Beyond Resemblance

Gabriel Greenberg
University of California, Los Angeles

0. Introduction

The scientific study of representation in the twentieth century was driven at almost every turn by the formal analysis of language. From logic, to computation, to cognitive science, our understanding of language served as a model for our understanding of representation generally. Yet there are other, nonlinguistic forms of representation as well, notably representation by *pictorial images*. In the realm of public communication, the use of imagistic representation is ancient in human societies and thrives without special training or tools in the form of iconic gesture. In modern industrial society, pictorial representations are ubiquitous, used to efficiently encode and transmit vast quantities of information—exemplified by maps, road signs, text book illustrations, architectural drawings, television broadcasts, and so on. In the private domain of cognition, the spatial organization of the human visual cortex strongly suggests that picture-like representation is one of the basic strategies for information manage-

This essay benefited from the generous conversation and comments of many friends, colleagues, students, and teachers: Tendayi Achiume, Luvell Anderson, John Carriero, Sam Cumming, Doug DeCarlo, Jerry Fodor, Edinah Gnang, Jeff King, Peter Kivy, Ernie Lepore, Dominic Lopes, Alex Morgan, Kevin Sanik, Jason Stanley, Matthew Stone, and an anonymous reviewer; participants in the NSF-IGERT Perceptual Science Program at Rutgers University, the Rutgers Language Group, Peter Kivy's graduate seminar in aesthetics at Rutgers in 2009, and my undergraduate class on the philosophy of art at Rutgers in 2010; audiences at the University of Witwatersrand, the American Society for Aesthetics Eastern Division, Rutgers University, and University of California, Los Angeles. This research was supported in part by NSF IGERT DGE 0549115, under the guidance of Matthew Stone.

Philosophical Review, Vol. 122, No. 2, 2013
DOI 10.1215/00318108-1963716

ment implemented by the brain, particularly in low-level perceptual processing.[1]

Like linguistic expressions, images can represent the world as being a certain way, and they can carry precise information about objects and situations potentially quite distant from themselves. But they do so in a way that differs radically from linguistic representations. Today there is near consensus about how sentences of public languages encode representational content, at least in broad outline: arbitrary conventions associate individual words with meanings; then compositional rules determine the meanings of sentences from word meanings, the way these words are combined, and the contexts in which they are employed. Complex formulae in a mental code are hypothesized to acquire their meaning in an essentially similar way, save that the arbitrary associations are determined by biology rather than social convention. But this linguistic model is inappropriate for pictorial representation for many reasons. The most fundamental is that successful pictorial representation does not seem to be arbitrary at all. The relationship between a drawing, photograph, or perceptual representation of a scene and the scene itself is one of intimate correspondence, nothing like the stipulative association between a word and its denotation.

Official White House photo by Pete Souza

What then is it for a pictorial representation to depict a scene? Orthodoxy holds that pictorial representation is grounded in *resemblance*. The observation that motivates this view is simple and incontrovertible. Suppose we were to return to the time and place at which the photograph at right was taken and view the scene from the original position of the camera lens. Obama would appear to us much as the photograph itself appears. His apparent shape would resemble a particular region of the image. His apparent surface color would resemble the surface colors of the image in the same way. In short, the scene and the picture would seem to be linked by many dimensions of similarity.

1. See Pylyshyn 2006, 392–93, for discussion. It remains a controversial question whether there are also picture-like representations in higher cognition. See, for example, Block 1982, 1–16; and Pylyshyn 2006, chap. 7.

Understood as the basis of pictorial content, resemblance also has a computational rationale. To the extent that depiction is grounded in resemblance, an interpreter may extract information about a scene simply by coming to understand the structural properties of a picture that depicts that scene, for each has the same properties. Recovering information about a sign's content by directly measuring properties of the sign itself is surely the simplest and most efficient interpretive computation available. We should *expect* to find representational systems governed by resemblance in this way.

Such considerations are the basis for all RESEMBLANCE THEORIES of depiction.[2] As I will argue in the first part of this essay, resemblance theories are best understood as theories of accurate depiction: they define necessary and sufficient conditions for a picture to accurately depict a scene. Though these theories are diverse in content and range from simplistic to highly nuanced, all more or less agree that the correct analysis of depiction takes the following form. For any picture P and scene S:

P accurately depicts S if and only if P resembles S (and X),

where the optional constraint X guarantees that any foundational preconditions for depictive representation are met but is itself compatible with both accurate depiction and its failure. It is the resemblance condition that does the work in resemblance theories; as a shorthand, I shall say that according to analyses of this form, accurate depiction is "grounded in resemblance." Parallel proposals have emerged from cognitive science and the philosophy of mind. Taking pictures as their model of representation, such theories hold that mental representation is grounded in *isomorphism*, a kind of similarity with respect to abstract relational structure.

This essay evaluates resemblance theories of depiction specifically as applied to pictorial images, considered as public objects deployed in communication. This curatorial decision is partly motivated by the subject's intrinsic interest and partly by the hypothesis that the study of pic-

2. In the twentieth and twenty-first centuries, resemblance theorists have included Peirce (1906), Wittgenstein ([1921] 1997), Morris (1946), Gibson (1960, 1971), Manns (1971), Novitz (1975), Lewis (1986, sec. 3.3), Neander (1987), Peacocke (1987), Malinas (1991), Files (1996), Hopkins (1998), Newall (2003), Kulvicki (2006), Blumson (2007, 2009a), and Abell (2009). Resemblance has also found adherents in the adjacent domains of mental representation and scientific representation, typically under the banner of "isomorphism"; see Cummins (1996), French (2003), and Giere (2004, sec. 4).

tures as public signs is both a necessary means toward understanding picture-like mental representation and easier than studying mental representation directly. The formal analysis of public language has been an essential ingredient in the study of mental representation, in part because expressions of public language, unlike thoughts, can be straightforwardly inspected and evaluated. My hope is that by developing a parallel formalization of pictorial representation, we will be able to make new strides in understanding all forms of imagistic representation. Henceforth I shall use the terms PICTURE and IMAGE synonymously to refer to those objects at work in public communication; so too for the relations of DEPICTION and PICTORIAL REPRESENTATION, which, with some caveats outlined in the next section, I will also use synonymously.

In proportion to their historical currency, resemblance theories of depiction have been the target of numerous objections, aimed at everything from their conceptual foundations to details of formulation. Far from dampening the spirit of resemblance, however, the proposed objections have only shown the way for more nuanced and plausible articulations of the same basic idea. Today, resemblance theories remain popular, resilient to objections, and defended at each turn by ever more sophisticated adherents.[3] In the first two-thirds of this essay, I will show how resemblance theorists have successfully responded to all of the most important charges pressed against them thus far.

But in the last third of the essay, I present a new and general argument that, I claim, undermines the best version of the resemblance theory and all of its ancestors. The argument emerges from a body of knowledge more familiar to artists and geometers than philosophers, according to which accurate images are produced by following particular recipes for projecting three-dimensional scenes onto two-dimensional surfaces. Yet there are many kinds of projection, corresponding to the myriad systems of depiction. I concentrate here on linear perspective and curvilinear perspective. I argue that linear perspective can be characterized in terms of similarity, but curvilinear perspective is an intractable counterexample to resemblance theory. For this system, the *differences* between a picture and a scene are as much factors in determining the success of pictorial representation as the similarities; it cannot be analyzed in terms of similarity alone. The same problem arises for many other common systems of depiction. Unlike some previous

3. See, for example, the dissertation by Blumson (2007) and a recent article by Abell (2009).

opponents of resemblance, I do not claim that resemblance between a picture and its object is irrelevant to accurate depiction, only that differences are relevant as well.[4] I conclude that depiction in general is not grounded in resemblance, but is better understood in terms of the more polymorphous and precise notion of geometrical transformation.

The remainder of the essay is organized as follows: section 1 argues that resemblance theories are properly understood as theories of accurate depiction. Section 2 takes the reader through four successive versions of resemblance theory, beginning with the naïve and implausible and building up to a sophisticated and compelling treatment of linear perspective. Section 3 lays out the argument against all such resemblance theories of depiction, by examining in detail the case of curvilinear perspective. I argue that although linear perspective can be accounted for in terms of resemblance, curvilinear perspective cannot. Section 4 is a conclusion: here I discuss how the same examples that undermine resemblance lead naturally to the positive view of depiction as geometrical transformation.

