

Depiction and Composition

Abstract: Traditionally, the structure of a language is revealed by constructing an appropriate theory of meaning for that language, which exhibits how – and whether – the meaning of sentences in the language depends upon the meaning of their parts. In this paper, I argue that whether – and how – what pictures represent depends on what their parts represent should likewise be revealed by the construction of appropriate theories of representation for the symbol system of those pictures. This generalisation, I argue, reveals a much cited disanalogy between depiction and description is illusory: the structure of pictures, like language, is compositional.

I. Introduction

Language has compositional structure. The meaning of ‘Theaetetus flies’ for example depends on the meaning of ‘Theaetetus’, the meaning of ‘flies’ and the order in which they are concatenated. In general, the meaning of a sentence depends on the meanings and arrangement of its parts. In contrast, depiction is supposed to lack compositional structure. The Mona Lisa, for example, is supposed not to be divisible into parts in the way that ‘Theaetetus flies’ is divisible into ‘Theaetetus’ and ‘flies’. In general, what a picture represents is supposed not to depend on what its parts represent in the way that what a sentence means is supposed to depend on what its parts mean.

Despite the platitudinousness of this observation, a sense in which what is represented by pictures doesn’t depend on what is represented by their parts is difficult to discern. The Mona Lisa represents Lisa, for example, in part because parts of the Mona Lisa represent parts of Lisa; if its left and right half didn’t represent her left and right half, for example, then it as a whole wouldn’t have represented her as a whole. If there’s a sense in which what pictures represent does not depend on what their parts represent in the way that the meaning of sentences does depend on what their parts mean, that sense cannot simply be that pictures do not have representational parts.

So the supposed disanalogy between depiction and description is not that pictures do not have representational parts, but that the division of a picture into representational parts is arbitrary in a way that the division of a sentence into representational parts is

not. The meaning of ‘Theaetetus flies’ depends on the meaning of ‘Theaetetus’ and of ‘flies’ in a way that what the Mona Lisa represents doesn’t depend on what its left and right halves represent, for example, because the division of ‘Theaetetus flies’ into ‘Theaetetus’ and ‘flies’ is natural in a way that the division of the Mona Lisa into its left and right halves is arbitrary.

This disanalogy between depiction and description is constantly cited in the literature. Roger Scruton (1987, 107), for example, claims “there seems to be no way in which we can divide [a] painting into grammatically significant parts – no way in which we can provide a syntax which isolates those parts of the painting that have a particular semantic role.” Likewise, David Braddon-Mitchell and Frank Jackson (1996, 180) write “there is no preferred way of dividing [a] map into basic representational units. There are many jigsaw puzzles you might make out of the map, but no one would have a claim to have pieces that were all and only the most basic units.”

Similarly, Jerry Fodor (2008, 174) writes “Iconic representations ... have no canonical decomposition; which is to say they have no canonical structure; which is to say that, however they are sliced, there’s no distinction between their canonical parts and their mere parts.” Even Roberto Casati and Achille Varzi (1999, 191), who are sympathetic to an extremely close analogy between depiction and description, ask “Suppose you have a uniformly coloured map region: is it composed of its left and right halves or is it composed of its top and bottom halves?” Despite this consensus, this paper argues this disanalogy is illusory: depiction is compositional in the same sense description is.

The refrain that the division of a picture into representational parts is arbitrary echoes an earlier refrain from the philosophy of language, according to which the division of sentences into meaningful parts is equally arbitrary. Almost exactly the same point is raised by Willard Quine (1970, 392), for example, about language when he writes “... suppose again a language for which we have two extensionally equivalent systems of grammar ... According to one of these systems, the immediate constituents of a certain sentence are ‘AB’ and ‘C’; according to the other system they are ‘A’ and ‘BC’. ... which is right?” (Quine, 1970, 392).

But whereas this point is still widely accepted in the philosophy of pictures, it's now rarely accepted in the philosophy of language. As Fodor (2008, 172-3), for example, writes in the same passage quoted above "... 'John', 'Mary', and 'loves Mary' are among the constituents of [John loves Mary] ... But 'John loves' isn't, and nor is 'John ... Mary'." If Fodor is right about this, there must be a way to decide between theories of English which agree about its sentences but disagree about their division into meaningful parts. That there is is supposedly uncontroversial: "Further details are available upon application at your local department of linguistics" (Fodor, 2008, 172).

Further details, I argue, do show there is a way to decide between competing theories of a language which agree about its sentences but disagree about their division into meaningful parts – a theory according to which the meaningful parts of 'John loves Mary' include 'loves Mary' but not 'John loves', for example, meets constraints that a theory according to which the meaningful parts include 'John loves' but not 'loves Mary' doesn't. But the same constraints which must be met for a theory of meaning for a language to properly reflect its structure, I shall argue, also reveal that there are non-arbitrary divisions of pictures into their representational parts.

