Modeling the Meanings of Pictures

Depiction and the philosophy of language

John Kulvicki

Table of Contents

3

Parts of pictures

Pictures have rich descriptive contents. These contents are determined to be what they are by pictorial characters, which can be described in terms of bare bones contents and rules for fleshing those contents out. This content talk makes it easy to forget that pictures are syntactically complex as well. Understanding pictorial syntax allows us to understand what the parts of pictures are, and how they conspire to make pictures syntactically and thus semantically whole. As we will see, in the right contexts different parts of a single picture can play different semantic roles, thus making pictures semantically much more complex and expressive than one might expect. This chapter's goal is to unpack and explain what the syntactic parts of pictures are, and illustrate how they can be deployed expressively.[1]

Pictures can be broken into parts by slashing them to pieces, dissolving them in acid, or burning them to ash and smoke. These violent partitions do not capture the notion of syntactic parts suggested here. That's significant because philosophers have claimed that cutting pictures to pieces is one way of generating their syntactic parts. Understanding the details, virtues, and limitations of those earlier proposals will be very important in what follows. The next two sections clear some ground, and set up the main claim, in section 4, about the nature of picture parts. Sections 5 – 7 then apply this idea to pictures and set the stage for the chapters that follow.

1. Syntax without grammar

Some features of pictures are relevant to their representing what they do, while others are not. We saw this in the previous chapter when we asked for the relevant sense in which all of the pictures in *No Chairs* were indistinguishable from one another. It was not

[1] The main theoretical tools deployed in this chapter are drawn from earlier work of mine, specifically Kulvicki (2006, 2007, 2010a, 2014 Ch8, 2015a). I'll be more specific about these debts as the chapter moves along.

important that the pictures were made of different materials, printed in different inks, or different distances from the sun. Those features of pictures are *incidental*, while the disposition of pale and dark patches on their surfaces are *syntactic*. They are relevant to the picture representing what it does. Change them, and you might very well change what the picture represents, but changing any of the other features just mentioned will not, unless the process also happens to change those surface features.

Different features are syntactically relevant to different kinds of picture. Black and white photos might be tinted in pink here and there, but hue does not affect what they represent and thus is strictly speaking incidental. Line drawings differ from those built around patches of pigment. That's not to say it is always easy to know which features matter. In many cases, consumers have to make an educated guess. Line drawings tend to stand out as such, just as photographs do, but we can make photographs of line drawings, and trace line drawings atop photos. A photo of a line drawing has a different set of syntactic features than a line drawing does, so before we can interpret a picture we need to know, or at least guess, what kind of picture it is. These decisions are akin to what Kaplan called "pre-semantic" judgments, in which one must decide what language is being spoken before being able to interpret it (Kaplan 1978, 296-297). In the pictorial case, these decisions are also pre-syntactic, since different kinds of picture make use of surface features differently in determining content.

We can say similar things about the syntactic parts of sentences. An inscription can be ripped up haphazardly, but only some ways of separating it into parts cleave at the joints. In one sense, the joints of sentences are the words .Those words are also assembled into important grammatical parts. Some phrases are in the subject position, some are verbs, some are objects of this or that sort. Some, like the words 'the' and 'a', do not have a meaning absent their use among other words. These grammatical positions matter for how we interpret sentences. Does John love Mary or Mary love John?

Pictures do not have grammar in the manner just sketched. There is no subject, verb, or object position within the picture. We talk about the subject of a picture, but this isn't a grammatical claim, so much as one about the important thing depicted. The lack of grammar doesn't mean there is nothing worth saying about pictorial syntax. To the extent that pictures break up into syntactic

parts, however, those parts will not be grammatically marked as playing different roles in constituting what the picture as a whole means. In brief, we should expect a good degree of uniformity among the syntactic parts of a picture.

To the extent that pictures' syntactic parts play different roles in constituting the meaning of a picture, it will not be because they come grammatically packaged as doing so. Instead, I'll suggest, context and conventional practices can make some parts salient as the things that play one role, rather than another. But that's to get ahead of ourselves, because we don't yet have a clear idea of what pictures' syntactic parts are.

2. Abstraction and content

The previous section distinguishes the syntactic features of pictures from their representationally incidental features like mass, distance from the sun, and the like. There is no distinctive pictorial grammar to speak of, so we should expect the syntactic parts of pictures to be uniform in that they can all contribute in the same range of ways to what a picture as a whole represents. So, how do those syntactic features, as it were, break up into parts that are relevant to what the whole represents? In order to understand this, we first need to consider how pictorial content relates to these syntactic features. This is something that makes pictures distinctive as kinds of representation. Specifically, there is an interesting connection between the syntactic features of pictures and their bare bones contents (Kulvicki 2006). By extension, there is an interesting, albeit more attenuated, relationship between pictures' syntactic features and their pictorial contents.

First, notice that different kinds of syntactic features seem responsible for representing different properties of scenes. Colors of the picture represent chromatic features of their subjects, for example. Black and white photos just represent patterns of light and dark, and they themselves tend to be monochrome. Spatial features of pictures represent spatial features of scenes. Temporal features of an audio playback represent temporal features of the scene, while its tonal features represent tonal features. And so on. This division of labor makes a lot of sense, and is deeply related to what makes pictures a distinctive kind of representation.[2] So far,

[2] In fact, I've argued (2006, Ch3) that pictures are instances of their bare bones contents. That is, pictures manifest all of the features included in

the idea is that we can isolate different kinds of syntactic features – colors, shapes, textures, etc. – as playing different semantic roles.

Keep in mind that these features play the semantic roles they do in a manner mediated by pictorial character. Character, as we saw in the previous chapter, amounts to a bare bones content plus a rule for fleshing that content out. The most straightforward connection between syntactic features of the picture and features that figure in its content will be found with respect to the bare bones content. So, the picture is dark green, but represents a bright green surface, in shadow. The pictorial content includes being bright green, and in shadow. The bare bones content, by contrast, includes only a more abstract feature, which is compatible with many combinations of surface color and illumination – bright green in shadow as well as dark green in bright light.[3] Similarly, a region of a picture is trapezoidal, even though its pictorial content includes a rectangular table. The picture manifests a trapezoidal shape, but also an abstract shape property that is shared by trapezoids, seen head-on, and rectangles, seen obliquely. The aspect of the picture responsible for it representing what it does is the more abstract one, which the table and picture can share. All of the details about how this works are open for argument and further articulation, but for now the point is just that we notice a fairly systematic division of labor among pictures' syntactically relevant features.

The second important point is that abstractions over syntactically relevant features map readily onto abstractions over what they represent. Let's say reds on a picture surface represent reddish chromatic features of scenes, and greens do a similar job for greenish features. Notice that if all you know is that a syntactic feature is red-or-green, then you know that the picture represents either a reddish or a greenish chromatic feature of the scene. Imagine also that a specific shade of vermillion represents a specific shade of red. If you know that vermillions are shades of red, then you also know that the vermillion picture represents a feature that is a shade of red. This point should not strike anyone as controversial. In fact, same can be said for linguistic

---

their bare bones contents. That's quite different from saying that pictures instantiate all of the features they represent pictorially, of course.

[3] I discuss such chromatically perspectival features, in relation to their spatial counterparts, in Kulvicki (forthcoming).

representations. If all you know is that some word is 'red' or 'green', then you know that it represents red or it represents green. And if you understand 'vermillion' then you know anything so designated is, if it's accurately characterized, something that's also a shade of red. This uncontroversial point plays an important role in the more controversial claim that follows.

When a picture surface is red, it is eo ipso red-or-green. Does the picture's being red-or-green represent a red-or-green feature of a scene? Similarly, any vermillion picture is red, so does any vermillion picture represent its scene as being red? The uncontroversial point made in the previous paragraph does not commit to an answer here. Just because all vermillion things are red doesn't mean that a representation of something as vermillion is also a representation of it as red. In fact, in the linguistic case, it very tempting to deny these claims. I can have the concept of vermillion, after all, without having a concept of red, so it's hard to see why we would build redness into the representational content of a claim about vermillion things. The third, and controversial, point is that in the pictorial case, things work exactly that way (Kulvicki 2007, 2010a).

A picture that represents something as vermillion does so, let's say, in virtue of being a specific shade of red (most likely vermillion).[4] But such a picture is a member of a class of pictures that can represent a large range of colors: blues, greens, yellows, and so on. If we consider the range of such pictures, we notice that many different shades of red are responsible for representing the range of red things. And because of this, the fact that a picture is red, rather than this or that shade of red, means that the picture represents something red. Similar remarks apply to spatial qualities of depicted scenes. A region of a picture is trapezoidal, and it depicts a rectangular table, from an oblique angle. That trapezoidal region of the picture is also quadrilateral, and the picture thus also depicts a quadrilateral shape.

Why think this is so? Perhaps pictorial content is quite specific, and anything less specific is not so much included within the content, but an inference away from it, which is typically the way we think about things in language? That is, the concept of vermillion might

[4] Remember that the main point here applies most straightforwardly to bare bones content, and only by extension to pictorial content. Because it's easier to talk about pictorial content, I'll do that here, and get more careful when it's important to do so.

not represent red, but anyone with a basic understanding of the concept should readily infer from the fact that something is vermillion that it is red. In pictures, things work differently, and a few points should suffice to show this.