1. Pictorial Representation

The subject of this essay is PICTORIAL REPRESENTATION. In this section, I specify what is meant by this term. To begin, not all representation by pictures is genuinely pictorial. In the seventeenth-century painting by Philippe de Champaigne on the right, the rendering of the tulips represents life, the painting of the hourglass represents the passing of time, and that of the skull, inevitable death (see Lubbock 2006). Yet none of these symbolic elements figure in the image's pictorial content,

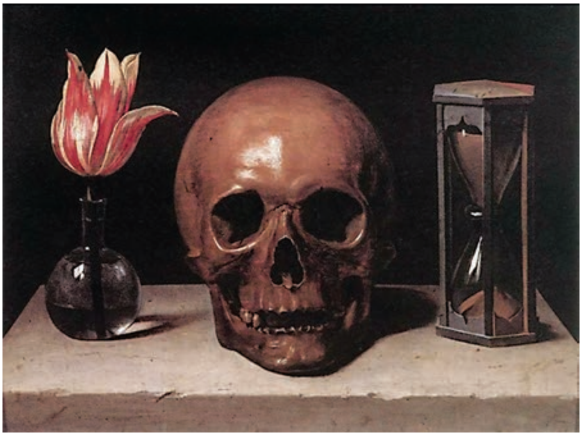

Philippe de Champaigne, “Vanitas,” ca. 1671

4. Goodman (1968, 5) notoriously rejects any role for resemblance in the analysis of depiction. Van Fraassen (2008) has also recognized that accurate depiction typically requires difference, though his reasons differ from those offered here.

in the sense intended here. Instead, the painting depicts a flower, a skull, and an hourglass, in a certain arrangement, at a certain time, and under certain lighting conditions, but nothing more.[5]

Note, next, that the topic of this essay is *pictorial* representation and not *artistic* representation. The interpretation of works of art is by nature unbounded, a consequence of the fact that artworks are characteristically tuned to expressive and metaphorical significance beyond literal content and often specifically intended to violate the interpretive norms that govern ordinary communication. This makes artistic meaning an admittedly unlikely candidate for the kind of systematic analysis attempted here. Just as linguists do not take poetry as their primary source of data, there is no reason that students of pictorial representation should take visual artworks as their point of departure. Instead, we should expect robust regularities to emerge only from the use of pictures deployed for practical information exchange, under conditions in which efficiency and fidelity are at a premium. Examples include road signs, maps, architectural and engineering drawings, textbook illustrations, police sketches, and so on. In this essay, I will confine my attention exclusively to pictures that are intended to answer to such standards; works of art are not my subject matter. (Still, I will use the term "artists" to refer generically to the creators of pictures.)

In the typical case of pictorial representation, a picture represents a particular scene that actually exists. But clearly there are pictures that represent purely fictional scenes; and there are pictures that appear to represent merely generic scenes (for instance, textbook illustrations). Both of these cases have been deemed problematic for resemblance theories.[6] Yet both fictional and generic representation are notoriously recalcitrant to analysis, quite apart from theories of resemblance. It seems unfair at this stage to attack resemblance theories for failing to account for phenomena that are poorly understood even in the more heavily traveled domain of language. We should give resemblance the benefit of the doubt: here I will consider only cases of pictorial representation holding between pictures and particular, actual scenes.[7]

5. The distinction between pictorial representation and other kinds of representation by pictures has been made by most authors on the subject; see especially Novitz (1975, sec. 2) and Peacocke (1987, 383). Goodman (1968, 5), for one, appears to reject the distinction altogether.

6. For discussion, see Neander 1987, 223, and Abell 2009, sec. 3.1.

7. To be clear, I will not exclude hypothetical pictures and hypothetical particular scenes, so long as they each exist in the same possible worlds.

Although there are many kinds of pictorial representation, in this essay I shall focus exclusively on the prototypical cases of photography and perspective drawing.[8] I will not consider photographs and perspective drawings in which symbolic elements such as textual labels or color coding is integrated into the image. I shall also ignore, however arbitrarily, such abstract expressions as Cartesian graphs, pie charts, and Venn diagrams—even setting aside much more pictorial styles of representation, such as the nonperspectival projections used in architecture, engineering, and cartography. These modes of depiction present obvious challenges to resemblance theories, but for this very reason I set them aside. Further, I will exclude alternative media such as film, relief, scale models, and audio recording. This narrow focus does not reflect my own views about the proper categorization of representational kinds; rather, it reflects my intention to rebut resemblance theories of pictorial representation in general, no matter how conservative their domain of analysis. As I shall argue, the problems with resemblance theory are deep and systematic and do not merely arise among controversial outliers.

Still, a theme of this essay is the diversity of representational styles even within photography and perspective drawing. There is, for example, color, black and white, and X-ray photography, as well as wide-angle and telephoto photography; among drawing styles, there are one-, two-, and three-point perspective drawings, as well as contour and cross-hatch drawings. All of these expressive techniques feature in human communication, from casual, day to day interactions to the high-stakes and high-precision contexts encountered by pilots, architects, and scientists. Even the most cautious theory of pictorial representation should be able to accommodate these phenomena.[9]

This study further presupposes that pictorial representation is at least a two-place relation that holds between pictures and the scenes these pictures represent. I remain neutral about the ontology of pictures—for example, whether they are concrete objects such as marked surfaces or abstract geometrical entities partially specified by marked surfaces. Pictures stand in the relation of pictorial representation to SCENES, conceived of as concrete, spatially, and perhaps temporally extended parts of physical reality. Informally, I will often talk of pictures representing *objects*, but only insofar as those objects are parts of scenes.

8. I intend the category of "drawing" to include paintings as well as line drawings.

9. I inherit this expectation from Lopes (1996, 32), who terms it the "diversity constraint."

Finally, following a widely observed distinction, we must prize apart two strands in the concept of pictorial representation: PICTORIAL REFERENCE and ACCURATE DEPICTION.[10] This conclusion is forced on us by the phenomena of pictorial *mis*representation.

Let's say that I attended the Obama staff meeting on November 20, 2012, where the photograph presented above was originally taken. The next day I decided to draw what I saw there at a particular moment during the day, from a particular vantage point. I produced the first image (above left). A week later I set out to draw the same scene, but this time poor memory and lack of common sense conspired against me, resulting in the second image (above right).[11] There is a certain sense in which both pictures depict the very same scene—a particular, real situation that occurred at a certain time and location. On the other hand, there is a clear sense in which they do *not* depict the same scene. The first picture accurately depicts Obama with short hair; the second inaccurately depicts him with longer, spiky hair.

10. The distinction is typically made with respect to pictorial reference and pictorial *content*, which amount to the same thing if content is defined in terms of accuracy-conditions. (I examine this connection at length in Greenberg 2012.) See, for example, discussions by Beardsley (1958, 270–72), Goodman (1968, 27–31), Walton (1990, sec. 3.2), and Lopes (1996, 151–52). Various ways of separating reference from content, including cases of misrepresentation, are discussed by Knight (1930, 75–76), Kaplan (1968, 198–99), and Lopes (1996, 94–98). Recent work on scientific representation has also marked the division clearly, for example Cohen and Callender (2006, sec. 2). Suarez (2003, 226) notes that his approach "presupposes a distinction between the conditions for *x* to be a representation of *y*, and the conditions for *x* to be an *accurate* or *true* representation of *y*. Both are important issues, but they must be addressed and resolved separately." Finally, Cummins (1996, 5–22) observes a parallel division with respect to mental representation, and Burge (2010, 30–46) elaborates the idea in detail for the central case of visual perception.

11. Both drawings are based on the photograph by Pete Souza, White House.

For a picture to bear the relation of PICTORIAL REFERENCE to a scene is for that picture to be *of* or *about* that scene, regardless of how imperfect a representation of the scene the picture is. (For short, we may call this relationship REFERENCE, and the referent of a picture its SUBJECT.) Both pictures here refer to the same scene. By contrast, for a picture to ACCURATELY DEPICT a scene is for it to represent that scene *correctly* and *without error*. Reference and accuracy are dissociable, for while accuracy is a norm of pictorial representation, this norm may go unrealized, as in the image at right above.

In the prototypical case of successful pictorial representation, the distinction between reference and accurate depiction is obscured because the two relations obtain simultaneously. Thus the short-hair picture is both a pictorial reference to Obama and an accurate depiction of Obama. We may illustrate this fact diagrammatically, using a thick arrow to indicate accurate depiction, and the thin arrow to indicate pictorial reference.