I'll consider three constraints on theories of representation – the finite axiomatization constraint, the mirror constraint and the structural constraint – and argue that only the structural constraint ensures that a theory of representation reveals how and whether what a representation is of depends on what's represented by its parts. Neither the finite axiomatization constraint nor the mirror constraint entail that theories of representation for depictive symbol systems should be compositional, but – I'll argue – the structural constraint does. Language has compositional structure. Pictures have compositional structure of the same kind.

Three clarifications. First, it's often argued that pictures cannot have a compositional semantics on the grounds that they do not have a compositional grammar. Scruton (1987, 107), for example, writes "While there may be repertoires and conventions in painting, there is nothing approaching a grammar as we understand it." And Flint Schier (1986, 66) writes "Pictures, by contrast, have no grammatical rules, natural or conventional." But whether or not pictures have compositional grammar, I will argue

in the conclusion, turns on the same kind of consideration as whether or not they have compositional semantics, so the question cannot be resolved so quickly.

Second, I will focus on two highly simplified examples – chess diagrams and maps. The example of chess diagrams is intended to simply and uncontroversially illustrate the thesis, whereas the example of maps extends the argument to a more controversial case – one that is taken as paradigmatically non-language-like (Braddon-Mitchell & Jackson, 1996, 180; Lewis, 1994, 310). Just as semantics for natural languages begin with simplified fragments of those languages, semantics for pictures must begin with simplified fragments of depictive symbol systems (Casati & Varzi, 1999, 187). But it is hoped that the arguments generalise in principle to other kinds of depiction.

Third, different analyses disagree over which representations to classify as depictions. Structural theories such as Nelson Goodman's (1968) and John Kulvicki's (2006), for example, exclude chess diagrams. Experiential theories such as Richard Wollheim's in terms of seeing-in or Robert Hopkins in terms of experienced resemblance exclude maps (Wollheim, 1987, 60-61; Hopkins, 1998, 30). Dominic Lopes' recognitional theory (1996, 5) includes maps, whereas Michael Newall's (2011, 1) recognitional theory does not. I presuppose a version of the resemblance theory which classifies both chess diagrams and maps as depictions (author's papers; Abell, 2009).

But however one distinguishes depiction from other kinds of representation, there is an interesting question about the how far the phenomena of compositionality extends beyond the traditional example of language. As the quote from Braddon-Mitchell and Jackson above illustrates, conventional wisdom is that compositionality does not even extend to the relatively simple example of maps, let alone to more complex varieties of depiction. And, as the quotes from Scruton and Fodor illustrate, the arguments for denying that depictions are compositional are the same as the arguments for denying that maps are. If what I say below is correct, all of these arguments are unsound.

II. Theories of Meaning

A theory of meaning for a language is a theory which entails, for each sentence in the language, a statement of the meaning of that sentence (Davidson, 1967, 22). A theory

of meaning for English, for example, should entail statements such as that ‘snow is white’ in English means that snow is white and that ‘grass is green’ in English means that grass is green. Likewise, a theory of meaning for German should entail statements such as that ‘es regnet’ in German means that it’s raining and that ‘schnee ist weiß’ in German means that snow is white. If a theory of meaning for a language is adequate, it should reveal the structure of sentences in the language.

In general, a theory of representation for a symbol system is a theory which entails, for each character in the symbol system, a statement about what that character represents. A theory of representation for the symbol system of Arabic numerals, for example, should entail statements such as that ‘1’ represents one, that ‘2’ represents two, that ‘3’ represents three, ... and so on. Likewise, a theory of representation for the symbol system of traffic lights should entail that red represents stop, that orange represents slow and that green represents go. If a theory of representation for a symbol system is adequate, it should reveal the structure of that symbol system.

A theory of representation is compositional if and only if the statement of what each character represents is derived from axioms which state the contribution of the parts of the character, and the significance of their arrangement. The statement that ‘snow melts’ in English means that snow melts, for example, might be derived in a compositional theory of English from an axiom stating that ‘snow’ refers to snow and an axiom stating that a referring term followed by ‘melts’ means the referent of that term melts. Likewise, in a compositional theory of representation for chess diagrams, the theorems which state what each diagram represents might be derived from axioms stating what the two colours, six figurines and sixty-four squares represent.

So just as a theory of meaning may reveal the structure of a language, a theory of representation may reveal the structure of a depictive symbol system. According to the compositional theory of representation for maps proposed by Roberto Casati and Achille Varzi, for example, what a map represents depends compositionally on what its atomic map stages represent, which in turn depends compositionally on what its colours and regions represent. If this theory of representation for maps were adequate, then – like a theory of meaning for a language – it would reveal that the structure of maps – like the structure of language – is compositional.