First, our ordinary practice with pictures suggests that they represent features across levels of abstraction. We can use pictures as aids to finding the orange things just as readily as we use them to help us find things that are specific shades of orange. The reason we can do this so easily is that the picture helps (Kulvicki 2010a). Not all words for shades of orange wear their semantic links on their sleeves, but all the syntactically relevant features that represent shades of orange do. This point will return when we discuss some objections to this claim in the next section.

Second, more abstract color features, like being green rather than chartreuse, work in the same way to determine pictorial content as the more specific shades do. We can ask what work a specific shade like vermillion does semantically, but this question gets exactly the same kind of answer for the more abstract shades like red. Red pictures, per se, represent red things, and vermillion pictures represent vermillion things. So, red and vermillion pictures represent red and vermillion things. Similar remarks apply to spatial features that scenes are represented as having.

Third, it's implausible that we could interpret pictures at all absent such a mechanism being in place. Many pictures, though not all of them, have contents that depend in an indefinitely subtle manner on their syntactic features. And many more pictures have contents that depend on their syntactic features in a manner too subtle for ordinary consumers to appreciate, even if they are not indefinitely sensitive. Observers simply don't have access to such subtleties, even under conditions one might think of as ideal for viewing the picture. But if we can't determine, to such high degrees of precision, what the syntactic features of pictures are, how on earth do we manage to interpret them? Put differently, how could we interpret a representation without knowing which representation it is? If abstractions over syntactic features – I'm not sure which shade it is, but it's definitely red – map readily onto abstractions over the most determinate contents, then interpretation can be accurate, and blameless, even though it's not capturing all of the details (Kulvicki 2015a).

Remember that the points about abstraction and content apply, first and foremost, to bare bones contents. That is, pictures have

syntactic features responsible for their having the bare bones contents that they have. Some abstract chromatic features, spatial features, and so on, specify a bare bones content, which, coupled with a rule, yields a fleshed out pictorial content. Abstractions over the most determinate features responsible for this bare bones content map onto abstractions over bare bones contents, which, in turn, yield differing fleshed out, or pictorial, contents. This is important because if we don't restrict attention to bare bones content, the claim about abstractions seems implausible. Pictorial content can include a feature, for example, like being a turkey. But it's implausible that any picture of a red thing also represents that thing as being red, or a turkey. It is plausible, by contrast, that the picture represents the thing as being red or warm-colored, as well as representing it as being red, warm-colored, vermillion, and so on.

This feature of pictorial content is not limited to pictures. Diagrams, graphs, maps, and all sorts of non-pictorial images work this way as well, and I have suggested that this is the true mark of *analog* representation (Kulvicki 2015a). Most of the non-pictorial examples, we learned in the previous chapter, have constant characters. We do not, therefore, need to worry in those cases about whether the relevant abstractions are drawn over the features of bare bones or fleshed out contents. In many ways, those representations are easier to understand than pictures. Their non-constant characters are what distinguish pictures from all other non-pictorial, imagistic representations.

3. Two clarifying objections

Wait a second, anything that is red is eo ipso red-or-Martian. So, any picture that is red is also red-or-Martian. Does any picture of a red thing thus depict it as being red or Martian? That's madness, but it's not obvious how to avoid the conclusion, given what has been said so far.

The way to avoid this conclusion is to look to the most specific things a picture in some system can say as a way of limiting the most abstract things it can say. Imagine a given kind of picture can represent a range of specific color properties. Any abstract feature you can build up by disjoining those properties is also part of the pictorial content. So, red-or-green works, as does red in the vermillion case, given that the system of depiction can represent

the full range of reds. But red-or-Martian is not a possible content because the system of depiction cannot represent being Martian. The idea is that not any abstraction over a picture's syntactic features yields a ready abstraction over its content. The limitation here is built into the most specific syntactic features the picture might possess. Abstractions constructed from those are also represented, while those not constructible in such a fashion are not.

Okay, but if we're looking to the power of the system as a whole before determining what abstractions count as legit, it seems as though languages work in exactly the same manner. Consider all of the names for colors in English. They constitute the lexical resources devoted to representing colors, and we can imagine abstractions over those words, written or otherwise. 'Blue', 'green', and 'puce' are all representations of colors. Any word that is 'blue' is *eo ipso* 'blue'-or-'green'. Ditto for 'green'. So, being 'blue' or 'green', in English, represents being blue-or-green, and so on for any abstractions that can be built up from the most specific terms we have in the language. If languages work this way, then the feature seems not to distinguish pictures from other kinds of representation. But, more importantly, it doesn't seem as though just making the abstraction yields a representation of a new, abstract content in language, so perhaps it doesn't do that in the pictorial case, either.

Yes, we can think of languages in this manner. But no, doing so is not helpful for understanding how we interpret them. By contrast, this fact is central to how we understand pictures, what makes them a distinctively perceptual kind of representation, and what makes them useful in ways linguistic representations are not. Trivially, if you have a range of representations, you can generate abstractions over their contents by considering abstractions over the representations themselves. That's just what we did in the previous objection. But learning a language amounts to learning meanings of terms. These terms are the meaningful units of more complex expressions, and abstractions over those lexical expressions typically play no significant role in linguistic interpretation.[5] Language provides another tool for generating abstractions over the contents of representation: disjunction. It's

[5] Fully unpacking this point will help us distinguish representations with propositional content from those that have non-propositional content. So, if this is compelling, or the opposite, head right to chapter 8 after this one.

easy to represent something as being red-or-green in language, in fact, we just did it. Explicit disjunction is the tool we use in language, while in the pictorial cases the key to interpretation is found in making abstractions over syntactic features of representations.

By contrast to the linguistic case, if we were unable to lean on abstractions over the most specific syntactically significant features of pictures, we could hardly interpret them in the first place, as the previous section suggests. But how do these abstractions help in the pictorial case, if they do not in the linguistic case? Pictures' syntactically relevant perceptible features are such that *perceptually salient abstractions over them are the ones that are important to interpreting them*. We can see many shades of blue, and see them all as shades of blue. We can see many kinds of quadrilateral and see them all as being quadrilateral. Perceptual access to things' properties extends across levels of abstraction (Kulvicki 2007, Green 2015). Pictures are readily interpretable because those perceptually salient abstractions map readily onto abstractions over content.

Some linguistic practices exploit this feature of pictures, and when seen in this light they seem like the most pictorial parts of language. Using dashes or lines to count, for example, is somewhere between a linguistic and a pictorial practice, and all of the points about abstraction made above apply to it. Arabic numerals are more complex than dashes, but they are built to allow ready abstractions over their features that map onto abstractions over their contents. Focus on the tens' place in a list of numerals, or the hundreds' place, gives you a sense for how values vary within that range. If all you care about is whether a value falls between 100 and 200, it's easy to find this out by ignoring the ten's and one's places.

4. The parts principle

The points about abstraction and content are essential to understanding what makes for a syntactic part of a picture. Earlier I mentioned that the notion of part I have in mind relates in interesting ways to some of the more violent things we can do with pictures, like cutting them up into pieces. In fact, it's been common to think about the syntactic parts of pictures just in those terms.

Elliot Sober (1976) suggested that one interesting thing about pictures is that parts of pictures are themselves pictures. He imagines cutting up a picture into smaller and smaller pieces, and notices that "successive snippings never destroy representationality." (1976, 124) I think this is one of the earliest, explicit suggestions to this effect, but the idea has been implicit in talk about pictures in the West going back at least to Alberti's *Della pittura* (1425). This idea was also important to participants in the imagery debates from the late '70s through the '90s. The topic here is not mental images, what they are, or whether they exist, but those who were focused on that topic needed to know what might make something in the mind or brain an image, and one thing that suggested itself is the way in which spatial parts of an image are images themselves.[6] Not only are spatial parts of pictures themselves pictures, they seem to be pictures of spatial parts of whatever the uncut representation depicts. Jerry Fodor (2007, 2008) proposed a "picture principle" that almost makes these points explicit: "If P is a picture of X, then parts of P are pictures of parts of X." (2008, 173) All we need to do is qualify "part" as spatial, and we get the point just made. It's clear from his text that this is what Fodor has in mind.

The parts referred to by these earlier accounts fit the bill for being syntactic parts of the picture because they isolate some syntactically relevant features of the picture and disregard others. Also, any way you break up a picture via this cutting method, you are able to reassemble the whole, which is another thing we want to capture with our notion of a syntactic part. Suitably recombined, they give us the whole, and in that sense the semantic identity of the whole depends on the semantic identities of the parts and their method of combination.[7] Elisabeth Camp goes so far as to say that "insofar as we can discern syntactic 'parts' of a picture at all, these are either just points in a two-dimensional

---

[6] See Kosslyn (1980, 33; 1994, 5), Tye (1991, 44), and also Kulvicki (2014, Ch 8), where the ideas presented here are developed in some detail.