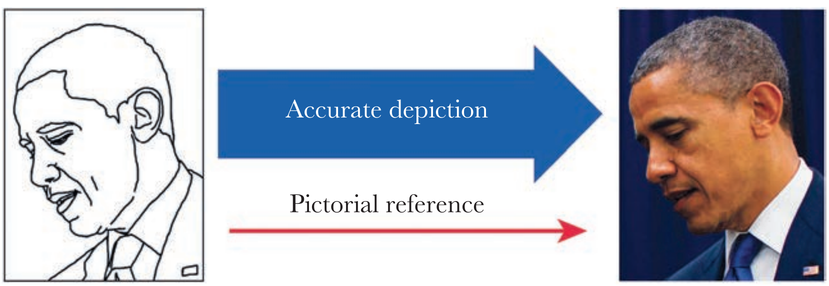

Picture Scene

Only in cases of pictorial misrepresentation do reference and accurate depiction come apart. Hence the spiky-hair picture is a pictorial reference to Obama, but it is not an accurate depiction of Obama.

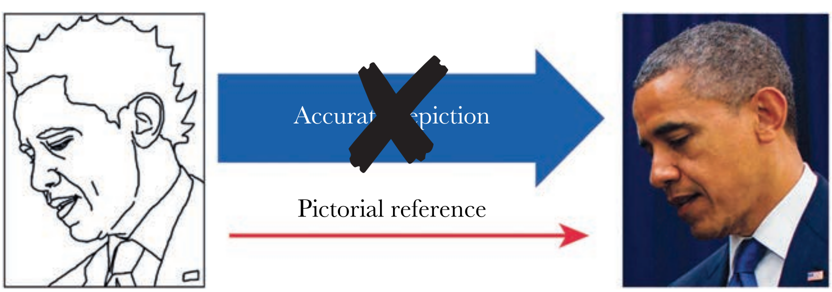

Picture Scene

The concept of accuracy at play here is the subject of direct and robust judgments, and it plays a central regulatory role in information exchange. For example, accuracy is the standard of pictorial fidelity that governs high-stakes communicative acts with pictures, like engineering and medical drawing. In such cases, where the information contained in the drawing may form the basis for decisions with life and death consequences, accuracy is the standard that artists characteristically strive for and viewers characteristically expect them to achieve. When the stakes are lowered, it is permissible for artists to infringe on this standard proportionately. Accuracy in this sense is not the exclusive denotation of the word "accuracy" in colloquial English, but it is one of them.[12] In addition, although accuracy comes in degrees, for practical purposes I will focus on the binary property of *perfect* accuracy. Here and throughout when I speak of a picture simply being "accurate," I mean that the picture is perfectly accurate, or very nearly so.

It may help to clarify what accuracy is not. By accuracy, I do not mean *realism*. Realism is a style of painting that approaches illusion: under the right conditions, an observer cannot discern a realist rendering of an object from the object itself. Yet the selective line drawings used by architects, engineers, and medical artists are often perfectly accurate, though hardly illusory.[13] Nor does accuracy imply *precision*. The first drawing of Obama is perfectly accurate in the black and white system despite the fact that it is wholly indeterminate with respect to color. It is also accurate even though the lines that compose it are wobbly; this does not mean that the shape of Obama's face is correspondingly wobbly, only that the standards of accuracy determined by this system of depiction are insensitive to a certain level of detail. Nor does accuracy imply *closeness to reality*, in any straightforward sense; a full-scale animated model of Obama is arguably more similar to the man than a black and white line drawing, but both may be perfectly accurate representations. Finally, accuracy does not entail *actuality*. A picture may accurately depict a merely possible scene just as well as an actual one, as illustrated by the use of architectural drawings in the evaluation of merely proposed building plans.

12. More or less the same concept of accuracy is widely invoked in discussions of pictorial realism (see, for instance, Lopes 1995; Abell 2007; and Hyman 2006, 194–97). Yet such accounts typically aim to analyze the richer concept of *realism*, taking some form of (nonfactive) accuracy as a merely necessary but insufficient condition. Accuracy also corresponds to what Walton (1990, sec. 3.2) calls the relation of "matching."

13. At any rate, this is *one* sense of the term "realism" among many. See Jakobson [1921] 1987; Lopes 1995; and Hyman 2006, chap. 9, for discussion.

If reference and accuracy can come apart, then what determines each? The examples above suggest that pictorial reference is determined by the etiology of the picture, unfettered by the degree of "fit" between the picture and its subject. What makes the spiky-hair picture a picture of *Obama* despite its significant misrepresentation, rather than of someone else more appropriate, or of no one at all, seems to be the referential intentions of the artist *that it be of Obama*, along with a sufficient causal connection between artist and subject to warrant such intentions. By contrast, once pictorial reference has been established, the accuracy of the picture is independent of its history, determined instead by the degree of "fit" between picture and subject. Thus even if I had produced the spiky-hair picture with the intention of making an accurate drawing, my intention would have been thwarted. The formal properties of my drawing, on the one hand, and those of Obama, on the other, inflexibly determine that the drawing is inaccurate.[14]

The dependence of pictorial reference on the history of a picture's creation is illustrated even more starkly when honest creative intentions are combined with impaired skills or inhospitable drawing conditions. For example, in an amusing parlor game, participants first look at a scene and then draw it while blindfolded. At right is my own, admittedly pathetic, attempt to render Obama using the same technique.[15]

Anyone familiar with the circumstances of the drawing's creation would agree that this is a drawing of Obama—that is, it refers to Obama. At the same time, the picture is a grossly inaccurate representation of Obama. Since the features of this blindfold-drawing that correspond to Obama's actual features are negligible at best, the only plausible explanation for why this picture is a pictorial reference to Obama, instead of to something else or

14. On the intention-relativity of pictorial reference, see Cummins 1996, chap. 2; Cohen and Callender 2006, 74; and Van Fraassen 2008, 23. If Goodman's concept of denotation for pictures is understood as pictorial reference, then he can be read as making much the same point (Goodman 1968, 5). It must be admitted that, in the case of photography, it is less apparent that pictorial reference and accurate depiction can come apart. I suspect that, in fact, they can; but a defense of this claim is beyond the scope of this note. See Costello and Phillips 2009, 16, for related discussion.

15. This type of example was suggested to me by Jeff King (pers. comm.). Walton (1973, 315n23) argues for the same conclusion from other considerations.

nothing at all, must go largely via the intentions of the artist and the context in which the drawing was produced. Meanwhile, once we know that the image *does* depictively refer to Obama, we need only a little knowledge of Obama's actual appearance to infer the image's obvious inaccuracy.

The central subject of this essay is accurate depiction; a theory of depiction, as I understand it, is a theory of accuracy conditions.[16] Further, I take it that accurate depiction, and not reference, is what has been at stake in the debate over resemblance. This is because resemblance is at best a dubious analysis of pictorial reference but an interesting and plausible analysis of accurate depiction.[17] For consider again the first pair of drawings above: the short-hair drawing at once accurately depicts Obama, depictively refers to him, and significantly resembles him.

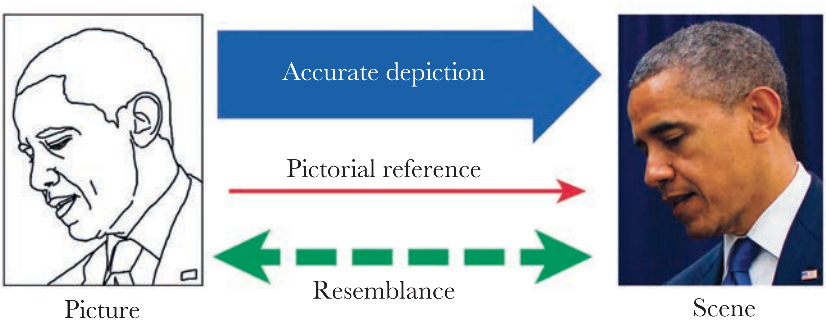

By contrast, the spiky-hair drawing, despite depictively referring to Obama, is not an accurate depiction, and also resembles its subject much less than its counterpart.

16. For those authors who do not explicitly distinguish pictorial reference from accuracy, there is an exegetical question about what they mean by "depiction" and "pictorial representation." My policy is to assume that other authors mean by "depiction" roughly what I mean by "accurate depiction," except where there is decisive evidence to the contrary.

