of the picture is not a semantical or "syntactical" unit.

Consider the following analog of arguments for the intuitiveness of PP(2)–supporting respectively the Goodmanian and contiguity notions of part. Someone might argue as follows.

> (Goodmanian) The scattered part of the sentence 'The dog nuzzled the cat' that consists of the words 'The dog' and 'the cat' represents the dog and the cat. The dog and the cat, together, make up a part of the state of affairs that the sentence represents. So the part of the sentence consisting of 'The dog' and 'the cat' represents that part of the state of affairs consisting of the dog and the cat. Any combination of any two words or word-combinations

> in a sentence, each of which represents a part of a state of affairs that the sentence represents, is a semantical unit and itself represents a part of that state of affairs. So 'the dog the cat' is a semantical unit that represents a part of the state of affairs.
>
> (Contiguous) The part of the sentence 'The dog nuzzled the cat' that consists of 'The dog nuzzled' represents a part of the state of affairs represented by the sentence. It represents the dog and the nuzzling. These make up a part of the state of affairs that the sentence represents. So the part of the sentence consisting of 'The dog' and 'nuzzled' represents the part of the state of affairs consisting of the dog and the nuzzling. Any combination of any contiguous words in a sentence, each of which represents a part of a state of affairs that the sentence represents, is a semantical unit and itself represents a part of that state of affairs. So 'The dog nuzzled' is a semantical unit that represents a part of the state of affairs.

Naively, both arguments are intuitive. But anyone who knows anything about the semantics of language knows that these are bad arguments. Neither 'The dog the cat' nor 'The dog nuzzled' is a semantical or syntactical unit in the sentence. 'The cat' is embedded in a verb phrase that is independent of 'the dog'. 'Nuzzled' dominates that verb phrase, and is again not a part of any semantical or syntactical unit with 'The dog', except the unit of the sentence. The sentence is built from a noun phrase and a verb phrase. Decomposition of the sentence does not cut across these units. 'Nuzzled' and 'the cat' are embedded in one unit, with 'nuzzled' dominating 'the cat'. 'The dog' is another unit. One cannot mix and match. We know these things via reflection on patterns of linguistic usage, competence, production, and understanding.

Fodor recites such facts about language. But analogous points undermine his claims about pictures. The idea that there are no semantically natural joints in pictorial representation that depict natural joints in the scene has nothing to be said for it. Arbitrary combinations of picture parts can seem naively to represent parts of the scene. But most such combinations correspond to no units that figure in usage or understanding. Usage and understanding ground any serious semantics for pictures. They ground postulating picture parts or aspects that correspond to natural perceptual, intended, or conventionally demarcated units in the scene.

Parts of the picture are involved in representation of a giraffe's body and body parts. Parts of the picture represent each natural constituent of the background and that constituent's parts. Each of these scene parts is represented in the picture–just as nuzzling, the cat, and the dog are each represented in the sentence 'The dog nuzzled the cat'. One can compose arbitrary "parts" of a scene out of such materials. But there is no representational *unit*, short of the whole picture, formed by combination of arbitrary picture parts–anymore than 'The dog the cat' is a representational sub-unit of a sentence. Representational sub-units in the picture correspond to units in usage and understanding. Sub-units that represent the giraffes are thus grounded. Sub-units that represent leaves in the background are thus grounded. But an alleged butchered sub-unit that represents an amalgam of a giraffe's upper half and arbitrary slivers of sky and foliage is not thus grounded. The whole picture parses into representational sub-units grounded in patterns of usage and understanding.

As indicated, some representational sub-units induce a part-whole structure. For example, the picture will have sub-units that represent parts of a giraffe's visible

flank's surface, as well as sub-units that represent the whole visible surface. But the picture has no sub-unit that represents an amalgam of a flank part and a part of a background leaf.

Competencies associated with perceptual, intentional, and conventional patterns in making and interpreting realist paintings do not support the idea that every complex part of the picture is a representational unit. Representational units are determined by psychological use, processes, functions, and competencies. Only picture parts that correspond in some way to psychological kinds are constituents.

Similar points apply to maps. Suppose that cities larger than a certain size are represented on a road map by a standard-sized circle. Cut a piece of the map that includes part of one of the circles, and conjoin it with a non-contiguous, arbitrarily chosen piece also cut out from the map. The conjunction forms a Goodmanian part of the map. Alternatively, cut one of the circles in half; then include in the same cut an arbitrary part of the region contiguous with the circle half. In both cases, one would have map parts that intuitively map parts of the terrain. But those parts would not be representational units. Representational units are representations of cities, roads, and spaces between roads–and parts of roads, spaces, and (sometimes) cities.