[7] See Blumson (2014, Ch 6). In one sense, Blumson seems to suggest that the parts of pictures are spatial parts. In another sense, he suggests that parts might be non-spatial. He compares pictures to chess diagrams and maps, on which one can place markers here and there. The markers are parts, but they are not spatial parts, per se. See Chapter 7 for a discussion of maps. I flag this point here because Blumson might be more on board with this suggestion than he might seem at first glance. Blumson never comes out and says exactly what a picture part is, however. Abell (2005b §6) suggests that the parts of pictures are their syntactic features, which is more in line with what I have in mind here.

array, or else regions whose boundaries are given by salient boundaries in the scene represented…" (2007, 156).

This proposal captures the thought from section 1 that the syntactic parts of pictures, if any, will look quite uniform, in the sense that there are no grammatical markers to single out this of that part as playing this or that role in the whole. All of these spatial parts represent spatial parts of the scene, and the way in which they might combine to form a whole depends on nothing more than how you spatially combine them. This proposal also reveals an interesting feature of picture parts: there is no privileged way of forming them. We can cut up a picture in indefinitely many ways. All such divisions yield bona-fide syntactic parts. So, not only do pictures not come grammatically prepackaged, with set roles for the different parts, they do not come pre-divided. There are indefinitely many, equally good choices for how to break up a picture into parts.

My own account of picture parts preserves the advantages of the Sober proposal, while being a good bit more open about the range of possible parts. The earlier claims focus on breaking a picture in to pieces. That's a useful way to go because the pieces are the kinds of things that we understand as existing independently of one another, and they are the kinds of things we can, mechanically, assemble into larger pictures. But the thing that makes such pieces parts of pictures in the syntactic sense, I suggest, is that the pieces abstract over spatial detail. Instead of the whole picture, which represents features across a given range of space, the part we get from cutting out a bit just represents things happening in a subspace of the whole. And it does so because its syntactic features are a subset of all the picture's spatially relevant syntactic features. So, the core notion behind syntactic parts of pictures is abstraction. We focus on some syntactic features to the exclusion of others. Doing so yields an interpretable representation, and in that sense, a part of the whole. Let's review a couple of ways to generalize the Sober proposal, keeping in mind that abstraction is the key thing behind it.

For example, one can focus exclusively on high-spatial-frequency features of a picture or low-frequency features, depending on one's interests. One can ignore the hue of a color photograph to focus more on what one would ordinarily find in a black and white photo. One can ignore color details in favor of focusing on where the scene is depicted as warm, or cool. All of these acts in effect break the

picture into parts that matter interpretatively and those that do not. Cutting up a picture is just a vivid example of this, which also tracks the sense of a part as a separable piece of something. But when we are interested in syntactic parts, being a separable piece is not the main notion. You can't cut a picture into its high-frequency and low-frequency parts, even though in the sense just outlined those are syntactic parts of the whole. In the final chapter, I'll suggest that a variation on this idea can help us distinguish propositional from non-propositional representations and contents.

The Parts Principle, as I call it, just says that abstractions over the syntactic features of a representation, R, are representations of abstractions over R's content (Kulvicki 2015a). Cutting in Sober's sense is a special instance of this principle, because cutting up a picture amounts to focusing on the parts of it that represent only some spatial parts of the whole scene. The Parts Principle is superior to the previous claims because it focuses on abstractions over syntactic features of all sorts. They need not be spatial, so we need not understand the chief syntactic operation as cutting something to bits. Pictures have high-spatial-frequency parts and low-spatial-frequency parts, they have hue parts and brightness parts, but those parts completely overlap one another spatially.

The appeal of Sober's claim is that it identifies parts in a manner that's perhaps more familiar to us from other cases, as well as from the philosophy of language. Parts, for Sober, are the kinds of things that have lives of their own. I can cut off a part of a picture, and we are left with two distinct parts, which can be assembled with other parts in indefinitely many ways. This sounds a lot like the lexemes of a language. By contrast, the parts I have identified by appeal to abstraction are not necessarily separable from one another. You cannot create a picture that only has indeterminate shades of color, though you can create a picture for which indeterminate shades of color are the most determinate syntactic features. It's important to liberate the notion of a syntactic part from the more ordinary notion of a separable part, but the details will not be important until Chapter 8. As we will see there, all of the syntactic parts of linguistic representations are, in a technical sense I will specify, *separable* from one another, while none of a picture's syntactic parts is separable from any of its others.

Let's say you don't like the point about abstraction, and thus want to reject the parts principle in its most general form. For much of

what follows, this won't be a disaster. I've defended the parts principle and its usefulness elsewhere (2015a), but the more restricted claim from Sober (1976) and those who followed him, like Fodor (2007), will serve many of the purposes served by the parts principle. We won't be able to do everything I would like to do, as we will see, but many of the points made will stand. Because I'll try to convince you that there are practices in which it seems as though we want the more expansive notion of a pictorial part, perhaps the cases that show up in subsequent chapters will bring you closer to seeing things my way. In addition, Chapter 8 shows that we get a particularly appealing, if surprising, way to distinguish imagistic from linguistic representation if we adopt the parts principle in its general form.

5. Syntactic parts and semantic roles

So far, we have understood pictorial content to be purely descriptive. In effect, a picture is a complicated description, though not one marked as definite or indefinite. As such, pictures are not themselves expressive of propositions for the same reason your average (in)definite description is not. Properly embedded in a larger conversational act, however, perhaps one involving a linguistic utterance, the picture can be part of the expression of a proposition.

While that point seems plausible enough, one is likely left with the intuition that sometimes pictures themselves are treated as expressing propositions. But if pictures are not embedded within linguistic contexts, and, by themselves, are just rich descriptions, how could a picture be the kind of thing that expresses a proposition? The short answer: pictures have parts, and in the right contexts those parts can be deployed to play different semantic roles. In some cases, for example, one part of a picture plays the role of identifying a subject, while other parts predicate things of that subject. In that case, the picture is expressive of a proposition. Let's see how this is supposed to work.

Consider what happens when your friends want to show you a bunch of vacation photos. There they are at the Taj Mahal! There they are at India Gate! There they are in Connaught Place! Tedious, but revealing. The family is represented by each of these photos. In fact, a part of each of these photos is devoted to representing the family. The pictorial content specifies a set of features that the

members of the family, and presumably only the members of the family, satisfy. The rest of the picture represents other things, like patterns of features satisfied by the Taj. We can break up the pictures like this, and these different parts can play different semantic roles. When the context is one in which you are looking at vacation photos, the part representing family members plays a subject role, while the other parts of the picture play the role of predicates.

If we were to give voice to what such pictures say, it would be something like "the people of this sort were at a/the place like that." Notice that I've expressed this in terms of definite and indefinite descriptions, even though the pictures themselves has no way of marking such things. They have no grammar for identifying one part as a subject, another an object, or predicate. Pictures have no 'is' of predication. All they provide is an articulate, and complex descriptive content. Because the complex content is articulate, in the sense that there are many parts that contribute to the complex whole, different parts can play different roles. They do play different roles when the pictures are used within a highly conventionalized practice, such as the display of vacation photos. Find the family, and notice what the rest of the picture says about it. In one sense, the picture doesn't do this all by itself. But then again linguistic expressions are not, all by themselves, responsible for what is expressed linguistically. Some languages do not mark a difference between definite and indefinite descriptions, for example. And all languages are such that linguistic content can depart significantly from what is expressed in context.

It might be tempting to say that these contexts force pictures into something like a linguistic mold. But saying that presupposes that predication, description, and the like are uniquely linguistic, and I don't see why we need to say that they are. Part of the point of modeling meanings is explaining the range of conversational uses to which representations can be put. Understanding pictures as having articulate, rich descriptive contents does just that. Parts of these contents can serve different roles in conversation, which is just what we find with the vacation photos, and let's consider what happens when the conversational context changes.

The local tourism office collects photos of people posing at nearby monuments. In their hands, the parts of the picture that isolate subject and predicate are inverted. They are not concerned with

this or that family, in particular, but with the place, and those who visit it. Is it mostly families? Singles? Tourists, or locals? In such situations, we can understand strings of photos presented as saying of the place like this, that it has visitors like that.

With vacation photos, you might imagine the relevant subject parts can be understood on the Sober model. Cut them out, and we have pictures of family members, which then occupy the places in question. Not quite. After all, the family occupies the space, and cutting them out doesn't preserve those features of how they occupy it. Similarly, the tourist board photos can't quite cut out their subject parts. After all, the place in question is represented in different ways by every part of the picture including the ones occupied by representations of the families. Indeed, this thought should make it clear that just cutting out the people from the vacation photos won't suffice to identify them as parts, since it's unclear what remains of the predicative part. The family members, for example are represented as standing in full sun, but the illumination is a feature of the scene as a whole.

Product catalogs are keen to show you what, for example, clothing looks like when worn by different people in different contexts. Those photos take the clothing as a subject, and predicate much of it. Advertisements for paint show it *in situ*. Those pictures treat the paint, in fact usually its color, as the subject, and show what paint like that is like in situations like these. In these cases, too, it's hard to think about the relevant part of the picture in terms of something you could cut out of the whole. The part is the features responsible for representing the colors on the walls, or the features of the clothing. The other parts play a role predicating things of them.