17. The irrelevance of resemblance to pictorial reference but plausible relevance to accurate depiction is noted by Cohen and Callender (2006, sec. 4) and Van Fraassen (2008). Suarez (2003) argues against resemblance theories of referential depiction, though on different grounds than those cited below. It is also possible to read Goodman's (1968) attack on resemblance as an attack on resemblance as a theory of pictorial reference. See Walton 1990, 122–24, for detailed exegetical discussion.

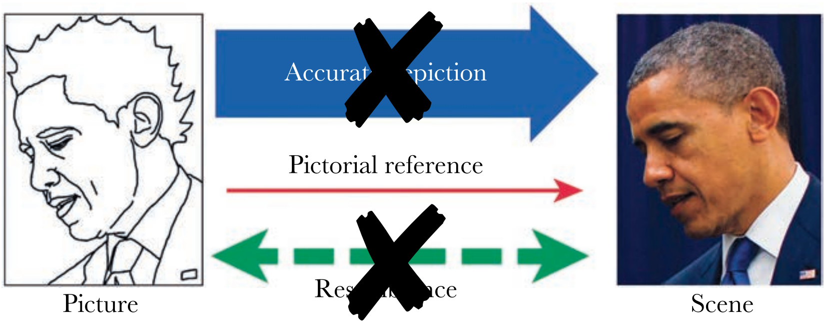

Thus, on one hand, depiction appears to covary closely with resemblance: it holds where resemblance holds and fails where resemblance fails. On the other hand, pictorial reference is relatively insensitive to resemblance, obtaining with or without significant resemblance.

It might be thought that, nevertheless, there is a minimum lower bound of resemblance necessary for pictorial reference—once the loss of resemblance becomes too great, reference cannot survive.[18] But such a view is belied by cases of extreme misrepresentation such as the blindfold-drawing above. For that is a misrepresentation of Obama, hence a pictorial reference to Obama, despite resembling him in almost no relevant respects.[19] In any case, all parties will agree that, if resemblance matters to pictorial representation at all, it matters *more* to accurate depiction than to pictorial reference. This much is enough to justify confining my discussion of resemblance theories to those targeted at accurate depiction. If resemblance does play a role in determining reference by pictures, it is not my concern here.

18. In the words of Hopkins (1998, 30), "Pictorial misrepresentation . . . has its limits." See also Abell 2009, 212.

19. The primary evidence cited in defense of minimal resemblance accounts of pictorial reference is that, out of context, a viewer looking at the blindfold-picture could not identify it as *of* Obama, even when familiar with Obama's actual appearance. But this evidence is compatible with, for example, a wholly intention-based approach to pictorial reference. For though the image may be of Obama, a viewer unfamiliar with the context of the image's creation could not guess as much. Indeed, what is right about the minimum lower-bound view is that, unless there is a minimum amount of resemblance, viewers will not be able to correctly guess the artist's referential intentions. Nevertheless, correct guessing on the part of uninformed viewers is not a necessary criteria for pictorial reference.

2. Resemblance Theories of Accurate Depiction

For a picture to accurately depict a scene is for the picture itself to resemble the scene. This is the intuition that guides all resemblance theories of depiction. The concept of resemblance may then be extrapolated in one of two ways. On the first approach, for two objects to resemble one another is for them to have the same "look," to be similar in appearance to the relevant audience. This way of understanding resemblance defines it as similarity with respect to properties that are viewer dependent.[20] The second approach analyzes resemblance in terms of viewer-independent properties. Such properties include intrinsic features such as size and shape, and relational properties like causal connection to an event or spatial relations to a point. In either camp there is space for a huge variety of approaches, corresponding to the myriad potentially relevant dimensions of similarity. But for the purposes of this essay, a resemblance theory of depiction is a theory that grounds accurate depiction in *any* kind of similarity, be it viewer dependent, viewer independent, or some admixture.[21]

In the remainder of this section, I develop a series of versions of the resemblance theory, at each stage raising what I take to be the most basic problem with that account and then revising the analysis to circumvent this problem. At the end we will be left with a version of the resem-

20. The useful distinction between characterizations of resemblance in terms of viewer-dependent and viewer-independent properties is due to Newall (2006, 588).

21. Resemblance theories also divide up according to whether the similarity in question is real or merely experienced. For example, even if we focus upon similarity with respect to shape, theories diverge over whether accurate depiction requires *actual* similarity with respect to shape, or just the *experience*—on the part of a prototypical viewer—of similarity with respect to shape. One reason for introducing this subjective element has been as a way of overcoming challenges associated with the pictorial representations of fictional and generic entities. (This is the strategy adopted by Peacocke 1987 and Hopkins 1998; see Abell 2009, 188, for discussion.) But since in this essay I concentrate exclusively on the depiction of extant particulars, there is no reason to involve ourselves with these complexities. For our purposes, resemblance is always real resemblance.

A second reason for this subjective turn is its resonance with the intuitive motivation for resemblance theory: that a picture resembles an object *to an agent* from a certain viewpoint. But I take it that even on experienced similarity accounts, the observer in question must be significantly idealized, in such a way that both the observer's perception and judgment is guaranteed to be accurate. Ultimately, I develop an enriched notion of imagery that captures this idealized experience by relativizing resemblance to abstractly defined viewpoints for both scene and picture. This kind of resemblance can then be defined as real similarity. So, for my purposes, nothing is gained by focusing on experienced similarity as opposed to real similarity. Thanks to Matthew Stone for this point.

blance theory that is at once flexible and compelling. Since my overall goal is to critique resemblance theories *in general*, at each step I will incorporate only those specifics that are absolutely required to avoid counterexample. As I will emphasize in the subsequent section, it is not necessary to elaborate a resemblance theory in much detail in order to appreciate that the whole approach faces fundamental problems. But to get to the point where we can see this, we will have to be precise about those details that we are forced to include along the way.

2.1. Fixed Resemblance

The cornerstone of contemporary resemblance theories is the philosopher's concept of similarity, according to which similarity is defined simply as sharing of properties.[22] Yet if similarity is defined this way, depiction cannot be grounded in total similarity, for pictures and the scenes they depict nearly always differ in *some* respects. For example, the picture of Obama is flat, while Obama himself is spatially extended. Since total similarity is too demanding, a more reasonable theory should employ a more restricted notion of similarity.[23] Such a notion may be defined relative to a conscribed set of properties relevant to comparisons of similarity. Let X be such a restricted set of properties. Then we may say that any objects A and B are SIMILAR WITH RESPECT TO X just in case A and B have all the same properties, of those included in X. More explicitly:

> ***Restricted Similarity*** for any A and B, and any set of properties X:
> A and B are similar with respect to X if and only if for any property F in X,
> A is F if and only if B is F.

The most basic resemblance theory hypothesizes that a common, but suitably restricted, notion of similarity underlies all instances of

22. A consequence of this definition is that everything is similar to itself, since every object has the same properties as itself. And if A is similar to B, then B is similar to A, for parallel reasons. In these respects the philosopher's concept may diverge from the layman's. Other conceptions of similarity are considered as bases for a theory of depiction by Eco (1979, sec. 3.5) and Van Fraassen (2008, 17–20).

Here and throughout I shall assume a liberal conception of properties, allowing that for every possible distinct predicate, there is a corresponding distinct property. I will also quantify over properties freely, with the assumption that nominalistic reconstructions of the various definitions provided below can be articulated if necessary.

23. See Neander 1987, 214, for the same observation.

depiction. For example, this might be similarity with respect to shape, or similarity with respect to the visual impression elicited among viewers.[24] To express this analysis, the resemblance theorist posits a fixed set of properties $\mathcal{F}$ that are those properties relevant to determinations of depiction in general. The initial resemblance theory of depiction is then formulated as follows.[25]

R1 *Fixed Resemblance*
for any picture P and scene S:
P accurately depicts S if and only if P is similar to S with respect to $\mathcal{F}$.

We may illustrate this theory diagrammatically. Where a picture P accurately depicts a scene S, we shall draw: $P \rightarrow S$. Where P and S are similar with respect to $\mathcal{F}$, we draw: $P \overset{\mathcal{F}}{\leftrightarrow} S$. According to R1, the drawing of Obama is an accurate depiction of Obama if and only if the drawing itself is similar to the man with respect to $\mathcal{F}$. This condition is displayed below.