A map stage, according to Casati and Varzi (1999, 192), is any colouring of a map's regions. A map stage is atomic if and only if it colours all and only the regions of a single shade (1999, 192). If, for example, the whole of the region representing France is coloured purple, then the colouring of that region is an atomic map stage. In contrast, if the region representing Vichy France is coloured purple, and the region representing occupied France is coloured red, then the red and purple colouring of the region representing France is not an atomic map stage. And if the region representing the British Empire is coloured pink, then this colouring is also an atomic map stage.

An atomic map stage is true, according to the theory, if and only if (a) it colours a region of the map which represents a region of the world which has the property represented by its colour and (b) the region of the world represented by the rest of the map does not have that property (Casati & Varzi, 1999, 194; Rescorla (2009) defends (b)). So the maximal blue colouring of the world map, for example, is true if and only if the region it colours represents a region covered by ocean, and the rest of the map represents a region which does not have the property of being covered by ocean.

A map, according to the theory, is true if and only if all its atomic map stages are true (Casati & Varzi, 1999, 195). So the world map, for example, is true if and only if its maximal green and blue colourings are both true. Casati and Varzi's theory reveals compositional structure in the symbol system of maps: map regions are like names referring to world regions, colours are like predicates representing properties, atomic map stages are like atomic sentences predicating properties of regions, and whole maps are like complex sentences conjoining atomic map stages. So the structure of maps, according to the theory, is closely analogous to the structure of language.

One clarification. Strictly speaking, Casati and Varzi's theory is not a theory of representation which entails, for each map, a statement of what that map represents but a theory of truth which entails, for each map, a statement of the truth-conditions of that map. The theory entails that the world map, for example, is true if and only if the region represented by its blue part is covered by ocean and the region represented by its green part is covered by land, but not that the world map represents that the region represented by its blue part is covered by ocean and the region represented by its

green part is covered by land.

An interpretive theory of truth is one in which ‘is true if and only if’ in its theorems can be correctly replaced by ‘represents that’ (Davies, 1981a, 34). A theory which entails that ‘snow is white’ is true if and only if snow is white and the earth moves, for example, is not interpretive, because ‘snow is white’ does not represent that snow is white and the earth moves. So if a theory of truth is interpretive, one should not be able to infer from ‘snow is white’ is true if and only if snow is white that ‘snow is white’ is true if and only if snow is white and the earth moves. As long as a theory of truth is interpretive, it may serve as a theory of representation.

If Casati and Varzi’s theory of truth for maps is interpretive, then from its statements about the truth-conditions of maps one can infer statements about what those maps represent. So as long as one is not able to infer from, for example, the fact that the world map is true if and only if the region its blue part represents is ocean and the region its green part represents is land that the world map is true if and only if the region its blue part represents is ocean, the region its green part represents is land and the earth moves, one may infer that the world map represents that the region its blue part represents is ocean and the region its green part represents is land.

One objection. It might be argued that because some pictures cannot be finitely paraphrased, what they represent can’t be stated, so that no theory of representation could entail statements about what those pictures represent. A theory which entailed, for example, simply that a photograph of a cup on a table represents that a cup is on the table would be incomplete, since the photograph would also have to represent that the cup is smaller than the table, whiter than the table, curvier than the table, ... and so on. So although it may be possible to construct theories of representation for some symbol systems such as maps, this may not be possible in general.

This objection can be avoided in two ways. First, a theory may completely specify what a picture that can’t be finitely paraphrased represents if it is allowed to entail an infinite number of statements about what that picture represents. A theory of representation could entail, for example, an infinite number of theorems which combine to state that the photograph of the cup on the table represents that the cup is

smaller than the table, whiter than the table, curvier than the table, ... and so on. Even if what a picture represents cannot be paraphrased by a single sentence, it may be paraphrased by an infinite number of sentences.

Second, a statement of what a picture represents may be made using that very picture, just as what a sentence means may be stated by using that very sentence. So instead of entailing a theorem which states in English what is represented by a photograph of a cup, a theory of representation may entail a theorem of the form: _____ represents _____, where the second blank is replaced by the picture of the cup, and the first is replaced by a picture of that picture (framed, instead of in quotes). Since the picture itself is used in the statement of what it represents, it cannot fail to be an accurate paraphrase, and no problem is posed by its being unparaphrasable in English.

III. The Finite Axiomatization Constraint

Just as a theory of meaning for a language should reveal how and whether the meanings of sentences in that language depend on the meanings of their parts, a theory of representation for a symbol system should reveal whether and how what its characters represent depends on what their parts represent. A theory of representation for chess diagrams, for example, should reveal whether and how what diagrams represent depends on what's represented by the figurines and their arrangement and a theory of representation for maps should reveal whether and how what maps represent depend on what their parts represent.