This theme will reappear. Because the syntactic parts of pictures are uniform, they all have the same range of possible uses. So, if a picture as a whole can be used referentially (Chapter 4), so can a part of a picture, to the exclusion of the rest. Iconographic (Chapter 5) or metaphorical (Chapter 6) interpretation can apply to all, or just some, of a picture. But which parts should be used this way? That is up to contexts, in a one-off manner, or highly conventionalized practices. They can not only make it clear which parts are treated in which way, but which parts play the role of subjects and objects. They can, that is, allow pictures to go well beyond their simple descriptive contents and be expressive of propositions, commands, and so on.

6. Revisiting indirect pictorial reference

Toward the end of Chapter 2, I raised a worry about indirect pictorial reference, or pictorial reference via denotation. The worry was that pictorial contents are so rich and detailed that we might never be able to say that a picture, in fact, denotes any scene. There must be some way in which the content of any picture comes apart from the scene it is purported to represent. So, in those cases, which are most if not all of them, denotation is bound to fail.

The response to this worry was that in most cases we do not tie denotation to the most specific things a picture might say. A worry about that response was that is it far from clear how that might happen. No such thing seems to happen with richly detailed descriptions, for example, so why should we expect it to happen in the pictorial case? The discussion of parts hopefully clarifies how we might ignore the most specific things a picture says. The most specific things a picture says are, after all, but a part of what the whole represents. That's so because pictures also represent things at many less detailed levels of abstraction. Pictures can denote, even if the most determinate things they say fail to match the scene. This is so because pictures can be embedded within a practice that happily disregards those most determinate details.

Sometimes, the kinds of pictures we encounter strongly suggest that we are to disregard the details. A blurry photo is not usually interpreted as representing a blurry scene. Understanding the photo despite its blurriness is, in part, being willing to ignore some of its most specific details in favor of focusing on its less determinate features. Such a picture depicts a person with a reddish shirt, but it doesn't go into detail regarding his or her physiognomy. That is much more plausible a reading of the blurry photo than one according to which it represents a blurry individual. Precisely the same kind of interpretive move is available when interpreting photographs in perfect focus. Just like the blurry photos, they can be read as going into lots of detail about a scene, but they needn't be. Chapter 7 will explore another way in which pictures might refer to a scene even though they fail to accurately describe it.

7. Summary

Pictures have syntactic parts. Those parts are found by abstracting over the most specific syntactic features a picture manifests. Those

parts are also representations, just like the syntactic parts of linguistic expressions are. They are different from linguistic representations, however, because these parts are not grammatically marked as contributing in this or that manner to the whole. All the parts of a picture are, in that sense, the same. This explains the appeal of earlier discussions that suggest parts of pictures are just themselves pictures, of parts of a scene. In effect, that rule holds true, as long as we accept a more abstract characterization of pictorial parts.

The fact that all picture parts are, grammatically, the same does not mean that all parts must play the same role in communicative acts involving those pictures. Instead, it means that all parts have the same range of possible contributions to the whole. Within suitably conventionalized practices, certain parts can play the role of singling out a subject for comment, while in others those same parts might constitute the comment on a subject singled out some other way. Because there are no grammatical cues, practices must be highly conventionalized for any such separation of roles to take place. Alternatively, as when this happens in a one-off fashion, a fairly elaborate setup of the context is required.

Once parts are broken up and treated differently, it becomes clear that pictures, properly embedded in communicative practices, can succeed in expressing propositions. One part identifies a subject, while the other constitutes a comment on it. Yes, these parts fail to do such work absent the right communicative context, but this is not to say that pictures only manage to express propositions when combined suitably with linguistic representations. Within the practice of showing vacation photos, for example, it's the pictures themselves that say the family is in such and such a situation. We can also linguistically express such contents – the previous sentence did just that – but that is not to say that the picture only expresses a proposition in virtue of language. In the cases described above, pictures are the only representational parts of the communicative acts.

So far, we have considered only a couple of ways in which parts might be interpreted differently. Specifically, because pictorial contents are descriptive, we can understand them in a definite or an indefinite manner. And different parts of the picture can be read in each of these ways. The following chapters consider different ways in which pictures, understood primarily as unspecified descriptions, might acquire contents. The operations

in question will all be defined over character and content, and all of the operations defined in what follows can apply differently to different parts of pictures, once a suitably rich practice is developed.

8
Distinguishing kinds by parts

Pictures are an important instance of a broader kind of representation, one that includes maps, graphs, and diagrams. In the previous chapter we saw, for example, that maps are in many ways similar to pictures, but in many ways they are not. Elsewhere I've offered an account of what makes this class of representations distinctive (2014), but the present work puts us in a position to improve on that. Specifically, the distinction between linguistic and imagistic, propositional and non-propositional representation comes down to how they have syntactic parts. Linguistic representations are such that all of their syntactic parts are separable. Imagistic or non-propositional representations are such that none of their syntactic parts are separable. The challenge is explaining what separability is, and then arguing, against what seems to be compelling evidence, that the rather extreme all-none claims hold for each kind of representation. Separability is a new notion, though, like any idea, it has ancestors.[1]

When language-minded philosophers have turned their attention to pictures and other non-linguistic representations, they looked for separable syntactic parts. That's what we find in language, after all, and separable parts are very important for understanding how languages can do what they do. In fact, it even seems as though pictures and maps have separable parts, and some very interesting philosophy came out of attempts to understand them in this way.

The view presented here suggests that the break between linguistic and imagistic representation is fairly profound. This makes it all the more interesting that the semantic models stolen from Kaplan and others work well in both the pictorial and linguistic cases. It suggests that we can understand representation, generally speaking, using a fairly unified syntactic and semantic vocabulary, despite the significant differences between its subclasses.

---

[1] As we will see, John Haugeland (1998) came closest to articulating the point I make below.

The chapter begins by clearing ground, identifying what we mean when we talk about syntax, what a separable part is, and how philosophers have tried to force this notion onto pictures. Worries about compositionality turn out to be involved here. After that, the notion of an inseparable part is introduced, and I argue that all syntactic parts of imagistic representations are inseparable. We will then be in a position to see how questions of compositionality and look different when asked of pictures, and why it makes good sense to say that pictures and their ken are non-propositional.

1. Two ways to think about syntax

It's somewhat rare for philosophers working on pictures to use the notion of syntax at all, but when they do they focus on which features of a representation can make a difference semantically.[2] For example, one might say that some feature is syntactically relevant to a representation because representations that differ with respect to it can differ in meaning. In some pictures, hue does not matter, while in others it does. When it does, you can change what a picture means by changing its hue here and there. When shapes matter at a given level of detail, changing them can also change meaning. Thought of this way, syntactically relevant features are a special subset of a representation's features overall. The mass of a photo is completely incidental to what it represents, as are the colors on its backside (Kulvicki 2014, Ch 4).

Thinking of syntactic features this way can help account for what makes pictures a distinctive kind of representation. Colors on a picture, for example, can stand for colors, shapes for shapes, and so on. The details get complicated, and they are worked out differently by different people, but the point is that there seems to be something special about how pictures pair their syntactic features with meanings.[3] By contrast, intuitions go, there is

---

[2] See, e.g., Goodman 1976, Bach 1970, Stern 1997, Perini 2005, and Kulvicki 2006.

[3] Many versions of this intuition fit the model suggested here. Even theories that do not explicitly mention syntax, like most resemblance theories (Hyman 1989, 2006, Hopkins 1995, 1998, Abell 2009, Blumson 2014) can be understood in these terms (see Kulvicki 2014). And this is also true of Goodman (1976) who is considered the main proponent of pictorial conventionalism. His view was that pictures' syntactic features were distinctive in being dense, such that the only way to order all pictures is that between each two there is a third. He thought that the

absolutely nothing interesting about the way in which the syntactic features, so understood, of linguistic representations are paired with meanings. We can ask about how changes in the shapes of letters and words affect their semantic identity, but we don't think that anything interesting or systematic will be revealed by such an investigation. So, talking about syntax in this manner can be quite revealing when it comes to pictures but is unhelpful as far as language is concerned.[4]

What makes languages interesting as representational systems is that words can be combined into highly structured expressions. There are few if any similar operations defined over pictures. One cannot mark off a subject from an object or a relative clause from a verb within a picture. If you look to grammar, then, you will find little that makes pictures interesting.

This story about syntax is misleading because it likens pictures to the lexemes of a language. There are interesting relations between the syntactic features of pictures and what they mean, while there are no interesting relations between syntactic features of linguistic lexemes and what they mean. There is no grammar within pictures, but a lot within language. This ignores the fact that pictures are highly structured syntactically. They have syntactic parts that play roles in determining the meanings of whole pictures. Better to compare pictures as complex, structured wholes to complex linguistic expressions. When we do that, we will see that pictures and maps differ from languages in a more interesting and profound manner than has been noticed. To do this, we will focus on compositionality, which is a claim about how the meanings of parts collaborate to establish the meanings of wholes. This will, in turn, force us to articulate two kinds of syntactic part: separable and inseparable.

---

semantics of pictures worked similarly. So, even for Goodman there was an interesting relationship between pictorial syntax and semantics that played an important role in establishing what makes pictures distinctive as a kind of representation.