Lines that represent roads represent, iconically, the roads' length and direction. Spaces between the lines map into spaces between the roads. Line parts map parts of roads. Parts of spaces other than lines and circles map parts of terrain not occupied by roads or cities. But no map usage, perceptual capacity, convention, intention, or representational understanding takes combinations of circle parts and parts of surrounding space as a representational unit, any more than 'the dog nuzzled' is a representational unit. Arbitrary map parts, whether scattered or contiguous, are not expressions of the conventions or the perceptual or conceptual competencies that ground representational content for the map.

Similar points apply to iconic perceptual representation. Perceptual structure is not determined by intentions, conventions, or understanding. It is determined by perceptual processes, functions, and competencies. The representational units in any iconic perception must correspond to natural psychological kinds. For example, the processing of the edge that forms a representation of the boundary of a giraffe's ear is a different process from the process that forms representation of the color, or the size and shape, of the branches behind the giraffe. The computation of a representation of the farther-than relation between the branch-bodies and the giraffe-bodies hinges on forming separate representations of giraffe surfaces and branch surfaces. Representational units mark representational competencies, states, and formation processes. The idea that one can cut representational states and competencies in perception "any way one likes" is out of touch with the computational and kind-explanation practices of perceptual psychology. Representational content marks representational states and competencies. Perceptual states and competencies are explained in perceptual psychology via principles for forming units of representational content. There is massive evidence that perceptual

representation is iconic–makes use of natural mappings between the format of perceptual representation and elements in the physical environment.[16]

The master mistake in Fodor's reasoning is methodological. It is the mistake of reasoning to a semantics for pictures from armchair reflection on the format of pictures, rather than from the psychological and conventional capacities that underlie their use. One cannot understand the semantics of *anything* by starting from reasoning about its format–iconic or otherwise. One must begin by reflecting on usage and underlying psychological competencies.

Before criticizing Fodor's conclusions more specifically, I set out the most generic representational kinds that ground semantics for perception, maps, and realist pictures. The representational function of all these types of representation is referential identification.[17] All function to pick out particulars (specific surfaces, bodies, places, and so on) partly through contextual, causal relations to those particulars and partly by discriminating them from other contextually relevant particulars by characterizing them in terms of properties, relations, or kinds.

In all cases, this function is grounded in representational psychological competencies.

These two functions–picking out and characterizing–are fulfilled by the three basic types of semantical primitives in iconic representations that we have been discussing. One type is comprised of *referential applications*. Referential applications are event types whose representational function is to apply the characterizers so as to pick out or refer to particular entities. Referential applications are analogous to specific, referential uses of demonstratives or indexicals. They are individuated through occurrent events. In that sense, they are not freely repeatable.[18]

A second basic type of semantical primitive marks general, freely repeatable competence in referential application. This type consists of *referential schemas* for demonstrative-like or indexical-like application. The repeatable competence to apply 'this' or 'here' on particular occasions is marked by repeatable words 'this' and 'here'. These words do not in themselves refer to anything. They refer through events of referential application. Similarly, referential applications in pictures, maps, and perceptions are exercises of schematic, repeatable referential competencies that are also marked by referential schemas.

---

[16]Citing and explaining in detail why visual psychology routinely takes visual representations to have the format of a picture-like array would take up too much space for this article. For examples of work that either illustrate or help motivate the approach, see S. Murray, H. Boyaci, and D. Kersten, 'The Representation of Perceived Angular Size in Human Primary Visual Cortex', *Nature Neuroscience* 9 (2006), 429–434; M. Silver and S. Kastner, 'Topographic Maps in Human Frontal and Parietal Cortex', *Trends in Cognitive Sciences* 13 (2009), 488–495; T. Poggio, 'The Computational Magic of the Ventral Stream: Towards a Theory', *Nature Precedings* (2011), http://dx.doi.org/10.1038/npre.2011.6117.1.

[17]Not all iconic representations represent particulars. Some graphs represent only correlations among properties. Such iconic representations lack referential applications.

[18]For more discussion see my 'Five Theses on *De Re* States and Attitudes' in J. Almog ed. *The Philosophy of David Kaplan* (Oxford: Oxford University Press, 2009); *Origins of Objectivity*, *op. cit.*, 83–84.

The third type of semantical primitive is comprised of *attributives*. Attributives are representational contents that are kinds of abilities, states, or events that function to indicate a property, relation, or kind, and to attribute it to entities referred to by the referential applications of attributives. Attributives are characterizers. When they are applied, attributives function to characterize entities picked out by the referential applications.

For example, a visual perceptual state can single out two surfaces, characterize them as surfaces (applying the attributive surface), and characterize one as farther than the other (applying the attributive farther than):

(that x_1)(that x_2)(ego-here$_3$)[Surface(x_1)Surface(x_2)Farther-than(x_2, x_1, ego-here$_3$)]

Read: that$_1$ surface farther away than that$_2$ surface from ego-here$_3$.[19] The subscripts stand in for referential applications. 'that x' and 'ego-here' stand in for referential schemas. 'Surface' and 'Farther-than' stand in for attributives. In perception, and in a picture, the attributives (unlike the linearly ordered, convention-dependent language) occur in an iconic format, mapping naturally to a subject matter. Similarly, a map can refer iconically to three stations, characterize them as stations (with iconic markers, and perhaps, but not necessarily, also with words), and characterize one iconically as between the others.