But a theory of meaning for a language may entail what each sentence in the language means without revealing how or whether the meanings of sentences depend on the meanings of their parts. A theory of meaning for English, for example, might simply list an infinite number of axioms which state separately the meaning of each English sentence (Davidson, 1970, 56). Likewise, a theory of representation for pictures might simply list an infinite number of axioms of the form: _____ represents _____, where the second blank is replaced by each picture, and the first by a picture of that picture (framed, instead of in quotes).

It's sometimes suggested that to exclude trivial theories of meaning or representation

of this kind which fail to reveal the structure of a language or symbol system, they should be constrained to a finite number of axioms (Davidson, 1970, 56). Since, for example, both the theory of meaning for English which simply lists an infinite number of axioms which state separately the meaning of each English sentence and the theory of representation for pictures which simply lists a separate axiom which states what is represented by each picture both possess an infinite number of axioms, both are rightly excluded by imposing this constraint.

The finite axiomatization constraint reveals a lacuna in Casati and Varzi's theory of representation for maps. To entail a statement of what each map represents, the theory must entail statements about what each atomic map stage represents. And to entail what each atomic map stage represents, the theory should entail what property each colour represents and what world region each map region which is coloured by an atomic stage represents. The natural way to do so is to add an axiom for each colour stating which property it represents and an axiom for each map region coloured by an atomic map stage stating which world region it represents.

To entail what's represented by the world map, for example, four axioms could be added: an axiom stating that green represents the property of being covered by land, an axiom stating that blue represents the property of being covered by ocean, an axiom stating which part of the world the blue coloured part of the map represents and an axiom stating which part of the world the green coloured part of the map represents. The theory would then entail that the world map represents that the region represented by its blue part is covered by ocean and the region represented by its green part is covered by land.

If maps use only a finite number of colours, then only a finite number of axioms will be required to state which property each colour represents. But since every difference in shape, size and location is a different region, there's an infinite number of regions which atomic map stages may colour, so an infinite number of axioms would have to be added to state which world region each map region which may be coloured by an atomic map stage represents. If this were the case, then Casati and Varzi's theory would not meet the finite axiomatization constraint, despite being a compositional theory which states what maps represent in terms of what their parts represent.

Perhaps with this problem in mind, Casati and Varzi suggest two constraints on which map regions represent which world regions. First, a map region is part of another if and only if the world region the former represents is part of the world region the latter represents (Casati & Varzi, 1999, 194). So the part representing France, for example, must be a part of the part representing Europe. Second, one map region is connected to another if and only if the world region the former represents is connected to the world region the latter represents (Casati & Varzi, 1999 194). So the part representing Italy, for example, must be connected to the part representing France.

But these constraints don't resolve the problem, because they don't entail which regions of the map represent which regions of the world, but only which regions of the map represent which regions of the world given which other regions of the map represent which other regions of the world. The first constraint entails, for example, that a map represents the world if and only if its halves represent the hemispheres, but not that it does represent the world nor that its halves do represent the hemispheres. Likewise, the second constraint entails that the parts representing Italy and France are connected, but not which represents which.

Despite revealing a lacuna in Casati and Varzi's theory, the finite axiomatization constraint does not ensure that a theory of representation for a symbol system reveals how or whether what its characters represent depends on what their parts represent. A theory of meaning for English might, for example, have a single axiom consisting of an infinitely long conjunction, the conjuncts of which are separate statements, for each English sentence, of what that sentence means. Such a theory would be finitely axiomatized, but still fail to reveal how or whether the meanings of English sentences depend on the meanings of their parts (Davies, 1981a, 61).

Likewise, a theory of representation for pictures might have a single axiom consisting of an infinitely long conjunction, the conjuncts of which are separate statements of the form: _____ represents _____, where the second blank is replaced by each picture, and the first by a picture of that picture (framed, instead of in quotes). Such a theory would be finitely axiomatized, but would still fail to reveal whether and how what pictures represent depends on what their parts represent. So a theory of representation

for a symbol system may meet the finite axiomatization constraint, without revealing the structure of that symbol system or whether it is compositional.

Excluding infinitely long conjunctions as well as infinitely many axioms does not resolve the problem. Substitutional quantification, for example, can be used to state a theory of meaning for English whose only axiom is: $(\prod\phi) \lceil\phi\rceil$ in English means that ϕ (Davies, 1981a, 62). Likewise, a theory of representation for pictures could be given by a single axiom which states that all statements of the form _____ represents _____, where the second blank is replaced by each picture, and the first by a picture of that picture (framed), are true. Both theories are finitely axiomatized, but don't reveal how or whether what complexes represent depends on what their parts represent.