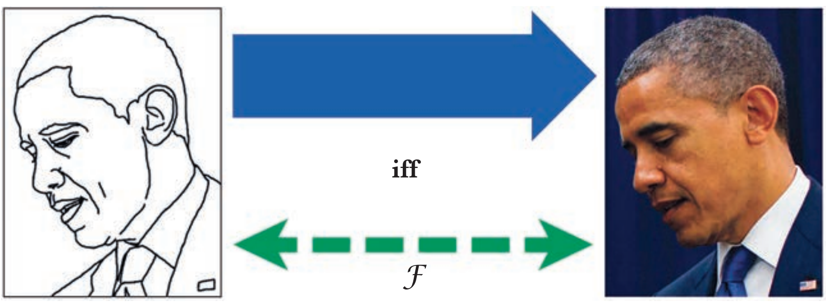

Picture — Scene

According to R1, one kind of similarity grounds all instances of accurate depiction, hence the "fixed" character of the analysis. Of course, this species of similarity must be sufficiently lenient to allow that a flat, black and white line drawing may be relevantly similar to an extended, full-color scene. But it must also be sufficiently stringent that, for example, the inaccurate spiky-hair drawing of Obama is not similar to its subject in the same way.

R1 might be better classified as a *schema* for theories of depiction, rather than a theory itself, because it does not specify the contents of $\mathcal{F}$.

24. The latter is the view of Peacocke (1987) and Hopkins (1998).
25. Such a theory is described by Lopes (1996, sec. 1.2).

(Still, I will often refer to it as a "theory.") In principle, any number of more specific accounts could be based on it. Yet despite this range of potential articulations, it is possible to show that *any* analysis based on R1 will fail. It will fail because R1 treats depiction as if it were a single, invariable relation. But in fact, the definition of accurate depiction varies according to system of depiction, and different systems require that we interpret images in incompatible ways.[26]

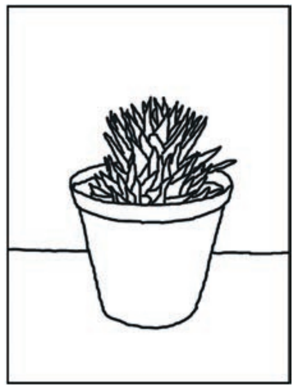

To see this, consider the first image at right. There is a system of depiction for colored line drawings, call it *C*, according to which this is an accurate pictorial representation of my plant. It is accurate, evidently, in part because of certain resemblances between the specific colors in the picture itself and the colors in the scene depicted; very different colors in the picture would have yielded inaccuracy. Meanwhile, according to a different system of depiction for black and white line drawings, call it *D*, the second drawing above is accurate. *D* clearly does not require similarity with respect to color, for this image is black and white while its subject is multicolored. If *D* did require similarity with respect to color, then the picture would be accurate only if the plant and pot were paper-white, on a paper-white desk, against a white wall—which of course they are not.

Thus we see that (i) some systems of depiction (such as *C*) *must* require similarity with respect to color for accuracy, but (ii) other systems of depiction (such as *D*) *cannot* require similarity with respect to color for accuracy. The conclusion is that R1 must be wrong, for there is no single

26. The objection below follows the reasons in Neander (1987, 214–15) for rejecting any resemblance theory defined in terms of a universal notion of resemblance. But whereas Neander answers the objection by making resemblance flexibly context dependent, I make it system relative. The same criticism of unified approaches to resemblance is developed by Lopes (1996, 20–32), albeit with respect to a particular version of the resemblance theory in Peacocke 1987.

class of properties, fixed across all systems of depiction, such that similarity with respect to that set of properties is necessary for depiction. Instead, if resemblance theory is to succeed at all, it must make similarity relative to systems of depiction.

2.2. Variable Resemblance

In order to accommodate systems of depiction, let us revise the resemblance theory so that a picture accurately depicts a scene only relative to a system, and each such system specifies its own kind of similarity underwriting accurate representation.[27] Thus a color line drawing is properly evaluated relative to a system of depiction for which similarity with respect to color, as well as shape, are conditions on accuracy; meanwhile a black and white line drawing is properly evaluated relative to a different system of depiction for which similarity with respect to shape but *not* color is a condition on accuracy. The tension caused by the examples above is dissolved.

To state such a theory compactly and explicitly, let us introduce the following nomenclature. When a picture P accurately depicts a scene S relative to a system of depiction I, let us say that "P accurately depicts$_I$ S." Next, let us call the kind of similarity relevant to I, "similarity with respect to $\mathcal{F}_I$"—where $\mathcal{F}_I$ is a set of properties determined by I. Then the improved, variable view can be stated as follows:[28]

R2 *Variable Resemblance*
for any picture P, scene S, and system of depiction I:
P accurately depicts$_I$ S iff P is similar to S with respect to $\mathcal{F}_I$.

Again, we may illustrate this proposal diagrammatically. Where a picture P accurately depicts a scene S relative to a system I, we shall draw: $P \xrightarrow{I} S$. Where P and S are similar with respect to $\mathcal{F}_I$, we draw: $P \xleftrightarrow{\mathcal{F}_I} S$. According to R2, the black and white drawing accurately depicts the plant-scene

27. How is the relevant system selected for a particular communicative act? Novitz (1975) suggests that the artist's intentions are determinative. Abell (2009) provides a more nuanced account that accommodates the influence of both artistic intentions and communicative conventions.

28. While none since Goodman (1968) have defended this version of the resemblance theory—the reason why is discussed below—the move to system relativity was clearly anticipated by Manns (1971) and recapitulated by Novitz (1975). Lopes (1996, sec. 1.5) describes but does not endorse the variable approach. Variants on the view are defended in detail by Malinas (1991, 282–91) and Abell (2009).

relative to the system of line drawing D iff the drawing and the scene are similar with respect to $\mathcal{F}_D$.

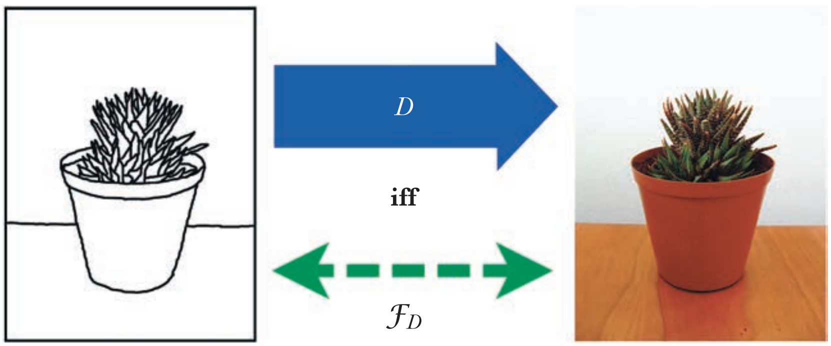

Picture Scene

Whereas accurate depiction in the system of color drawing C is grounded in a different kind of similarity—similarity with respect to $\mathcal{F}_C$.

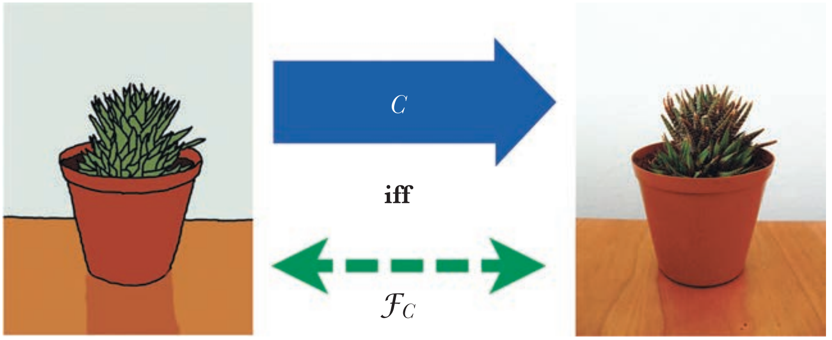

Picture Scene

So understood, the resemblance theory puts a substantive constraint on the structure of *all* systems of depiction, not merely *one* or *some* systems of depiction. It is uncontroversial that there are some systems that are governed by resemblance. (The system by which paint chips are used to depict the color of paint from a bucket is a plausibly obvious example.) It remains to be seen whether all systems of depiction can be analyzed in terms of resemblance—even when the species of resemblance in question varies from system to system. Though this is an ambitious thesis, R2, by design, leaves much to be filled in, thus allowing considerable flexibility in the elaboration of more detailed accounts. Line drawing and color systems each determine their own appropriate kind of similarity. The same approach could be extended to black and white and color photography, and plausibly such challenging cases as color-negative or X-ray photography.