All iconic representations that we have discussed– perceptions, pictures, maps– that function to represent particular entities have a noun-phrase-like representational structure. The structure is scope- dominated by one or more applied referential schemas, including at least one applied demonstrative-like schema. The applied referential schemas apply one or more attributives.

Thus, in the blocked off illustration, the applied referential schemas are (that x_1), (that x_2), and (ego-here$_3$). There are two applied demonstrative-like referential schemas–(that x_1) and (that x_2). The applied referential schema (ego-here$_3$) is indexical-like, rather than demonstrative-like. As noted, there must always be at least one applied demonstrative-like schema in every perceptual state, in every map, and in every picture that functions to represent a real subject matter. In the illustration, the applied demonstrative-like referential schemas function to pick out a surface, each a different surface. The applied indexical-like referential schema functions to pick out a place. These three applied schemas scope-dominate the whole structure–insuring that the whole structure is noun-phrase-like, not propositional. The applied demonstrative-like schemas each apply the attributive Surface.

[19] Ego-here$_1$ is an application of an *ego-centric index*, marking the position of the perceiver. A spatial ego-centric index marks the origin of a spatial mapping from a perception to spatial structures, and does so in a way that privileges the origin as being of special psychological ("ego") significance. Nearly all perception contains such applications of spatial or temporal ego-centric indexes. See *Origins of Objectivity*, *op. cit.*, 187, 199, 287. Most commercial, paper maps lack ego-centric indexes and are *allocentric*. They map space in a way that is independent of the position of the map or the map's user. Many allocentric maps still have origins established by referential applications. See note 25.

The applied indexical-like schema, (ego-here$_3$), applies the attributives Place and ego. The three referential schemas jointly apply Farther-than.

The identificational function of iconic representations is embodied in the demonstrative-governed noun-phrase-like structure. The relevant iconic representations constitutively have an identificational function. Noun phrases in language are usually not iconic. But many share the abstract representational structure just articulated. This noun-phrase-like structure first arose in perception, which pre-dates language.

I believe that the points just made about the representational constituents and representational structure of realist pictures, maps, and perception are apriori. They derive from reflecting on the representational functions and competencies involved in these types of representation. In the case of perception, the practice of perceptual psychology accords with the account. The science of perceptual psychology takes perceptual states both to refer to particulars and to characterize them. Specific representational units with these functions are delineated empirically. The science postulates no propositional states.

Present purposes do not demand explaining this structure in detail.[20] The important point here is that any such noun-phrase-like structure has a canonical decomposition, no less in iconic representation than in non-iconic language. The structure decomposes into its constituent referential applications, referential schemas, and attributives. Complex attributives decompose into their constituents. Farther-than is almost certainly primitive. Surface (like edge) is probably also primitive. There are interesting issues here about how representations of surfaces and edges relate to representation of surface- and edge-parts, issues to which I shall return. Perceptual primitives are determined not by intuitions, but by science's discovery of basic perceptual competencies. Both primitives and complexes are non-arbitrary representational units.

Since I am less interested here in elaborating a semantical analysis than I am in discrediting the arguments that attempt to show that there is no semantical structure in iconic representations, I do not expand on these remarks. The main point is that realist paintings, maps, and perceptual states function to identify their *representata* by picking them out referentially, and to characterize them by indicating and attributing attributes. The attributives function to characterize particulars, which are picked out occurrently and contextually. Representational units are not arbitrary, but correspond to psychological competencies.

Let us return to (a)–(h), the remaining claims inferred from the false principles PP(1) and PP(2):

(a) iconic representations have no constitutive structure;
(b) constituents of iconic representations are homogeneous ('each constituent contributes in the same way');
(c) iconic representations lack logical forms;
(d) iconic representations lack a distinction between semantic constituents that contribute individuals and constituents that contribute properties;

[20] I do so in a coming book, tentatively titled *Perception: First Form of Mind*.

(e) iconic representations lack truth [or accuracy] conditions;
(f) iconic representations lack ontological commitments;
(g) iconic representations do not impose principles of individuation on domains in which they are interpreted;
(h) iconic representation is not representation-as.

Contrary to (a), (b), and (d), the constituents have constituent structure, in that they are typed as singular representations applying attributive representations. Contrary to (a), (b), and (c),[21] they have the grammar-like and semantical forms of contextual determiners dominating attributives. Contrary to Fodor's understanding of (e), they have accuracy conditions. Accuracy conditions are instances of this general scheme for pictures, maps, and perceptual states: If every singular referential application in a given representation picks out a particular and every attributive is accurate of particulars to which they are attributed, the representation is accurate. Otherwise, it is not accurate. It is obvious that realist pictures, maps, and perceptions–paradigms of iconic representation—can be accurate or inaccurate.