Infinitary conjunction and substitutional quantification raise controversial issues, but the same point can be made by considering finite languages and symbol systems. Take, for example, a language with just ten names and ten predicates. The finite axiomatization constraint cannot be used to decide between a theory of meaning for this language with one hundred distinct axioms which state the meaning of each sentence separately, and a theory of meaning with just twenty axioms which state the contribution made by each name and each predicate to the meaning of sentences which contain them (Evans, 1981, 326-328).

Likewise, since there is only a finite number of positions in chess, the finite axiomatization constraint cannot decide between a theory of representation for chess diagrams with a large but finite list of axioms which state what each diagram represents separately, or a smaller number of axioms which state the contribution made by each figurine to what is represented by the diagrams which contain them. So a theory of representation for a symbol system may meet the finite axiomatization constraint, without revealing whether or how what the characters of that symbol system represent depends on what is represented by their parts.

IV. The Mirror Constraint

The failure of the finite axiomatization constraint to favour a twenty over a hundred axiom theory of meaning for a language with just ten names and ten predicates is

often taken to motivate imposing a constraint according to which a theory of meaning for a language, or of representation for a symbol system, should mirror the structure of the ability of interpreters to understand it. According to this mirror constraint, the axioms of a theory which entail what $s_1...s_n$ represent should entail what s represents if and only if people with the ability to understand what $s_1...s_n$ represent can understand what s represents without further training (Davies, 1983, 15).

Since, for example, people with the ability to understand what ‘John is happy’ and ‘Harry is sad’ mean can understand what ‘John is sad’ means without further training, the mirror constraint entails that the axioms of a theory of meaning which entail what ‘John is happy’ and ‘Harry is sad’ mean should also entail what ‘John is sad’ means. So in the case of the language with just ten names and ten predicates, the mirror constraint favours a theory of meaning with just twenty axioms stating the contribution made by each name and each predicate over the theory of meaning with one hundred distinct axioms which state the meaning of each sentence separately.

Likewise, since people with the ability to understand the chess diagram illustrating the opening position can understand a diagram illustrating any other position, the mirror constraint entails that the axioms of a theory of representation for chess diagrams which entail what the diagram of the opening position represents should also entail what is represented by the diagrams illustrating every other position. So the mirror constraint favours a theory which states the contribution made by each figurine to what is represented by the diagrams which contain them over a theory with a large but finite list of axioms which state what each diagram represents separately.

It might be suggested that people with the ability to understand one picture can understand any picture without further training (Schier, 1986, 43). In this case, the mirror constraint would favour a theory of representation for pictures with a single axiom, such as the theory consisting of an infinitely long conjunction, the conjuncts of which are separate statements of what each picture represents or a theory of representation for pictures with a single axiom which states that all statements of the form _____ represents _____, where the second blank is replaced by each picture and the first by a picture of that picture (framed, instead of in quotes), are true.

But two qualifications are required to the suggestion that people with the ability to understand one picture can understand any picture without further training. First, one may have the ability to understand one picture, without being able to understand pictures in other styles or symbol systems (Schier, 1986, 46-48). Someone with the ability to understand chess diagrams, for example, may be unaware of the conventions governing contour lines and thus be unable to understand topographical maps. And someone with the ability to understand impressionist paintings, for example, may still lack the familiarity required to understand cubist paintings.

If it's true that people with the ability to understand a picture can understand every picture in the same symbol system without further training, then the mirror constraint would favour separate single axiom theories of representation for each symbol system. For the symbol system of chess diagrams, for example, the mirror constraint might favour a theory of representation consisting of a long but finite conjunction, the conjuncts of which are separate statements of what each chess diagram represents. So even in this case, the mirror constraint would still favour non-compositional theories of representation for depictive symbol systems.

Second, one may have the ability to understand a picture, without having the ability to understand every picture in the same symbol system, because some pictures in the symbol system depict things one lacks the ability to recognise (Schier, 1986, 44). If you don't have the ability to recognise armadillos, for example, then you may not be able to understand a picture of an armadillo either, even if you have the ability to understand other pictures in the same symbol system. So it's not the case that people with the ability to understand a picture can understand every picture in the same symbol system, and the mirror constraint may not favour single axiom theories.

If people with the ability to understand a picture have the ability to understand any picture in the same symbol system which depicts something they're able to recognise, the mirror constraint will sometimes favour compositional theories. If one's able to recognise a chess piece, for example, one's able to understand the figurine which represents that piece. And if one understands one chess diagram, one can understand other chess diagrams which contain only figurines one understands without further training. So the mirror constraint would favour a theory which states what diagrams

represent in terms of the contribution of each figurine.