[4] It will tell you something, just not much for present purposes. Languages are digital representational systems, and being so is a matter of how syntactically relevant features contribute to the identification of words (Kulvicki 2017b).

2. Compositionality

Barbara Partee unpacks what philosophers of language and mind mean when they say a representational system is compositional:

> The meaning of an expression is a function of the meanings of its parts and of the way in which they are syntactically combined. (Partee 2004, 153)

This definition does not require that the expression be a bit of language (see Szabo 2017 for a brief discussion of road signs), though languages are its main target. Key here is that some expressions are complex, in that they are built out of other expressions that can be meaningful. Some expressions are simple in that they are meaningful, but have no meaningful parts. Parts are meaningful on their own, and those meanings, plus rules based on the ways in which they are combined, yield the meanings of complex expressions. 'John loves Mary' and 'Mary loves John' have the same basic constituents but they are combined differently in each case, resulting in complex expressions with different meanings. We say that a representational system is compositional when the meanings of its complex expressions are determined by the meanings of their parts and the ways in which they are combined syntactically. Compositionality is thus not a purely semantic condition.

Parts of an expression are those things over which syntactic operations are defined. They are identified as the kinds of things that can occur in many different expressions, and often the kinds of things that can be combined in different ways with other expressions. 'John' plays different roles in different sentences, for example. Camp (2018, 25)suggests that the operations defined over propositional representations like languages are digital and universal in that they have discrete elements that can come together in indefinitely many ways.[5]

It's natural to formulate questions about whether pictures are compositional by using language as a model. As a result, one's understanding of the syntactic parts of pictures will be shaped by how things work in languages as well. That, I claim, is a mistake, and clearing it up will help us understand the distinction between

[5] She also suggests that they are asymmetrical and recursive. We will have occasion to talk about those features a bit later.

linguistic and imagistic representation. It will also show that the question of whether pictures are compositional is in some ways more and in some ways less interesting than has been noticed.

3. Separable syntactic parts

Remember that syntactic parts are the kinds of things over which syntactic operations are defined. In language all the operations are such that they work with separable parts. A syntactic part is *separable* just in case it can occur in complex expressions independently of whether any other specific syntactic parts are present. That's not to say that any syntactic part can, by itself, constitute a complete expression. The only claim is that a given syntactic part does not demand the presence of any other syntactic part within expressions it constitutes. The syntactic parts are meaningful, and they are combined via a kind of structured assembly. Any given part can occur in indefinitely many structured assemblies, and being a part is precisely being one of those meaningful things that can be so assembled. When a putative syntactic part seems to demand the presence of another, as in The Bronx, or hot dog, we tend to regard the ensemble as an indivisible whole.

A good test for separability is whether a putative syntactic part keeps its syntactic identity across the complex expressions that it seems to constitute.[6] When we talk of parts, syntactic or otherwise, we think of things that can be independently handled, and then assembled into a whole: building blocks, Legos, puzzle pieces. The building block keeps its shape whether it is part of a wall, a car, a letter, or whatever else you make out of it.

Export this linguistic thinking to the pictorial case, and you get the Sober (1976) line of thought about picture parts, which was discussed in some detail in Chapter 3. What could a part of a picture be, after all, but the kind of thing that can figure in many other pictures? Cut up a bunch of pictures, and you can mix and match the pieces to form new ones. The pieces are just the same as they were when they were parts of other pictures. So, pictures'

[6] Notice that I am not saying that the part keeps its *semantic* identity across the contexts in which it appears. That is a matter of some controversy. See, e.g., Fodor (1998a,b), Robbins (2005), and Johnson (2006).

syntactic parts are, like those in language, separable parts. We also see this line of thinking at work in recent discussions of maps (Camp 2007, 2018; Rescorla 2009; Blumson 2012). There, marks that indicate cities and towns, mountains, and the like, are treated like separable elements that can be placed anywhere you like.

The Sober line of thought is quite powerful, in that it suggests at least three convincing points of contrast between pictures and language. First, being a picture is, to borrow a term from Nelson Goodman (1951, 38), a dissective property, in that their parts are also representations, of the same sort as the whole is. In a fairly idealized sense, the parts of pictures are *pictures* in their own right. Linguistic representations do not work this way. A part of a sentence need not be a sentence, parts of noun phrases need not be noun phrases, and so on.

Second, there is only one rule for combining parts: connection. Make the border of one part congruent with the border of another and the result is another picture. The new picture is, in a way, a conjunction of the earlier two (Sober 1976). There is no rule of overlay. Though the previous chapter suggests that something like overlay is important to understanding maps, it is not overlay of spatially disjoint parts of maps. As with pictures, you can break a map into parts, by cutting it up, reassemble the pieces and wind up with a new map.

Third, any spatial part of a picture is a part. There is no privileged way of breaking up a picture into its constituents. Those in the Sober line typically have connected parts in mind, but most of their claims apply to non-connected parts as well.[7]

These three features play well together. Pictures have only one kind of part, so to make sense out of different ways of combining them one would need a privileged way of breaking up a given picture into parts. For example, if we want something akin to subject and object positions, we would need to be able to say what the parts are such that one is a subject and the other an object. Without a special way of breaking up pictures, it's hard to imagine how to implement different ways of combining the parts. Camp (2018, 25) suggests that, in addition to being digital and universal,

[7] A part is connected in the relevant sense when one can travel from any point within it to any other point within it along a path all of whose points are also within the relevant region.

propositional systems of representation are asymmetric and recursive. This means that, inter alia, the system comes with more than one type of part, that there are rules for which kinds of part can combine with which other kinds, and that such procedures can be applied recursively, to form more and more complicated representations. Maps and pictures, she claims, are structurally "flat" in that it just seems as though they say of places, what's going on there.

Because the Sober proposal makes pictures seem so different from language, one might think that it provides an account of the distinction between these representational families. There is thus good reason for this line of thought animating the imagery debate from the late 70s until the early 90s. What I want to challenge in the next section is the view that finding the syntactic parts of pictures amounts to finding separable parts. That move is motivated by an analogy with language that is strained at best. The easier task in what follows is convincing you that pictures have some inseparable syntactic parts. The harder task, discussed in sections 5-7, is convincing you that pictures only have inseparable syntactic parts.

4. Inseparable syntactic parts

Chapter 3 presented a Parts Principle according to which abstractions over a representation's syntactic features are representations of abstractions over the original representation's content. Ignore the hue of a picture surface, focusing instead on patterns of light and dark, and the result is also a representation, though one with a different content than the original. The content is different in that it is an abstraction over the first. It says nothing about hue, but still says quite a bit about patterns of light and dark in space. Similar syntactic operations – leave out the high-frequency spatial detail of the picture, focus exclusively on its hue, excluding saturation and brightness, etc. – also result in (syntactically and semantically diminished) representations.

Admittedly, it's a bit odd to say that an abstraction over the syntactically relevant features of a representation is itself a representation, distinct from the original. Abstraction is not a procedure that produces any new objects in the ordinary sense. Even though abstraction doesn't yield distinct objects, it can yield different representations because the result has a different set of

syntactically relevant features than the original. In the original, let's say, all aspects of color matter, while in the abstracted one hue is irrelevant. So, the latter representation is syntactically the same as others covered in different hues because hue is irrelevant. In any case, at best a subset of an objects' features contributes to its syntactic identity.

It is one thing to give a formula for making new representations out of a given one, and another to give an account of pictures' syntactic parts. After all, I could write a rule for transforming an inscription of the word 'green' into an inscription of the word 'red', but that does not suggest that 'red' is in any sense a syntactic part of 'green'. Why, then, is this principle a *parts* principle?

Abstracting over syntactic detail yields syntactic parts because the abstract features – red, instead of scarlet, quadrilateral instead of square, etc. – play their own roles in determining the content of the representation with which you started. Pictures and maps are such that they represent qualities *across* levels of abstraction. The picture does not just represent a specific shade of red, but instead represents the specific shade, as well as less specific shades, all at once. Simplifying a bit, red is responsible for the picture representing things as red. The more determinate shade, say scarlet, is responsible for representing things as scarlet. Pictures and maps have "vertically articulate" contents (Kulvicki 2007). As such, all of these more abstract features are syntactic, and thus deserve to be called syntactic parts. Understanding syntax helps us understand what we can do with representations. What we want from pictures, maps, graphs, and diagrams, is often to find information at some remove from the most determinate things they tell us. They help us find that information because of how abstractions over their syntactic features represent abstractions over their contents (Kulvicki 2010).

The Parts Principle suggests, therefore, that some syntactic parts of pictures and related kinds of representation are *inseparable*. They are parts, because they make their own contribution to the meaning of the whole expression, but they are not the kinds of parts that can occur independently of whether any other specific part is present. You can't peel the red off of scarlet.

5. The main claim

We now have a sense of what separable and inseparable syntactic parts are. In addition, I've offered what I take to be at least some convincing examples of inseparable parts in pictures. This puts us in a position to state a simple proposal for distinguishing imagistic from linguistic representation.

> Imagistic representations only have inseparable syntactic parts, while linguistic representations only have separable syntactic parts.