Where iconic representation is appropriately committal, it carries "ontological commitments"–contrary to (f). Perception is *committal*, in that it presents the world as being a certain way, and undergoes a certain sort of failure–a representational failure–if the world is not as the perception represents it.[22] Whereas perception is constitutively committal, maps and realist paintings are not. They can be presented whimsically or as fictions. But most maps, and paintings presented as depicting real entities are committal.

Fodor does not explain 'impose principles of individuation'. On any normal construal, contrary to (g), some iconic perceptual states pick out bodies as bodies, and are capable of tracking them over time. Picking out and tracking bodies as such requires operating according to principles that determine when bodies are the same and when they are different.

Contrary to (h), iconic representation is representation-as. *Every* attribution to a particular in a realist painting, map, and in a perceptual state, is a form of representation-as.[23] Attributions are characterizations. Characterizations are representations-as. Attributions just *are* forms of representation-as.

---

[21]Fodor does not explain his notion of 'logical form'. See note 9. I regard logic as an account of propositional validity by virtue of propositional structure. Pictures and perceptions are not propositional. Regardless of how one uses the term 'logical form', there are certainly forms of pictorial and perceptual representation that have veridicality conditions and a semantical structure, together with something analogous to a grammar.

[22]*Origins of Objectivity*, *op. cit.*, 74–75.

[23]For detailed discussion of attribution and representation-as in perception, see *Origins of Objectivity*, *op. cit.*, 379–381; 'Origins of Perception', *Disputatio* 4 (2010), 25–28. Fodor makes the further fundamental error of conflating information registration with genuine representation. He calls both 'representation'. He assimilates all iconic representation to information registration. See note 2. For detailed discussion of the distinction, see *Origins of Objectivity*, *op. cit.*, chapter 8; 'Origins of Perception', *op. cit.*, 2–5; 'Perception: Where Mind Begins', *Philosophy* 89 (2014), 385–403; reprinted in T. Honderich ed. *Philosophers of Our Times* (Oxford: Oxford University Press, 2015).

I noted the master methodological error underlying PP(1) and PP(2). It is that of inferring semantical structure directly from iconic format, instead of inferring it from underlying use and representational competencies.[24]

[24]Fodor argues that some psychological states, for example those in what he calls an 'echoic buffer', are non-perceptual iconic representations with semantic content. He applies (a)–(h) to such states. I will not discuss this argument in detail. But it is as off-hand and unsound as the main argument that I have discussed. To begin with, the argument confuses information registration and representation. (See notes 2 and 23.) Although Fodor does not cite examples of what he means by an "echoic buffer", one can assimilate what he says about it to information-registrational states like the first registration of the retinal image. The first registration of the retinal image is iconic, but it is not representation. His claim ('The Revenge of the Given', *op. cit*., 113) that such states must have semantical content because categorization (attribution) is extracted from them is clearly mistaken. Perceptual categorizational attribution is extracted from the initial registration of a retinal image. But that registration lacks semantical content. It is purely information registration.

Fodor argues that an "echoic buffer" is not subject to the "item effect". The *item effect* is 'the rule of thumb that, all else being equal, the "psychological complexity" of a discursive representation (for example, the amount of memory it takes to store it or to process it) is a function of the number of individuals whose properties it independently specifies', *Ibid*, 110–111.

Fodor does not explain 'discursive' clearly, but his explanation, *Ibid*, 107, takes discursive representation constitutively to have all the properties (a)–(h) that he denies of iconic representations. I believe that his accounts of iconicity and discursiveness are both defective. If Fodor's description of the "item effect" were correct, one would expect perceptual representation as well linguistically expressed conceptual representation to show it in memories of such representation. Then Fodor's argument would divide non-representational, information-registrational states (registration of the retinal image, the "echoic buffer")–which do not show an "item effect"–from representational states–both iconic and non-iconic, both perceptual and conceptual–which do. The argument would then fail to bear on the distinction between iconic and non-iconic psychological states.