But if people with the ability to understand a picture can understand any picture in the same symbol system which depicts something they're able to recognise, the mirror constraint will not always favour compositional theories. If one is able to recognise the sex of chickens, for example, and one understands a picture of a female chick, one may be able to understand a picture of a male chick in the same symbol system with no further training. This may be true even if the picture of the male and of the female chick have no parts in common, because the two pictures differ in their shape, colour and other representational properties.

In this case, the mirror constraint wouldn't favour a compositional theory with axioms which entail a statement of what is represented by the picture of the female chick and a statement of what is represented by the male chick by stating the contribution of the common parts of those pictures, because those pictures have no parts in common. In general, the mirror constraint does not always favour compositional theories, because sometimes people with the ability to understand some pictures are able to understand another picture in the same symbol system which represent things they can recognise without further training, even if none of the pictures have parts in common.

However, the mirror constraint is not the appropriate constraint to impose on theories of representation, since if a language is spoken by a psychologically unusual population, a theory which meets the mirror constraint may nevertheless fail to reveal how, or whether, the language is structured. If the language with just ten names and ten predicates, for example, were spoken by a dim-witted population who had to learn the meaning of each sentence in the language individually, then the mirror constraint would favour the hundred axiom theory which states what each sentence means over the twenty axiom theory which states the contribution of each name and predicate.

Similarly, if a language were spoken by a population which was so wired-up that familiarity with any one sentence of the language triggered knowledge of every sentence of the language, then the mirror constraint would favour a theory of meaning for that language such as that with a single long conjunction, the conjunctions of which are separate statements of the meaning of each sentence. So the mirror

constraint favours theories of representation which reflect the psychology of a symbol system's users over theories of representation which reveal whether, and how, what its characters represent depends on what their parts represent.

V. The Structural Constraint

The failure of the finite axiomatization constraint and the mirror constraint to favour structure revealing theories motivates the imposition of a structural constraint, according to which the axioms of a theory which entail what $s_1...s_n$ represent entail what s represents if and only if what s represents can be inferred by rational inductive means from what $s_1...s_n$ represent. The axioms of a theory of meaning which entail what 'John loves Mary' means, for example, should entail what 'Mary loves John' means because it's possible to infer what 'Mary loves John' means from what 'John loves Mary' means by rational inductive means (Davies, 1981a, 56).

The structural constraint favours theories of meaning which reveal how and whether the meanings of sentences in a language depend on the meanings of their parts, even when those languages are spoken by psychologically unusual populations. Since the inference from what 'John is sad' and 'Harry is happy' mean to what 'John is happy' means is inductively strong, for example, the structural constraint favours a theory of meaning with separate axioms stating the contribution made by 'John', 'Harry', 'is sad' and 'is happy' to the meaning of sentences containing them, so the structural constraint favours compositional theories of meaning for languages.

Likewise, since it's possible to go by rational inductive means from knowledge of what the chess diagram of the opening position represents to knowledge of what any other chess diagram represents, the structural constraint entails that the axioms of a theory of representation which entail what the opening position represents should also entail what is represented by any other chess diagram. So the structural constraint favours a compositional theory of representation for chess diagrams, with axioms which state the contribution of each figurine to what's represented by the diagrams which contain them, over non-compositional theories.

If people with the ability to understand a picture can understand any picture in the

same symbol system which depicts something they're able to recognise, the structural constraint will not cease to favour compositional theories, because this ability is not explained by the possibility of going by rational inductive means from knowledge of what a picture represents to knowledge of what other pictures represent. Rather, this ability is explained by unusual features of the psychology of picture interpreters: for example, by the thesis that the ability to understand a picture of a thing is underlain by just the ability to recognise the thing and competence in the picture's symbol system.

If one is able to recognise the sex of chickens, for example, and one understands a picture of a female chick, one may be able to understand a picture of a male chick in the same symbol system with no further training, even if the pictures of the male and of the female chick have no parts in common. But because one does not proceed from understanding the picture of the female chick to understanding the picture of the male chick by rational inductive means, the structural constraint doesn't favour a theory the axioms of which which entail what the picture of the female chick represents entail what the picture of the male chick represents, but may favour a compositional theory.

It might be objected that in going from knowledge of what a picture represents to knowledge of what another picture in the same symbol system of something one's able to recognise represents, one does proceed by rational inductive means, since if a picture engages one's ability to recognise something, it probably depicts that thing. If one knows, for example, what a picture of a female chick represents and has the ability to sex chickens, and if that ability is engaged by a picture in the same symbol system as the picture of the female chick, then one would be rational to infer that the picture which engages one's ability to recognise a male chick depicts a male chick.