Unfortunately, despite its apparent flexibility, R2 is susceptible to a pair of objections put forth by Nelson Goodman (1968, 4). Simply put, Goodman's complaint is that similarity is a reflexive and symmetrical relation, but depiction is neither, so depiction cannot be equivalent to any kind of similarity.[29] To begin, consider a line drawing that accurately depicts a scene involving Obama in system D.[30] According to R2, similarity with respect to $\mathcal{F}_D$ grounds depiction for drawings in D. So the picture and the scene are similar in that respect.

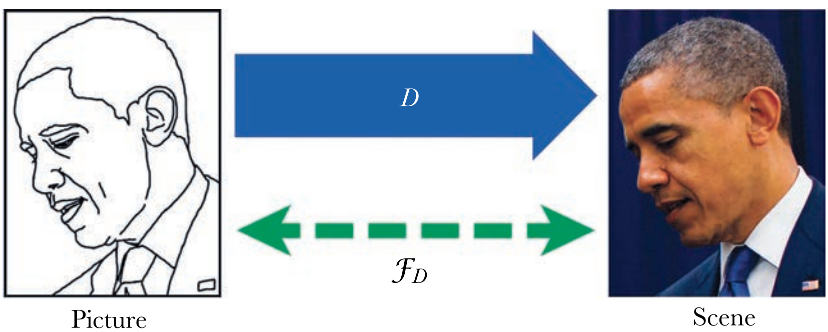

29. It is not entirely clear that Goodman would approve of this application of his argument because he may have been attacking resemblance theories of pictorial reference rather than of accurate depiction. Walton (1990, chap. 3, 122–23n11) surveys the textual evidence both for the view that Goodman's notion of pictorial representation is most like Walton's notion of "representation" (≈ pictorial reference) and for the view that it is more akin to "matching" (≈ accurate depiction). Walton concludes that the textual evidence is contradictory. In a subsequent discussion, Lopes (1996, sec. 3.2) suggests that Goodman had in mind both relations and believed they were independent. For the sake of storytelling, if nothing else, I shall continue to speak as if Goodman's objection was targeted at some theory of accurate depiction like R2.

30. Several of Goodman's original cases are susceptible to the complaint that resemblance theory aims only to define depiction for *pictures* not objects in general. Here I have been careful only to use cases in which depiction, or its failure, is ascribed to pictures. This has the unfortunate result that the counterexamples provided here are more involved than those described by Goodman.

Yet similarity is reflexive: for any A, A is similar to A.[31] This holds no matter how $\mathcal{F}_D$ is defined. Thus the picture of Obama is similar to itself with respect to $\mathcal{F}_D$. Yet this picture does *not* depict itself (see image below). Instead, it depicts Obama. In this case, depiction and similarity come apart, contrary to the claims of R2.

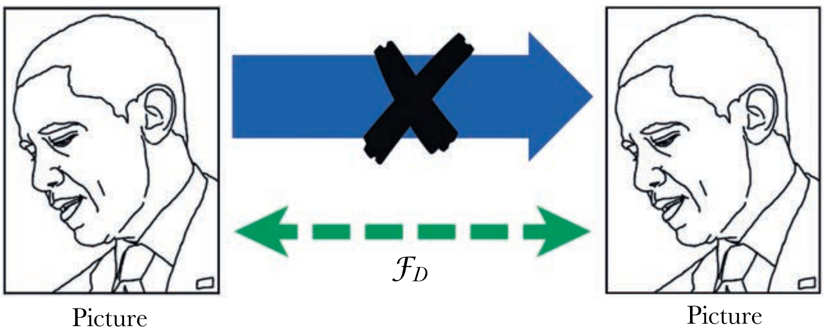

As it is possible to have similarity without accurate depiction, it follows that similarity alone is not sufficient for depiction. And since the argument does not assume any particular way of defining similarity with respect to $\mathcal{F}_D$, it follows that no kind of similarity is sufficient for accurate depiction. The objection cannot be met by prohibiting pictures from depicting pictures; for obviously pictures *can* depict other pictures. Alternatively, one might try to block the objection by requiring that the picture and the scene be distinct things. But this response comes at too high a cost since plausibly pictures *can* depict themselves.[32] In any case, a closely related objection based on the symmetry of similarity cannot be defeated in this way.

Suppose I create a color drawing of Obama, and a black and white line drawing of the color drawing—a picture of a picture. In the line drawing system D, the black and white drawing is an accurate depiction of the color drawing. So according to R2, the black and white drawing is similar to the color drawing with respect to $\mathcal{F}_D$.

31. That similarity is reflexive and symmetrical follows from the definition of similarity in terms of the sharing of properties.

32. Newall (2003, 386–87) suggests that a picture titled "Picture 12" with the caption "A picture of Picture 12" could accurately depict itself.

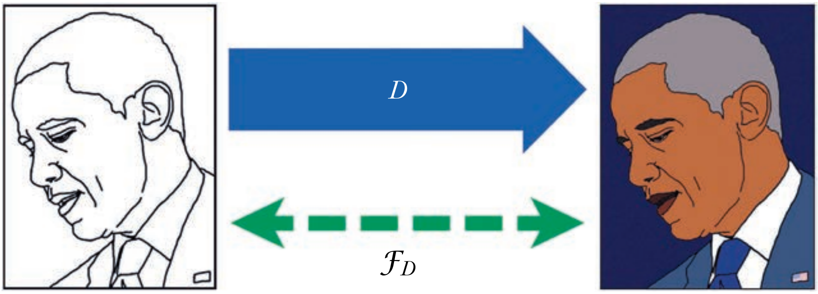

Line drawing Color drawing

But similarity of any kind is symmetrical: for any A and B, if A is similar to B, then B is similar to A. Thus the color drawing is similar to the black and white drawing with respect to $\mathcal{F}_D$. Yet the color drawing does not depict the line drawing. Instead, it depicts Obama. Once again depiction and similarity come apart, contrary to the claims of R2.

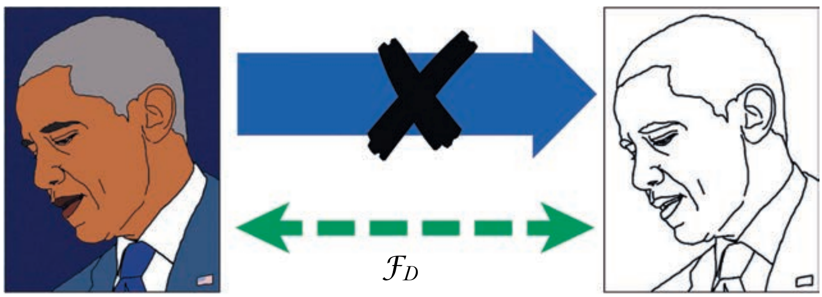

Color drawing Line drawing

Here again we find a case of similarity without depiction; thus similarity alone is not sufficient for depiction.[33] The resemblance theory must be revised.

2.3. The Reference Condition

While Goodman's objection is valid and illuminating, it is easily deflected, as Goodman (1968, 6) himself observed in passing. The proffered

33. While this is the same conclusion I advocate in section 3, Goodman's reasons and mine are fundamentally different. Furthermore, my argument targets all resemblance theories, even those that are amended so as to circumvent Goodman's objection.

counterexamples to resemblance theory all involve pairs of objects in which one object does not pictorially represent the other *at all*, let alone represent it accurately. Thus the line drawing of Obama is not a pictorial representation of itself, accurate or inaccurate, and the color drawing is not a pictorial representation of the line drawing, accurate or inaccurate. These are the facts that the objection trades on. One solution, then, is to augment the resemblance analysis with the requirement that the first object bears the relation of pictorial *reference* to the second, however accurately or inaccurately it depicts it.[34] This condition voids each of Goodman's counterexamples. For the line drawing does not refer to itself, and the color drawing does not refer to the line drawing.