But the argument has yet further defects, scientific defects. Limitations on memory, even in retaining complex representational states, including perceptual states, vary with the type of memory, not just the representational complexity of the state. Certain types of very short-term, iconic memory retain virtually the full complexity of perceptions' representational content. M. Coltheart, 'Iconic Memory and Visible Persistence', *Perception and Psychophysics* 27 (1980), 183–228. Certain types of unconscious iconic long-term memory are also virtually unlimited in their capacity to retain the complexity of perceptual states or of beliefs formed most directly from perception. Cf. T. Brady, T. Konkle, G. Alvarez, and A. Oliva, 'Visual Long-term Memory has a Massive Storage Capacity for Object Details', *Proceedings of the National Academy of Sciences of the United States of America* 105 (2008), 14325–14329. Even visual working memory, the original poster child for the item effect, is not limited in the way that Fodor assumes. Other factors besides the number of items and representational complexity determine even visual working memory's limitations. For reviews and explanations of why the item effect is not a basic explanatory notion, see C. R. Sims, R. A. Jacobs, and D. C. Knill 'An Ideal Observer Analysis of Visual Working Memory', *Psychological Review* 119 (2012), 807–830; D. Fougnie and G. Alvarez, 'Object Features Fail Independently in Visual Working Memory: Evidence for a Probabilistic Feature-store Model', *Journal of Vision*, 11 (2011), 1–12; G. Bae and J. Flombaum, 'Two Items Remembered as Precisely as One: How Integral Features Can Improve Visual Working Memory', *Psychological Science* 24 (2013), 2038 –2047; K. Hardman and N. Cowan, 'Remembering Complex Objects in Visual Working Memory: Do Capacity Limits Restrict Objects or Features?', *Journal of Experimental Psychology: Learning, Memory, and Cognition* 41 (2015), 325–347; T. Brady and G. Alvarez, 'Contextual Effects in Visual Working Memory Reveal Hierarchically Structured Memory Representations', *Journal of Vision* 15 (2015), 1–24. Fodor's argument that there are iconic representational states in perception is laced with both conceptual and scientific errors.

There is a corollary error–that of identifying representational constituents too closely with parts of a picture, where 'part' is understood in a commonsense way that is not guided by serious semantical reflection. I have criticized versions of this error that take parts of the picture to be representational units, when they are not. The Fodor-Kosslyn claim that one can cut a picture any way one likes embodies this error. Another version of the error is to assume that every representational constituent is a part of the picture, in an unqualified, intuitive sense of 'part'.

Recall some basic points about the relation between parts of *sentences* and semantical constituents. I stipulate that letters and parts of letters are not parts of sentences. Let us assume that words or morphemes are basic parts of sentences. Some parts, so understood, are not representational units. Words embedded in some idioms are not semantic constituents of sentences. Conversely, some representational constituents are not morphemes, words, or word combinations. In context-dependent uses of sentences, representational constituents include occurrent, contextual, referential applications. These are the referential events that constitute occurrent uses of demonstrative-governed phrases. No word or symbol, in the language, expresses or stands in for the occurrent, referential application.

Words like 'that' are schemas. Such words do not represent, demonstratively, any given entity. Occurrent use is needed to carry out such representation. The occurrent event of referential application, not the word itself, is the semantical unit that is central to referential representation. There is no separate part of the sentence that is specific to the occurrent event of application. Ultimately, language *must* rely on applications that are not themselves terms or parts of sentences. They are not themselves symbols. They are events of application.

The same point applies to pictures and perceptual states.[25] In a picture, reference to a particular is effected by an occurrent use. Singular reference to a particular is effected by the intentional act of putting a picture part onto the canvass or interpreting the picture, not by any picture part, token or type, taken on its own. The same colored shape could have represented a different particular, or no particular at all, in a different context. There are no two picture parts, one of which indicates a repeatable color property and another refers to a concrete instance of the property. In a perceptual state, reference to particulars is effected by occurrent events that instantiate a repeatable representational ability-type. No separate symbol or part of

[25]Allocentric maps can avoid context-dependent referential applications of the maps' spatial origin. But many allocentric maps use ordinary proper names, which do involve context-dependent determiners. And even non-context-dependent, canonical names must, in the end, be explained in terms of context-dependent referential applications. Ego-centrically anchored maps all involve referential applications to the "home" anchor position.

the content, of either a picture or a perceptual state, distinguishes the occurrent application from the representational, attributive type(s) that it applies.[26]

Depiction always involves a combination of occurrence-based singular application and a characterizing attribution. The intuitive notion of depiction is not a semantical primitive–at least in cases of pictures that pick out particulars. Similarly, for maps and perceptual representation. All perception is via combinations of singular applications and attributions.

A central respect in which parts in an iconic representation do not correspond to constituents resides in representation of *relations*. In a picture no *part* of the picture specifically represents the depth relation between an object in the foreground and the background. Yet the picture pictures the foreground object as in front of the background.