The important division between kinds of representation consists in whether they have separable or inseparable syntactic parts. Or, because parts are parts in virtue of the operations defined over them, these two representational families differ in the syntactic operations defined for them. In languages, the operations demand separable parts, while in images, maps, and so on, the operation is abstraction over determinate syntactic features. This operation demands inseparable parts.

So far, no argument has been given for the main claim. It is bound to strike readers as unnecessarily strong. Why not just say that pictures are distinctive because they have *some* inseparable syntactic parts? Might there not be another way of drawing a line between these representational kinds? The next two sections raise and reject a number of apparently strong objections to the main claim. With a defensible account in hand, we then return, at the end, to the question of whether pictures are compositional. This will help make sense out of why it makes sense to say that pictures and maps are not just non-linguistic, but non-propositional. Yes, maps are non-propositional, even though they express propositions!

6. Three objections to the main claim

Recall the exemplar of an inseparable syntactic part. Nothing can be scarlet without also being red. Both being scarlet and being red have their own semantic roles to play in determining a picture's content. So, they both count as syntactic parts, albeit inseparable

ones.[8] Now let's consider three cases in which syntactic parts of pictures and maps seem separable. One comes specifically from work on maps, the others apply to pictures and maps together.

First, perhaps the most obvious thought is that the Sober proposal about picture parts shows that some are separable. Cut up a picture, and the result is new pictures. Those spatial parts can be combined with other picture parts in any way one likes, so those parts are separable. This is as close to the building block analogy as one can get, and it applies to both pictures and maps. Though pictures and maps might have some inseparable parts, they also have separable ones.

Second, we can distinguish colors from the locations they occupy. Perhaps some color or other needs to be at each location in a picture, but the color's location plays a different role in determining content than the color itself does. The last chapter suggested, too, that locations can play the role of referring terms, while colors and the like play an attributive role. So this seems like another, fairly basic, case in which we have separable syntactic parts.

Third, Casati and Varzi (1999) offer a map semantics in which different colors are given distinct semantic values and can thus be evaluated independently of one another. You can layer as many colors as you like at any given location on a map, even though in practice we need to establish conventions for layering such features. On my view (Kulvicki 2015), maps have incompatibility classes of features. One of the important things about such classes is that they can be added or taken away from a map. Some maps tell you about cities, some do not. Both proposals yield fully compositional semantics for maps. So it seems as though colors, in Casati and Varzi's sense, or incompatibility classes of features, in mine, are separable syntactic parts.

All three of these cases suggest that the main claim is false. While pictures might have some inseparable parts, they also have separable ones. Separability might therefore be less theoretically interesting than advertised. In fact, the whole point of introducing

[8] We can certainly imagine kinds of representation in which being red is syntactically important, but being scarlet is not. And we can imagine cases in which scarlet matters but not red. In pictures, however, we find that that each one plays its own role in determining content. So, in the pictorial case, both are parts, albeit inseparable ones.

separability was to use it as a way to cut kinds of representation at an interesting joint. So, if we are to proceed, these objections need to be undermined, and the next section shows how to do that. These cases all collapse, I claim, to the exemplar of inseparability we find with scarlet and red.

7. Why the objections fail

Section 3.3 voiced a worry about abstractions over syntactic features and whether they could yield interpretable representations. Say that a picture, at a given location, is scarlet. That part of the picture is red, but it is also red-or-a-wombat. This might suggest that the picture *represents* the corresponding part of the scene as being red-or-a-wombat, which is deeply implausible. We can avoid this conclusion by insisting that abstractions are constructed out of the most determinate syntactic features of a picture. The picture can say that a scene has this or that color, to a high degree of precision, because the picture can be this or that color. It can also say a lot about abstractions over those determinate chromatic features. But it cannot say anything about wombats because there is no feature of the picture that represents wombats. The picture can say that there is a region of a certain specific shape present, but not that a region of a certain shape *or* a magnetic potential, is present.[9] We build abstractions by, as it were, disjoining the most specific things that a picture can say, which correspond to the most determinate syntactic features a picture can have.

This point about abstraction is central to understanding why there are no separable syntactic parts of pictures. Schematically, the responses to all three objections in the previous section go like this. The abstractions that matter are those built from the most determinate syntactic features of representations. But the most determinate syntactic features are typically much more specific than one might have thought. It's not being a specific shade of color, but being such a shade, at such and such a location. Likewise, on a map it's not a specific color, but a specific combination of specific colors, at a specific location, that is the

[9] A graph (or, map) of magnetic fields, which somehow also represents color properties of a scene, could be used in this way, but that's not what pictures do.

most determinate syntactic feature. Once we realize this, all of the putatively separable syntactic parts start to look inseparable, after all. We will be unable to understand the parts as keeping their syntactic identities if removed from the wholes they constitute.

The process whereby we found the exemplar of inseparable syntactic parts was simple. Find the most determinate chromatic features that can characterize the surface of a picture. Then ask whether less determinate features, understood as abstractions over those, also carry semantic weight. When the answer is positive, it looks like we have syntactic parts on our hands. These parts are inseparable because anything that has the highly determinate property is also scarlet, and anything scarlet is also red.

Let's apply that reasoning to locations. Locations are syntactically relevant features of pictures and maps. Where some color shows up is quite important to what the picture or map represents. Move all the colors around and you get a different picture. So, what is the most determinate location property on a picture surface? A location is syntactically significant to the extent that it is distinguished from and spatially related to other locations *on the picture surface*. The picture says nothing about locations that have no points corresponding to them on its surface. This is not a point about the grain of the picture, either syntactically or semantically. Whatever the grain, the smallest syntactically significant spatial part identifies a location vis-à-vis all and only the other spatial parts of the picture.

When we cut a picture in two, each half says nothing, spatially, about the other one. Locations on the new part no longer have anything to do with the rest of the picture from which they were excised. So, semantically each new picture is different in that it speaks of a smaller range of locations, but it has different syntactic qualities too. Locations on the surface are now part of a smaller ensemble, so they are distinct from and related to a mere subset of the points that they used to be distinguished from. The most determinate spatial feature on the cut-off part is less specific than the most determinate features on the whole picture (cf. Heck 2007, 127-128). So, when we cut off a piece of a picture, we don't leave the syntactic properties of either part intact. We change them.

An interesting case that was popular on the internet in early 2018 helps illustrate this point about space. It is a vivid example of what

has been called the Leaning Tower Illusion (Kingdom 2007). Someone posted the following pair of images on imgur.com:

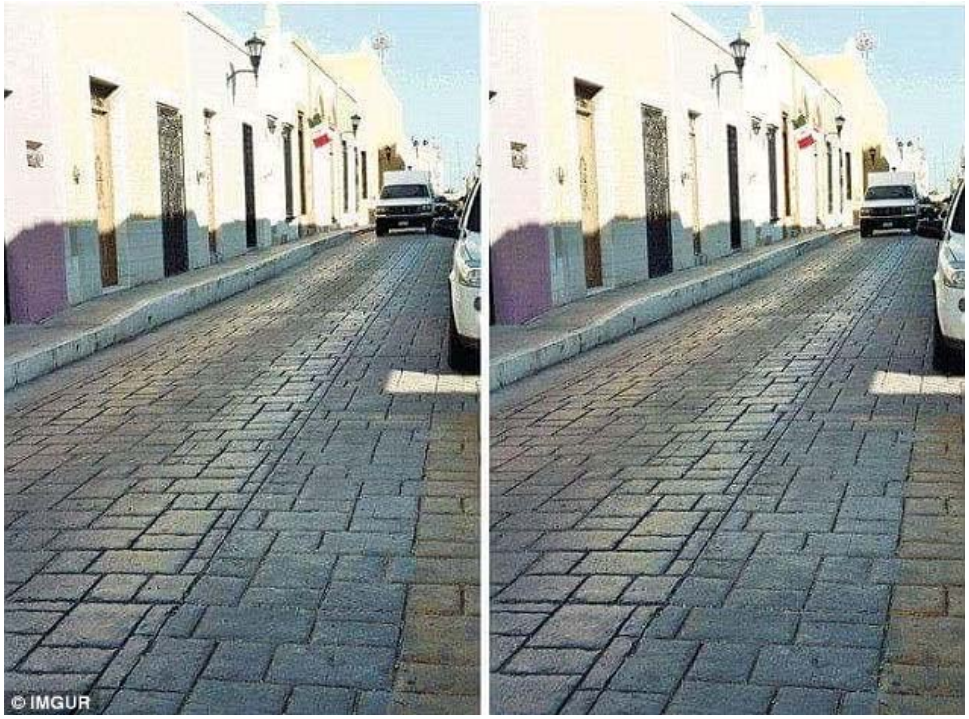

People are generally quite surprised to learn that the picture on the left is identical in its surface features to the one on the right. In fact, convincing yourself that they are the same might require moving one image so that it is below, or on top of the other one. What's the right way to describe this case? Because these pictures are close together, we are tempted to see them as one image (Kingdom 2007). But, as a whole, the spatial parts of the two images have quite different significance, syntactically and semantically, than those same locations do when they are in the separate pictures. With one picture, there is a road that forks. With two, neither picture has anything to say about the spatial parts of the other picture. When you put them close together, however, viewers are tempted to see two pictures as one, even though they also seem like two, because of the line separating them. It's hard to see the surface features of each picture independently of what they might represent. Though these cases are particularly vivid, much the same is true of any picture from which you excise a part. Do so, and you change the syntactic and semantic identity of the part you removed.