This is not the only solution available—any condition that captures the communicative asymmetry between Obama and the picture of Obama will do. For example, several recent defenders of resemblance pursue a Gricean account of pictorial communication, effectively requiring that, for successful depiction, not only must there be resemblance, it must also be the case that the artist intends that the picture be recognized as resembling the object (Abell 2009, 211; Blumson 2009a). I do not wish to try to adjudicate between the various alternative responses to Goodman's worry. All have the same schematic form of imposing a suitably asymmetric "reference condition" on the picture P and the scene S. I shall abbreviate this relation REF, writing the entire representation condition REF(P, S). We can define this schematic but improved version of the resemblance theory as follows:

> R3 *Variable Resemblance with Reference Condition*
> for any picture P, scene S, and system of depiction I:
> P accurately depicts$_I$ S iff P is similar to S with respect to
> $\mathcal{F}_I$ & REF (P, S).

The reference condition REF(P, S) can then be fleshed out in any way the best theory sees fit, so long as it is strong enough to answer Goodman's objection but weak enough to be compatible with both accurate and inaccurate pictorial representation. In the resulting version of the resemblance theory, the reference condition and resemblance condition divide their work. The reference condition guarantees the minimal

34. This reply is considered by Lopes (1996, 18) and defended, more or less, by Files (1996, 404–5). Blumson (2009a) argues convincingly that it is not enough to require merely that (the artist intend that) the first object *represent* the second—the solution suggested by Goodman (1968, 6)—but that (the artist intend that) the representation be distinctively pictorial.

requirements for both accurate and inaccurate pictorial representation. It is the task of the resemblance condition to determine the accuracy of the representation.

Thus, according to R3, the drawing of Obama accurately depicts its subject relative to system D if and only if the picture and the scene are similar with respect to $\mathcal{F}_D$, and in addition, the picture depictively refers to the scene (or the like). Here again we may illustrate the relation of pictorial reference between a picture P and scene S by $P \longrightarrow S$:

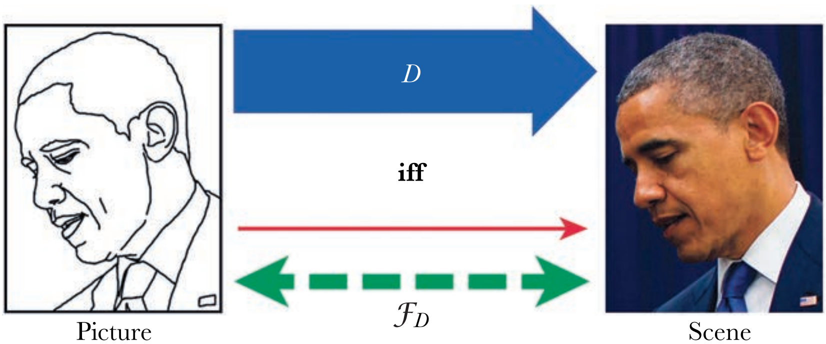

R3 presents us with a plausible and apparently flexible schema for theories of resemblance. At the same time, it expresses the substantive hypothesis about depiction that for any given system of depiction, the accuracy conditions for that system are grounded in similarity. As we will see in the following section, the demands imposed by such a theory are in fact rather stringent.

2.4. Linear Perspective

At the outset I identified photography and perspectival drawing as prototypical examples of pictorial representation, systems that any theory of depiction should be able to account for. In this section, we will consider a particularly common such system: line drawing in linear perspective. This system deserves special attention because it is ubiquitous and because it has already been largely codified through the development of projective geometry.

Linear perspective is based on *linear projection*, a geometrical technique for transposing a three-dimensional scene onto a two-dimensional surface, much the way a flashlight may be used at night to project the shadow of a spatially extended object onto a flat wall. Linear perspective projection is just one of infinitely many such 3D-to-2D mappings, but it has special human interest, for it was developed, with difficulty and over several centuries, by artists and scholars attempting to recreate human perceptual experience on paper.

This endeavor proved successful in several notable respects (Gombrich 1960, 4–30). For example, perceptual experience and linear perspective projection have in common that objects that are more distant from the vantage point of the artist/viewer are represented by smaller regions on the picture plane/visual field. A consequence of this effect is that parallel linear objects in the scene—such as the rails at right—are represented by converging lines on the picture plane. For while every rail tie is in fact the same width, they are represented by shorter and shorter lines on the picture plane as they recede from the camera.

Nevertheless, linear perspective is the product of geometry, not biology. So it is a matter of debate whether linear perspective, or some nonlinear alternative, best models the structure of visual perception. But for our purposes, it doesn't matter how this debate is settled. This is because, since the Italian Renaissance, linear perspective projection has become extraordinarily popular in its own right as a technique for making pictures. Today, drawings made in linear perspective are the norm, and nearly all mass market camera lenses are designed to mimic its results. All the accurate drawings exhibited in this essay thus far have been produced according to systems of linear perspective.

The challenge for resemblance theorists is to show that the antecedently discovered geometrical understanding of linear perspective can be reformulated in terms of resemblance. As we will see, the challenge is acute but surmountable.

I'll begin here by presenting an objection to resemblance theory that originates with Descartes. The objection shows that accuracy in linear perspective often depends on the *differences* in shape between the picture and the scene it depicts, so it cannot be analyzed straightforwardly in terms of intrinsic similarity. Instead, a successful resemblance theory must somehow incorporate the key notion of projection relative to a viewpoint. This in turn will require a revision of R3. The remaining discussion is devoted to working out the form of this revision. After rehearsing the Cartesian argument, I describe in outline the mechanics of linear perspective projection. I go on to show how this concept can be used to define a precise constraint on any theory of depiction for linear perspective. Finally, I'll show how resemblance theorists can meet this constraint. The development and success of such a projection-based account is a major triumph for resemblance theory. Only in the next part of the essay will we see that even this sophistication cannot save the theory from demise.

2.4.1. Linear Perspective and Necessary Difference

Consider again the drawing of Obama and the scene it accurately depicts. The challenge faced by the advocate of a resemblance theory like R3 is to show how accuracy in this case may be understood in terms of similarity between the picture and the scene depicted. Clearly, accuracy depends in large part on some kind of "fit" between the shape of regions in the picture and the shape of the scene it represents; if the picture is misshapen, it cannot be accurate. Letting L be the system of linear perspective line drawing to which the picture in question belongs, the most obvious way to define accuracy in term of similarity is something like the following. For any picture P and scene S:

(1) P accurately depicts$_L$ S iff P is similar to S with respect to shape & REF (P, S).

Proposition (1) aims to analyze accuracy in L in terms of simple similarity of shape. Here the requirement of similarity with respect to shape need not be understood, obtusely, as implying that the scene must be square because the picture plane is square. Instead, we should understand similarity of shape as the sharing of intrinsic properties of shape between the regions defined by the picture and the regions that make up the scene. For the sake of argument, let us assume that there is some systematic way of working out the regions defined by each entity in the intended way.

Nevertheless, it is not clear how (1) can be made to work. The theory fails, first of all, on the simplest and most general conception of shape. Every region on the picture plane is flat, but the scene it depicts and the objects in that scene are spatially extended. Thus the picture and the scene do not have the same shape in this general sense. But perhaps there is a more restricted notion of intrinsic shape that will do. For example, perhaps there is a planar slice of the scene such that the picture and this slice are similar with respect to shape. Yet this is also a nonstarter, for the regions specified by such a plane would perforce include the *internal* structure of the scene depicted—for example, the structure of Obama's skeleton and internal organs. By contrast, the lines of a drawing describe the *external* shape of the objects depicted; such shapes are extended in three dimensions, across many planes, not one privileged slice. Is there some other way of defining shape that vindicates (1)?

Adapting an argument from Descartes's *Optics* ([1637] 2001), I will argue in this section that there is not: on any way of understanding similarity with respect to shape, (1) is wrong. No notion of similarity in shape, no matter how restricted or nuanced, can ground depiction. But whereas Descartes took these considerations to undermine resemblance theory generally, I draw a different moral: accurate depiction in linear perspective must instead be defined in terms of similarity with respect to *relational* features of picture and scene. This conclusion will provoke a natural revision of R3, taken up in the following subsections, in which resemblance is defined in terms of projection relative to a chosen viewpoint.