It is a mistake to identify a specific part of the picture that serves as a representational constituent that represents *any* relation that a picture depicts. If one object is depicted as to the left of another object, with some distance between them, there is no answer as to what *part* of the picture specifically represents the relation *to the left of*. The spatial relation is depicted, but no part of the picture corresponds specifically and proprietarily to the space between the entities. The part of the picture between

[26]Both Fodor in *Lot 2: The Language of Thought Revisited*, *op. cit.*, 175, and J. Quilty-Dunn in 'Iconicity and the Format of Perception', *Journal of Consciousness Studies* 23 (2016), 255–263, take a *term* that symbolizes the singular reference to be required, if an iconic representation (the whole perception or picture) is to represent a particular. Quilty-Dunn writes, '…icons lack the representational apparatus to bind features by picking out an object and attributing those features to the object' (261). But that is exactly what paradigmatic icons, like realist paintings of individuals and perceptions, do.

Green and Quilty-Dunn, 'What is an Object File?', *op. cit.*, make much of the supposed non-iconicity of memory index files. Anaphoric applications in memory–in index files—that derive from referential applications in perception may or may not have a 'symbol' that effects the anaphora. There being such a symbol would not prevent perceptual memory from being iconic.

Indeed, if symbols occur in relations that bear natural correspondences to relations in a subject matter, the arrangement of symbols is iconic. Note the iconic representation of positions of light-rail stations by the arrangement of their names. As the light-rail map example indicates, many iconic representations have non-iconic symbolic elements. Moreover, being a symbol does not preclude being iconic, as hieroglyphs and some Chinese words show. But such a symbol is not needed in perceptual memory that relies on anaphoric retention of perceptual singular reference, any more than a symbol for a referential application is needed in perception, or indeed natural language. Anaphora can be effected through occurrent events that are causally and functionally connected to the occurrent referential application event involved in the original perception. Such a causal link can underlie the changing iconic perceptual attributives that support the application's use in the index file. There need be no symbol for the referential application in the index file. But since iconic representations can have symbolic elements that are non-iconic, the presence of such a symbol in index files would not prevent such files from being iconic.

A significant error in Quilty-Dunn's article is the mis-attribution to me in *Origins of Objectivity*, *op. cit.*, of the view that the difference between perception and cognition consists (sometimes he says 'partly consists', sometimes he does not) in perceptual representations' being iconic, and in cognitive representations' being discursive, or language-like. I think that some cognitive, even propositional, representations are iconic. In fact, I think that *all* propositional beliefs immediately formed from perception are iconic.

the parts that represent the two entities represents other particulars–a background wall, for example. That part is not specific or proprietary to the relation *to the left of*. No picture-part (scattered or not) is a representational constituent specific to the relation. The relation is represented through the relations among the parts, not through a symbol for the relation. The relation *to the left of* is depicted by situating the part of the picture that represents one object to the left of (from the viewer's perspective) the part of the picture that represents the other object. Language has a separate symbol for the relation *to the left of*. The picture does not.

Of course, the iconic arrangement in realist paintings never does the representing all by itself. The arrangement must derive from appropriate uses and competencies.

Analogous points apply for perception and maps. Trying to locate, in non-representational terms, the *part* of a perceptual state, or the underlying neural state–or the part of a map–that indicates and attributes even a simple spatial relation is a quixotic enterprise. The science of visual psychology has suggested no analog for a single symbol (a semantical or syntactical constituent) that in language commonly would indicate a relation and, in the context of a sentence, attribute it. Computations operate on representations' *being* in certain relations. Nothing in computational theory of perception requires that there be a symbol for every relation that is computed.

A simple identification of picture parts with semantical constituents does not work even for iconic indication and attribution of non-relational properties. If asked what part of a picture represents the color, or size, or shape of a surface, it can seem obvious that one can draw a boundary around a relevant picture part. It is intuitive that that part represents the color, size and shape of the surface, as well as the surfacehood of the surface. But this answer is only good enough for off-hand remarks about the relation between parts and semantic constituents. It is the *color* of the picture part that figures in representing the color in the scene. It is the scaled *shape* or *size* of the part that figures in representing the shape or size of the surface. None of these properties of a part of the picture is itself a part of the picture. Each is an aspect of a part.

A similar point applies to visual perception. Different aspects of the perceptual content represent different attributes. Some sensitivity to light intensity, reflected in a natural way in intensity of neural activity, figures in iconic perceptual representation of color. Some space-like arrangement, within a sensory-registration state or perceptual state, of the effect of neural firings figures in the iconic representation of shape. Computations work as well on properties or relations as on object-like "parts". Privileging parts is the effect of thinking of computation in perception too much on an analogy with linguistic computation.

In realist pictures, maps, and normal visual representation, properties are represented in packages. A realist picture and a normal visual perception represent a surface's color, size, shape, orientation, location, and so on, all at once. A road map does not represent a road without representing its position, length, and shape. Such packaging is a normal feature of many types of iconic representation. By contrast, a language user routinely represents a surface without representing its size or color, or a road without representing its position or length. Which properties one talks about is a matter of choice. The packaging-of-properties aspect of some types of perception is

a *type* of iconicity. It is not essential to iconic representation. A color-chit represents a color-shade iconically. It normally represents no other properties. Packaging is not even essential to perception. Certain pathological types of perception may iconically represent color without representing shape or any other spatial attribute.