However, the phrase "rational inductive means" in the structural constraint should not be construed to allow bringing to bear general knowledge, such as the knowledge that pictures which engage one's ability to recognise something probably depict that thing. Etymological knowledge, for example, would trivialise the constraint (Davies, 1981b, 141): if one knows what 'chickens roost' means, for example, then knowledge of etymology would allow one to infer what 'chooks roost' means, but the axioms of a theory which entail what 'chickens roost' means should not entail what 'chooks roost' means, since different axioms should state the contributions of 'chicken' and 'chook'.

The structural constraint also reveals the lacuna in Casati and Varzi's theory of maps revealed by the finite axiomatization constraint. If an axiom for each colour stating which property it represents and an axiom for each map region coloured by an atomic map stage stating which world region it represents is added to the theory, then the axioms which entail what some maps represent also entail what other maps composed of the same atomic map stages represent. The axioms which entail what the world map represents, for example, also entail what is represented by the map of just the ocean and what is represented by the map of just the land.

But it's also possible to proceed by rational inductive means from knowing what a map represents to knowing what other maps composed of different atomic map stages represent. If slightly more of the world map were coloured blue, for example, then its atomic map stages would be colourings of different regions, so the axioms stating what the world map represents would not entail what this slightly different map represents. But it'd be possible to proceed by rational inductive means from knowing what the world map represents to knowing what the slightly different map represents, so the structural constraint entails the same axioms should entail what both represent.

So the structural constraint favours not a theory of representation for maps which has axioms which state which world regions are represented by map regions coloured by atomic map stages, but a theory of representation for maps which states which map regions represent which world regions using a coordinate system. Since a coordinate system would entail what every point on the map represents, it would accommodate the fact that if one understands what the world map represents, one can go by rational inductive means to understanding what a map slightly more of which is coloured blue represents – that slightly more of the world is covered by ocean.

Braddon-Mitchell and Jackson (1996, 180) argue maps lack compositional structure because “there is no natural *minimum* unit of truth-assessable representation in the case of maps ... part of a map that stands for a city itself stands for part of that city.” In the same vein, Roger Scruton (1983, 107) writes that “... the parts themselves are understood in *precisely the same way*; that is, they too have parts, each of which is potentially divisible into significant components, and so on ad infinitum” and Gregory

Currie (1996, 130) writes that “There are no atoms of meaning for cinematic images; every temporal and spatial part of the image is meaningful...”.

But a theory of representation for maps which states which map regions represent which regions of the world using a coordinate system would accommodate the point that all the parts of maps are representational, while still revealing the compositional structure of maps. And a theory of representation for maps which states which map regions represent which world regions using a coordinate system isn't arbitrary but is forced upon us by the constraint that the axioms of a theory which entail what $s_1...s_n$ represent entail what s represents if and only if what s represents can be inferred by rational inductive means from what $s_1...s_n$ represent.

VI. Conclusion

I've considered three constraints on theories of representation – the finite axiomatization constraint, the mirror constraint and the structural constraint – and argued that only the structural constraint ensures that a theory of representation for a symbol system reveals whether and how what the characters of that symbol system represent depends on what their parts represent. The finite axiomatization constraint does not ensure that a theory of representation reveals whether or how what the characters of a symbol system represent depends on what their parts represent, because it can be met trivially by theories with only a single axiom.

The mirror constraint does not ensure that a theory of representation for a symbol system reveals whether or how what its characters represent depends on what their parts represent, since if the symbol system is used by idiosyncratic people, the theory will reflect their idiosyncratic psychology instead of the actual structure of characters in the symbol system. If the psychology of depictive representation is different from the psychology of descriptive representation, then theories of representation for depictive symbol systems conforming to the mirror constraint would reflect these differences, instead of revealing whether and how depiction is compositional.

The structural constraint does suggest that a theory of representation for a symbol system reveals whether and how what its characters represent depends on what their

parts represent. And just as theories of meaning which conform to the structural constraint reveal that the meanings of sentences depend on the meanings of their parts, theories of representation which conform to the structural constraint reveal that what pictures represent depends on what their parts represent, since one may proceed by rational inductive means from knowing what some pictures represent to knowing what other pictures composed of the same parts represent.

The same considerations drawn on above to support the thesis that pictures have compositional semantics support the thesis that they have compositional grammar. Just as, for example, there is both a twenty axiom and a one hundred axiom theory of meaning for the one hundred sentence language, there is a grammar for that language with ten names and ten predicates as its basic expressions and another grammar with all one hundred sentences as its basic expressions. The former grammar and not the latter reveals the grammatical structure of the language, since only the former meets the grammatical analogue of the structural constraint (Davies, 1981b, 158).