All of this should make sense, given the discussion that started this section. In interpreting pictures, we look for abstractions over the

most determinate things a picture can say. That's why scarlet pictures do not represent being red-or-a-wombat, even though they do represent being red. Similarly, a picture does not represent a color patch as being there, or on the rings of Saturn, or at any other location you might imagine. At its most specific, the picture says something about locations with respect to the other locations that the picture represents, no more.

The upshot is that spatial parts of pictures are not separable syntactic parts. They do not keep their syntactic identities when they are removed from the wholes that they partly constitute. As a result, they do not keep their semantic identities, either. And if you were to cut up a number of pictures and reassemble them in a hodge-podge fashion, the reassembly would change their syntactic and semantic features. We can certainly cut off parts of pictures and reassemble them into new pictures, just like we can cut up words into letters for reassembly. But to know whether those parts of pictures are separable requires more than that. Does the mechanical operation of cutting correspond to a syntactic operation that breaks the representation into parts that retain their syntactic identities? No. Excision changes the syntactic identity of the thing excised.

Any specific location on a picture, understood as above, is also an instance of innumerable, less specific locations, and all of those less specific locations have their own semantic work to do. The picture represents something as being exactly there, with respect to the other locations it can represent, but it also represents it as being sort of off to one side. And it does the latter by manifesting a less determinate spatial property. Those less determinate properties are syntactic, but inseparable. In that sense, locations work just like colors. Nothing can be scarlet without also being red. Likewise nothing can be this location, as opposed to any of these others, without also being innumerable other less determinate spatial features. Likewise, nothing can be red without being some more determinate shade of color, and nothing can be the abstraction over spatial features you get by focusing on a spatial part of this picture without being part of the ensemble of spatial features from which is it derived.

Cutting, then, is a physical operation on pictures that can result in new representations. And we can understand how to combine cut picture parts into new pictures. But having an operation that

shows how to transform one representation into others is not the same thing as showing what the syntactic parts of that representation are. We can change 'red' to 'green', remember, without revealing much of interest about either lexeme. Now let's consider the other cases, according to which the colors are separable, either from one another, as in maps, or from the locations at which they are placed.[10]

We generated the exemplar of inseparable parts by focusing on color. But insofar as we focused on the most determinate *chromatic* feature that can characterize a surface, we had already abstracted from a more determinate syntactic property the picture or map can have. The picture has the property of being this or that color *at location x*. Syntactically speaking, the picture is a collection of colors at locations. Every location gets a color, on pain of the picture having a hole in it. So, as far as the representation is concerned, anything with a color has a location, and all locations have colors. We can no more peel the color off of a color-location than we can peel the red off of scarlet. Colors are not separable from locations.

You can certainly think about color independently of location, and vice versa. But notice that you can do the same for red and scarlet. The point is not that we lack an articulate understanding of what the syntactic parts are, but that they are not separable, in that they are not parts that can occur independently of whether any other part is there. You need a place for colors, and you need colors for places. Talk of colors is convenient, because we have easy ways of thinking about them as determinate, indeterminate, and determinable. Adding locations to the mix leaves us with amagalmated properties that we don't often have cause to talk about. But syntactically speaking, these odd properties are the most determinate ones possessed by the picture. The scarlet-red case is griping because it's metaphysically impossible to have something scarlet that's not red. It's less gripping, perhaps, that it is syntactically impossible to have a location without a color, but in this case we are worried about syntax, so that's the kind of possibility that matters.

Now imagine that we have another feature like texture, and that this feature is independent of the other color properties at a given

[10] Camp (2018, §4.4) also notes that maps are *syntactically* holistic in the way they represent space.

location. In effect, let's say we have a map. Any texture can be combined with any color at any location. It might seem as though texture is a syntactic quality that's separable from color. So, even if we cannot separate colors from locations, we can perhaps separate located textures from located colors.

This move fails, for the same reason the simpler proposal about color and location does. If texture is another feature, even if it's independent of color in some metaphysical sense, the map commits to both color and texture values for each location. So, the most determinate syntactic features of the picture surface are amalgams of color, texture, and location. Every point has both, so it's combinations of these features that fill locations. You can't have some color-texture combination without both a color and a texture, and so in that sense color and texture are, within this system, inseparable. Schematically, this argument applies in the same way to cases that include many more than two distinct properties, so any map and certainly any picture, is such that the features that fill places are not separable syntactic parts.

The proposal is indifferent to whether one adopts the Casati-Varzi semantics for maps or the one I offer, though once one is in the weeds things might have to be done differently here or there. For Casati and Varzi, each location commits somehow with respect to every color that the map can place at a location. That's because lack of the color on their view has semantic significance. On my view, a map has a syntactic structure such that it places one feature from each incompatibility class at each location. In both cases, these complex, located, amalgamated features are the most determinate syntactic properties of the map, and while their parts can be thought of independently, they cannot be separated.[11]

The three objections to the main claim, though at first glance plausible, turn out to be misguided. Yes, we can think about the parts of pictures and maps independently of each other, but that's not enough to show that those parts are separable. [12]

---

[11] If Rescorla's (2009) criticisms of Casati and Varzi (1999) are correct, and they cannot model maps as employing predication, then my approach to maps will work in these cases, while Rescorla and Casati and Varzi need to think about things differently. This point is worth exploring, but it's a digression in the present context.

[12] John Haugeland suggested that we can distinguish what he called logical representations from iconic ones by appeal to the nature of their

An interesting result follows from these considerations. We have been talking about the most determinate syntactic features (plural) of pictures and maps. But the point about separability is just that at the most specific level all of these features coalesce into a highly specific whole. Color is inseparable from position, positions here are inseparable from those elsewhere, and so on. In reality, it's misleading to speak of a collection of most specific features for any given representation. It makes perfect sense to talk of a multitude of syntactic features, some more specific than others, but these feature form an amalgamated whole which is the sole most determinate syntactic feature of the picture, or map. In this sense, pictures are syntactically and semantically holistic (cf. Camp 2007, 2018, Kulvicki 2015), even though they have identifiable parts. They are holistic in that none of their parts are separable, and abstraction is the syntactic operation of interest.

Before moving on, let's return to a point raised in Chapter 3 as a worry about this approach. It should always be possible to define abstractions over representations within a system of representation, regardless of whether it is linguistic, maplike, or pictorial. Take all the color terms of English, for example, and then imagine abstracting over features of those words. Anything that happens to instantiate 'red' also instantiates 'red'-or-'blue'. So, an abstraction over these shape properties has predictable semantic consequences, viz. the abstraction represents red-or-blue. We can say the same for numerals, words for animals, and any mix of these we like. This is a worry because it seems as though exactly the same syntactic process identified as distinctive of pictures and maps works well in languages too.

---

contents. The elements of the contents of logical representations "can enter into atomic contents one-by-one, without depending on their concrete relations to one another, if any." (1998, 191) The elements of iconic contents "might be conceived as variations of values along certain dimensions with respect to locations in certain other dimensions." (1998, 192) This has much to recommend it, but here I need to repeat a complaint about this that I have made before (Kulvicki 2006, Ch 6): Haugeland should have considered the ways in which syntactic structure relates to semantic structure. By focusing exclusively on contents he missed much of importance. My discussion above focuses on syntax, and the way it affects semantics: structure and content. Cf. Heck (2007, 135), who says "We should instead regard the content of a mental state *S* as encoding not just *what S* represents, but also *how*, that is, as encoding *S*'s compositional structure."

One can imagine working with the lexemes of a language in this manner. In some cases, like color terms, it's relatively easy to imagine because there are only so many of them. In other cases, like animal names, it's hard because the list is huge and ever-expanding. The point is not whether we can imagine doing this, however. The point is that this kind of abstraction is not the syntactic operation at work in language, and, unsurprisingly, we rarely if ever work with linguistic representations in this way. It's easy, for example, to represent an abstraction over red and blue linguistically: red or blue. Alternatively, you can just invent a new lexeme like 'blured' (don't do this). In neither case is an abstraction over 'red' and 'blue' eo ipso a new lexeme. Yes, the expression 'blue' is 'blue'-or-'red', but no, being so is not a syntactically interesting feature of that lexeme. There is nothing within the linguistic system to suggest that such an abstraction gives you another representation.

8. Compositionality and inseparability

Compositionality is a claim about meanings of complexes and how they depend on syntactic combinations of their parts. Now that we have a clearer sense of what the parts of pictures and maps are, we can ask whether such representations are compositional. The question looks a bit different than it does for languages, precisely because of how pictures and maps have parts.