Even though packaging is usual in iconic representation, loss of some unit in the package, or even depackaging, does not in itself undermine the iconicity of a representation. A drawing could have its color removed and remain an iconic drawing. A map's use of larger circles for larger cities could be discontinued, and the map would remain iconic. Review of the explication of iconicity near the beginning of the article supports this point. The point is also intuitive.

Analogous points underlie the obvious fact that some aspects of an iconic perceptual representation may be better retained than others in perceptual memory, even though the memory is iconic, and even though the different aspects originally came in a "packaged" iconic perceptual representation.[27] Even the complete loss to

[27]E. J. Green and J. Quilty-Dunn, 'What is an Object File?', *op. cit.* state that PP(2) and the following principle are 'the signature markers of iconic format': (H) (for Holism) 'Each part of the presentation represents multiple properties at once, so that the representation does not have separate vehicles corresponding to separate properties and individuals'. Their exposition follows Fodor in conflating information registration and representation. Further, not all iconic representation represents multiple properties. So packaging is not a signature marker of iconic format. As noted, a color chit can be used to represent a color shade, and nothing else. As I indicated, it is problematic to claim that when multiple properties are represented in a package, the representation goes by way vehicles. If vehicles are (say, picture) *parts* or object-like entities, then it is not true that a single vehicle represents multiple properties. The parts are not, strictly, the representing units. If the properties or property-instances count as vehicles, then different "vehicles", not one, effect representation of different properties: the color of the picture part represents the color; the shape of the part represents the shape (scaled appropriately); and so on. In perception, different aspects of a perceptual state represent different environmental attributes, even though properties are represented in a package. So (H) is mistaken in various ways.

Green and Quilty-Dunn argue that since perceptual working memory sometimes retains different properties to different degrees, perceptual working memory cannot conform to (H). They conclude that perceptual working memory is not iconic. They further infer from considerations of simplicity that perceptual representation is not iconic either. Even apart from the mistakes in PP(2) and (H), this train of reasoning seems to me a *reductio* of their conception of iconicity and their conception of how representation takes place in pictures and perception. The fact that various properties are iconically represented in a package in perception does not begin to show that they cannot be remembered iconically to different degrees.

memory of some aspects of a packaged group of representations is compatible with the memory's remaining iconic.[28]

Iconic perceptual memory could, for example, retain iconic representation of shape better than color. A visual perception normally represents size, 2-D and 3-D shape, color, orientation, location, surfacehood, bodihood, all as a package. But since these properties are represented by different aspects of the perceptual state, and some derive from different, dissociable formation processes, remembering these properties to differing degrees, or even forgetting some, does occur. Such occurrence bears not at all on the memory's remaining iconic. Iconicity depends on there being one or more natural mappings. Even when the memory of one property in a package is weaker than that of another property, there remain natural correspondences between the aspects of the memory state that represent the properties and the properties themselves. Iconicity in memory does not depend on memory's preserving all aspects of pictorial representation.

Normal visual perception represents its basic attributes in packages. Doing so is an aspect of the iconic nature of normal visual perception. It is a way in which visual perception is normally picture-like or map-like. But neither perception nor iconicity is essentially picture-like or map-like in this respect. As noted, a visual perception and visual memory could be of a single unlocated color shade, with no packaging and no spatial representation. They would remain iconic, as a color chit is. Auditory perception is iconic but not at all picture-like. The etymology of 'iconic' is connected to visual imagery. But standard dictionary meanings of the term generalize beyond the visual.[29] The connection between iconic representation and pictures or maps is real and paradigmatic. It is not constitutive. What is constitutive is natural mapping from some units of representation to corresponding entities in the subject matter.

Let us return to the idea that because one can cut up a picture into small parts each of which represents something, the semantics of a picture is arbitrary. It is sometimes, though not always, correct to think of pictures, maps, and perceptions as

[28]There is evidence that iconic perceptual working memory in fact retains different features to different extents, as one would expect. Packaging is one factor in retention. But many factors bear on how well different iconically represented attributes are retained in memory. Attention might affect different attributes differently. Facts about how different properties are registered differently in neural coding can also ground differential retention. Relationships among the types of properties retained can affect retention. For papers that bear on these issues, see M. Wheeler and A. Treisman, 'Binding in Short-term Visual Memory', *Journal of Experimental Psychology: General* 131 (2002), 48–64; Y. Jiang, M. Chun, and I. Olson, 'Perceptual Grouping in Change Detection', *Perception & Psychophysics* 66 (2004), 446 – 453; D. Fougnie and G. Alvarez, 'Object Features Fail Independently in Visual Working Memory: Evidence for a Probabilistic Feature-store Model', *op. cit.*; G. Bae and J. Flombaum, 'Two Items Remembered as Precisely as One: How Integral Features Can Improve Visual Working Memory', *op. cit.*; K. Hardman and N. Cowan, 'Remembering Complex Objects in Visual Working Memory: Do Capacity Limits Restrict Objects or Features?', *op. cit.*; T. Brady and G. Alvarez, 'Contextual Effects in Visual Working Memory Reveal Hierarchically Structured Memory Representations', *op. cit.*.