Likewise, there is a grammar for chess diagrams which lists every diagram as a basic expression and another grammar for chess diagrams which takes figurines and squares as basic expressions. The latter grammar and not the former reveals the grammatical structure of chess diagrams, because only the latter meets the grammatical analogue of the structural constraint. So the idea that pictures lack compositional grammar can't be used to argue they lack compositional semantics, because the same considerations which support the thesis that they have compositional semantics support the thesis that have compositional grammar as well.

(A theory of meaning is sometimes defined as compositional if and only if its axioms and theorems correspond to axioms and theorems of the appropriate grammar for the language. If the theorem stating that 'John is happy' means that John is happy, for example, is derived from an axiom stating that 'John' refers to John, an axiom stating that 'is happy' expresses being happy, and an axiom stating that a name concatenated with a predicate means that the name's referent has the property expressed by the predicate, then the theorem of the grammar stating 'John is happy' is a sentence should be derived from three axioms stating that 'John' is a name, that 'is happy' is a predicate, and that a name concatenated with a predicate is a sentence.

I didn't adopt this definition above, because then the question of whether a theory of representation for a depictive symbol system is compositional would have depended on the question of which is the appropriate grammar for the symbol system. The theory of representation which simply lists what each picture represents with a separate axiom, for example, would count as compositional with respect to the grammatical theory which simply lists with a separate axiom for each picture that it is a picture. But if we concluded both that the appropriate theory of representation and grammatical theory for pictures are to be given by such lists, that wouldn't reveal the structure of pictures is compositional!)

References

Abell, Catharine (2009), "Canny Resemblance: An Account of Depiction", *Philosophical Review*, 118(2): 183-223.

Braddon-Mitchell, David and Frank Jackson (1996), *Philosophy of Mind and Cognitive Science* (Cambridge, MA: Blackwell Publishers).

Casati, Roberto and Achille Varzi (1999), *Parts and Places: The Structures of Spatial Representation* (Cambridge, Mass.: MIT Press).

Currie, Gregory (1995), *Image and Mind: Film, Philosophy and Cognitive Science* (Cambridge: Cambridge University Press).

Davidson, Donald (1967), "Truth and Meaning", *Synthese* 17: 304-23. Reprinted in Davidson (1984): 17-36.

Davidson, Donald (1970), "Semantics for Natural Languages" in *Linguaggi nella Società e nella Tecnica* (Milan: Edizioni di Comunità). Reprinted in Davidson (1984): 55-64.

Davidson, Donald (1984), *Inquiries into Truth and Interpretation* (Oxford: Oxford University Press).

Davies, Martin (1981a), *Meaning, Quantification and Necessity* (London: Routledge and Keegan Paul).

Davies, Martin (1981b), “Meaning, Structure and Understanding”, *Synthese* 48: 135-161.

Davies, Martin (1983), “Meaning and Structure”, *Philosophia* 13(1-2): 13-33.

Evans, Gareth (1981), “Semantic Theory and Tacit Knowledge” in Steven Holtzman and Christopher Leich (eds.) *Wittgenstein: To Follow a Rule* (London: Routledge and Keegan Paul). Reprinted in Evans (1985): 322-422.

Evans, Gareth (1985), *Collected Papers* (Oxford: Oxford University Press).

Fodor, Jerry (2008), *LOT 2: The Language of Thought Revisited* (Oxford: Oxford University Press).

Goodman, Nelson (1968), *Languages of Art: An Approach to a Theory of Symbols* (Indianapolis and New York: The Bobbs-Merrill Company, Inc.). Second Edition 1976.

Hopkins, Robert (1998), *Picture, Image, and Experience* (Cambridge: Cambridge University Press).

Kulvicki, John (2006), *On Images: Their Structure and Content* (Oxford: Oxford University Press).

Lewis, David (1994), “Reduction of Mind”, in Samuel Guttenplan (ed.) *A Companion to Philosophy of Mind* (Oxford: Blackwell Publishers). Reprinted in Lewis (1999).

Lewis, David (1999), *Papers in Metaphysics and Epistemology* (Cambridge: Cambridge University Press).

Lopes, Dominic (1996), *Understanding Pictures* (Oxford: Clarendon Press).

Newall, Michael (2011), *What is a Picture?* (London: Palgrave Macmillan).

Quine, W. V. (1970), "Methodological Reflections on Current Linguistic Theory", *Synthese* 21: 386-98.

Rescorla, Michael (2009), "Predication and Cartographic Representation", *Synthese* 169: 175-200.

Schier, Flint (1986), *Deeper Into Pictures* (Cambridge: Cambridge University Press).

Scruton, Roger (1983), *The Aesthetic Understanding* (Manchester: Carcanet Press).

Wollheim, Richard (1987), *Painting as an Art* (Princeton: Princeton University Press).