Remember the definition:

> The meaning of an expression is a function of the meanings of its parts and of the way in which they are syntactically combined. (Partee 2004, 153)

The one syntactic operation of interest in pictures and maps is, I suggest, abstraction over determinate syntactic details. No syntactic parts of pictures are separable. Once you fix the most specific features of pictures or maps, you eo ipso fix the rest. In fact, each picture or map has just one most specific syntactic feature anyway. Fix that, and the rest come along for the ride. So, though pictures have syntactic parts, there are not really any options concerning how to combine them. Every change is global. Questions of compositionality seem interesting and compelling only when there are parts that can be combined in many different

ways, but that's not what we find with pictures and maps. So, in a trivial sense, pictures are compositional. A picture cannot have the most determinate syntactic part it has without also having all the other parts. And pictures are meaningful, so of course they are compositional.

If pictures and maps are compositional, however, they certainly don't *seem* to be so trivially. And the fact is that we talk all the time of painters, photographers, cartographers, and the like *composing* pictures and maps. The process of doing so typically involves working with kinds of marks, colors, and so on, and assembling them into a whole. In fact, in some cases – chess diagrams (Blumson 2012), maps (Camp 2007, forthcoming), architectural diagrams (Haugeland 1998) – it feels very much like we have parts to assemble. Casati and Varzi break maps down into atomic stages, one for each color, which together yield a meaning for the whole. My view of maps is that we need to break them up into incompatibility classes of features, which, together give you a whole. Maps can have many different sets of incompatibility classes. Pictures are assembled by putting lines and color, in different combinations, on a surface.

Let's not mistake recipes for making and breaking representations for descriptions of separable syntactic parts. Yes, if you want to make a map of a certain sort, you need to place features of this or that sort at every point on its surface. If you want to make a line drawing, you have to decide, in effect, where the lines go, how they are weighted, and so on. But none of that speaks to whether the syntactic parts are separable. In effect, making a map or a picture is deciding what its most determinate syntactic feature is. Such determinate syntactic features are assembled from properties that, metaphysically speaking, can occur independently of one another, and for which we have independent concepts. But that's different from saying that they are separable.[13]

Let's also not mistake having an articulate understanding of a kind of representation for understanding its separable syntactic parts. We learn a lot about maps, for example, by trying to find their

[13] As mentioned in chapter 3, both Abell (2005b) and Blumson (2014, Ch 6) seem on the fence as to whether it is spatial parts or features like colors that count as the syntactic parts of pictures. They don't have the same notion of part in mind as presented here, but there is nothing objectionable about finding parts to be features, spatial or otherwise. It's just that in doing so one is abstracting over determinate syntactic detail.

incompatibility classes, and we can derive a compositional semantics from those classes. But the fact that we can offer an articulate description of incompatibility classes does not mean that they constitute separable syntactic parts of maps. There are pictures and maps of many sorts, and each kind has its own way of recruiting syntactic features to make a whole. So, while it might be uninteresting *that* pictures and maps are compositional, it is very important to try and give articulate compositional semantics for these different sorts of representations. Doing so shows us just how these different kinds of pictures and maps are different.

9. Why non-propositional?

Some have suggested that linguistic representations are *propositional* while pictures and maps are not. In one sense, that's just a name. I've offered a distinction between these kinds of representation, and you can call it what you like. In another sense, however, the name was meant to be significant, identifying a feature that sets each class of representation apart.

Unsurprisingly, some use the term because they don't think pictures or maps can express propositions, while linguistic representations can (Crane 2009), but others find the term apt even though they think maps or pictures can express propositions (Camp 2018). I am on Camp's side here. We've seen many ways in which pictures and maps can express propositions in the previous chapters. Nevertheless, there is an interesting sense in which it makes sense to call them non-propositional representations.

Giving articulate voice to the meaning of a declarative sentence comes quite close to unpacking its truth conditions. That is, if you understand the meaning of a declarative sentence you have a good sense for the conditions under which it would be true or false. Propositions have traditionally been thought of as the bearers of truth or falsity (see, e.g., Hanks 2009). Declarative sentences give voice to, or express, propositions.[14] So, what makes linguistic expressions propositional? Generally, each syntactic part and the

[14] Unsurprisingly, there is much controversy over the nature of propositions, and especially about how finely they should be individuated. Do two sentences, for example, with the same truth conditions express the same proposition? Hesperus is a planet. Phosphorous is a planet. Set these debates aside for present purposes.

way it's combined with others makes a distinctive contribution to the representation's truth conditions. That's why unpacking the meaning looks a lot like listing truth conditions. Raised on thinking about syntactic parts in language, this seems almost obvious. This is precisely what fails when we start talking about inseparable syntactic parts.

Pictures and maps have exactly one most determinate syntactic feature each. The parts are inseparable abstractions from it, and their role is representing abstractions over the most determinate content. As a result, the most determinate syntactic feature fixes truth conditions. The inseparable parts can only make a redundant contribution to determining them. If you've already fixed that a truthmaker for some representation must be scarlet, it's given that it must also be red. Abstractions over the most determinate content in no way add to, or change the truth conditions of the whole. The picture represents someplace as being that specific shade of red, scarlet, and red. But anything that shade of red is also scarlet, and red, so if we fix the first thing we have fixed the other two as well. Whenever you find, as it were, a proper syntactic part of a picture or map, you are finding something that, by itself, does not add anything to truth conditions already determined by some other part.

Many philosophers have pointed out that pictures tend to have much more specific contents, as a matter of course, than linguistic representations do.[15] Language, the thought goes, can choose the level of abstraction at which to represent things, while pictures must be quite specific. My point here is not just that pictures are specific. In fact, as Lopes (1996) points out, many pictures can be quite non-specific. The point is that pictures represent features at many levels of abstraction, all at once (Kulvicki 2007, 2010, 2015a). Most of what they represent articulately has little to do with fixing truth conditions. So, while it makes sense to say that pictures have truth conditions, and even that they express propositions when used in certain ways, there is much more to pictorial content than that. Of course, it's strange to say that there is *more* to pictorial content, because abstractions over determinate content say *less* than the whole does. But the point is that pictures are articulate

[15] For example, Chisholm (1942) and Dretske (1981) in the philosophy of perception, and Schier (1986) and Hopkins (1998) on pictures.

about the abstractions, so they say less, information-wise, in an articulate fashion, and in that sense they say more.

We can get a feel for how this difference works by considering linguistic examples that are articulate across levels of abstraction. Imagine describing something as scarlet, red, and reddish, all at once in a declarative sentence. "That's scarlet, red, and reddish!" What conversational contexts make saying that a reasonable thing to do? Perhaps the most plausible cases are those in which you are not trying to communicate anything about the color of an object, but about your understanding of a concept. "I know what scarlet is! It's a color that's red, reddish, and a bit yellowish as well." (See Kulvicki 2007, 366.) Otherwise, norms about giving the right amount of information kick in (Grice 1989, e.g.), and it would seem problematic to say so many things. If we need to know it is scarlet, say that, but if all we need to know is that it's red you could just say that instead. Why would you need to say both? Pictures and maps are not as selective in the information that they deliver. Yes, they say quite specific things, but they also say so many non-specific things as well. In that sense, I claim, it makes sense to say that they are non-propositional.

10. Summary

Representational kinds divide according to how they have parts. Despite appearances to the contrary, pictures, maps, and the like have no inseparable syntactic parts while linguistic representations only have separable parts. Parts come along with operations defined over them, so another way of putting this is that abstraction over syntactic detail is the only syntactic operation of relevance to pictures while grammatical (de)composition is the one relevant to languages.

We started by asking whether tools developed in the philosophy of language could profitably be applied to the study of pictures. Chapters 2, 4, 5, 6, and 7 constitute my defense of the claim that they can. Pictures have non-constant characters, which explains intuitions about the ways we can interpret them. They have attributive contents that are not marked as either definite or indefinite. In conversation, they can be used with either definite or indefinite force. Sometimes, particular individuals are constitutive of the contents of pictures, in at least three ways: via

dthat operations, iconographic interpretation, or as the contents of locations, which can serve as directly referring terms. Pictorial metaphor can even be modeled well within this frame.

Shortly after introducing the character/content distinction, however, it became clear that we needed to understand what the parts of pictures are. In conversation, different parts of pictures can play different roles, which makes pictures particularly useful kinds of representation. Sometimes, one part is responsible for singling out a subject, sometimes it's attributive, sometimes a part delivers a particular individual, while the other parts constitute a comment on it. In maps and photographs, in particular, locations, which are parts on the present view, play the role of delivering particular individuals while the other features tell us what's going on there. Chapter 3 made it clear, however, that picture parts were a bit strange. We found them by a process of ignoring syntactic detail. This had the consequence that the contents of pictures and maps were distinctive in being articulate across levels of abstraction. This last chapter uses that point to suggest a deep division between kinds of representation.

The parts suggest that pictures and maps are fundamentally different from those of linguistic representations. But both are profitably discussed using the semantic tools sourced in language. Semantic tools are thus not local to language. They apply as long as we have some kind of articulate syntax with meanings attached.

Though I have a theory of what makes pictorial representation distinctive, this book doesn't fully presuppose it. You might think that I get many things wrong in that theory and still think that the model introduced here for discussing pictures' meanings works quite well. There are probably other ways to model the phenomena discussed in this book. In fact, I hope there are. I've not attempted to build many models and choose the best one because that task goes well beyond what a single book should do. But it's hard to be confident that you're getting it right absent other sustained attempts to do things differently.