[29]Here is a definition from Merriam-Webster: a sign (such as a word or graphic symbol) whose form suggests its meaning.

containing smallest representational parts–pixel-like representations. It is easy to infer, mistakenly, that the semantics of pictures has no structure. Even if a realist picture is built up from small, primitive, representational parts, it does not follow that the parts can be combined in arbitrary ways. Representation depends on usage, and usage never evinces or allows wholesale arbitrary combinations among representations.

A further point is even more important. Presence of smallest, pixel-like representational parts in a picture, map, or perception does not imply that those parts are representationally primitive. Often the representational content of smallest-parts depends on the representational content of larger wholes. In such cases, pixels need not be primitives. For example, a pixel may have the content: part, in such and such a location, with such and such size and shape, with such and such color, of such and such surface. Such characterizations are consistent with a picture's having compositional, representational structure. (Of course, this characterization is intended as the non-iconic, linguistic counterpart of an iconic characterization.) I think that they often correspond to what is fundamental psychologically. Small parts of an edge are often represented only *as* such parts. So a representation's being analyzable into small representations of small parts (each with a potential use), where the representations occur in a part-whole structure that maps naturally into a part-whole structure in the *representata*, does *not* imply that the pixel-like representations of small parts are representationally, or semantically, primitive.

For example, smallest useable/discernible iconic representations of very small parts of a line, which certainly occur in maps, pictures, and visual perception, need not be primitive representational units. They need not be basic building blocks for composing representations of longer lines. Rather, representations of longer lines with natural end points can be representationally primitive. They accommodate iconic representations of smaller, contained line segments, which lack any natural endpoints, by representing them as parts of the "natural" longer line. Each represented part is distinguished by its represented position within the longer line.

I think that perceptual representation of spatial parts usually works that way. For what is representationally primitive depends on what representational competencies are fundamental. And perceptual competence to distinguish lines with natural endpoints is certainly more fundamental than competence to distinguish the various parts that lack natural boundaries. Representing spatial parts of natural wholes by representing them as parts of, or cuts in, the wholes is compatible with the known fact that, in visual perception, representations of lines with natural endpoints are formed through combining non-perceptual information registrations of smaller edge segments.[30] One must distinguish *information registration* of very small edge segments in proximal stimulation (in the visual image) from *representation* of very small parts of environmental lines with natural end-points. The former are probably

[30]See W.S. Geisler, J.S. Perry, B.J. Super, and D.P. Gallogly, 'Edge Co-occurrence in Natural Images Predicts Contour Grouping Performance', *op. cit.*; J. Frisby and J. Stone, *Seeing: The Computational Approach to Biological Vision*, *op. cit.*, chapter 6.

first in the order of perception formation. The latter are derivative in the structure of perceptual representation.

Iconic representation occurs in language and thought, as well as maps, pictures, and perception. Language sometimes incorporates pictorial elements. Propositional thought often incorporates perceptual elements. Although pictures and perceptions are not propositional, their iconic formats can be coopted in propositional structures, making units in those structures iconic. The most basic perceptual beliefs, for example, are iconic. The popular view that iconicity is disjoint with propositionality is mistaken.

Although most iconic representation is iconic partly in its relational structure, iconic representation does not essentially represent by way of relational structure: one could invent an iconic word for red that consists in a red color chit.

Iconic representations, like all representations, represent by being the product of a use and competence. Representational structure marks the representational functions of the psychological competencies of the users of representation. Representational structure is found by investigating the uses by and competencies of those who use the representations. Iconicity is an aspect of representational format, not a determiner of basic semantical functions. The basic semantical functions of iconic representation–reference and attribution–are shared with non-iconic representations in language and in non-iconic thought. Many, but not all, linguistic modes of presentation and representational content differ from iconic modes of presentation and representational content. Modes of presentation and representational content mark types of competence. Iconic and non-iconic competencies are psychologically different. But basic semantical functions and basic semantical or representational structures are shared.

What is distinctive of iconic representation is that it relies on natural relations between the representation and the represented subject matter. This reliance is evolutionarily very old. It is where representation begins. Representation begins in nature's stamping itself directly into a creature's capacities. The structural correspondences that result from the stamping–inasmuch as they are law-like and reliable–contribute not only to primitive representation, but also to primitive ancestors of epistemic warrant and knowledge. The effects of this stamping remain in language and abstract thought, which in many ways have outgrown iconic representation.