Descartes's original argument is disarmingly compressed:

> This resemblance [between a picture and its subject] is a very imperfect one, seeing that, on a completely flat surface [pictures] represent to us bodies which are of different heights and distances, and even that following the rules of perspective, circles are often better represented by ovals than by other circles; and squares by diamonds rather than by other squares; and so for all other shapes. So that often, in order to be more perfect as images and to represent an object better, they must not resemble it. (Descartes [1637] 2001, 90)

Descartes argues against resemblance theories on the grounds that, in perspective depiction, a more "perfect" image will often resemble its subject less than an "imperfect" one; hence pictorial perfection in such a system and resemblance cannot be equivalent. Taking "perfection" to be accuracy, I now turn to reconstructing Descartes's argument in greater detail.

Consider a flat white surface traversed by two parallel lines running on forever in both directions. From a bird's-eye view, they look like this:

Now suppose you and I are each charged to draw these parallel lines in linear perspective. Further, we have been instructed to draw exactly the same scene, from the same precisely specified vantage point, with the same position and orientation of the picture plane. We both intend to comply with these instructions. Taking turns, we each render the scene to the best of our abilities. Unfortunately, while you are an expert practitioner of linear perspective, I am a novice, and I misjudge certain crucial angles during the construction of my image. As a result, our pictures come out subtly different:

(2)

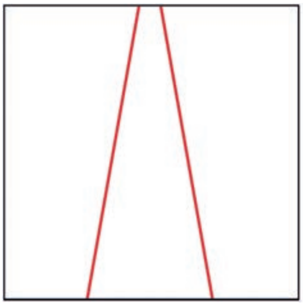

(3)

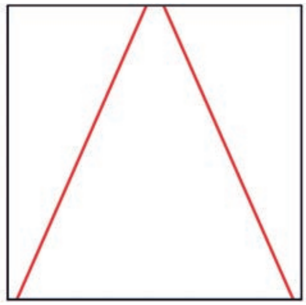

I have drawn (2), while you have drawn (3). The difference, visible above, is that in (2) the angle between the two lines is narrower than in (3). Intuitively, given the strictly specified vantage point of the drawing, your rendering is perfectly accurate, but mine is not. In the illustration below, I have included, first, my picture beside a bird's-eye view of the scene that it *would* accurately depict, and, second, your picture and the bird's-eye view of the scene it *actually* accurately depicts.

(2)

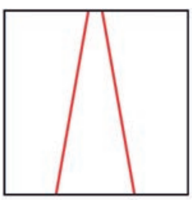

Perspective

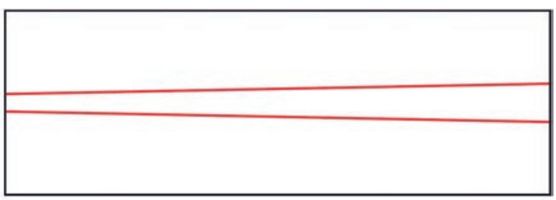

Bird's eye view

(3)

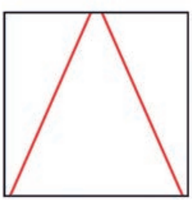

Perspective

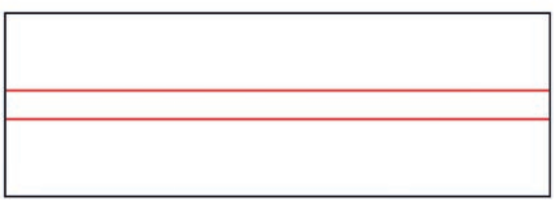

Bird's eye view

These examples now form the basis for a general argument that there can be no adequate analysis of perspective depiction in terms of similarity of intrinsic shape, no matter how nuanced or restricted the relevant notion of shape. The argument proceeds from three key premises, where *S* is the scene consisting of the two parallel lines described above.

Premise 1. Image (3) accurately depicts *S* in linear perspective.
Premise 2. Image (2) *in*accurately depicts *S* in linear perspective.
Premise 3. For every respect of intrinsic shape relative to which (3) is similar to *S*, (2) is similar to *S* in the same respect.

Informally, the argument is just this: by premise 1 and 2, image (3) and image (2) have different accuracy values; but by premise 3, they are on par with respect to intrinsic similarity to *S*. Thus an analysis in terms of intrinsic similarity lacks the resources to explain their divergence in accuracy. Formally, the argument proceeds as a reductio of (1), the proposition that depiction can be defined in terms of similarity with respect to shape. Let the shape properties invoked by (1) be arbitrarily selected. Begin by supposing (1); it follows from premise 1 that image (3) is similar to *S* in the relevant respects of shape. By premise 3, it follows that image (2) is similar to *S* in the *same respects*. By (1) again, it follows that image (2) accurately represents *S*. But this contradicts premise 2, that image (2) *in*accurately represents *S*. We conclude that (1) is false.

The third, crucial premise requires explanation, for it follows from a subtle feature of the scenario described above. The illustration below includes first the parallel lines of *S* itself, shown again in bird's-eye view, this time along the vertical axis, second the *in*accurate narrow image of the scene, and third the *accurate* wide image of the scene.

S (2) (3)

Begin by considering those shape features of image (3) held in common with *S*—for example both are planes, both planes contain only two straight line segments, both sets of line segments are nonintersecting, both sets of lines are tilted at symmetrical angles, and so on. But note that all of these shared features are also ones that image (2), the inaccurate image, holds in common with *S*. Of course, there *is* a difference between image (2) and image (3); the lines in the former define a more acute angle than those in the latter. But this does not make image (3) more similar to *S* than image (2) in any way. This is because, if anything, the lines of image (2) are more nearly parallel than those of image (3), while those in *S* are perfectly parallel; so image (2) shares the property of minimal degree of parallelism with *S*, but image (3) does not. Recall that the claim defended here is not that image (3) and image (2) have all the same intrinsic shape properties in common with *S*, but only that whatever such properties image (3) has in common with *S*, image (2) does as well. The difference in angle between the two does not undermine this claim.

Note that the argument was deliberately built around a case in which the image that is *more* similar to the scene is *in*accurate, while the image *less* similar to the scene is accurate. For if the scenario had been inverted, so that image (3) were inaccurate and image (2) accurate, it would have been possible to object that image (2) had something in common with the scene that image (3) lacked, namely, degree of parallelism. The objector could have claimed that this difference in degree of parallelism explained the accuracy of image (2) and the inaccuracy of image (3). But as the case actually stands, no such response is available.

The most obvious way of analyzing accurate depiction in linear perspective as grounded in similarity has failed. This failure illustrates an important lesson: accuracy under the system of linear perspective is not determined by merely copying the intrinsic properties of the scene, however selective the copying. As Descartes appreciated, this feature is a consequence of the fact that pictures in linear perspective are created and interpreted as projections from *viewpoints*. If similarity underlies accurate depiction, it must be a species of similarity that is in some way sensitive to this viewpoint relativity. In the remainder of this section, we will develop a resemblance-based analysis of linear perspective that incorporates these insights.

2.4.2. Linear Perspective Projection

In order to make headway we must secure a more detailed understanding of the mechanics of depiction in linear perspective. We begin with a short exposition of LINEAR PERSPECTIVE PROJECTION (for short, LINEAR PROJECTION), the geometrical technique for mapping three-dimensional scenes onto two-dimensional surfaces, which lies at its heart.

The basic mechanism of linear projection is vividly illustrated by a thought experiment originally suggested by Leon Battista Alberti and elaborated by Leonardo da Vinci (see Alberti [1435] 1991, book 1; and Da Vinci [c. 1500] 1970, 2.83–85). To begin, suppose you are looking out of a flat glass window onto a static scene. But instead of looking *through* the window to the scene, as one normally does, concentrate instead on the surface of the glass itself, where the rays of light reflected from the scene pass through the translucent medium of the glass. Now take a marker in hand and proceed to draw on the surface of the windowpane, according to the following rule: wherever you see the salient edge of an object through the window, draw a line on that part of the windowpane where that edge is reflected. Continue until you have marked every salient edge. (The exercise is much easier to conceptualize than to actually perform.) At the end, you are left with a drawing of the scene on the surface of the glass—and you have effectively constructed a representation in linear perspective.

Linear projection simply formalizes and abstracts from this thought experiment: the windowpane is replaced by a geometrical plane, rays of light by lines, and the eye by a point. The result is an algorithm for mapping three-dimensional scenes onto two-dimensional picture planes.[35] To illustrate exactly how such a projection is defined, suppose the scene that we wish to represent consists only of a grey cube in white space. We begin by fixing a VIEWPOINT—literally a geometrical point in space—at some distance from the cube.[36] The viewpoint determines the "perspective" of the picture. We then introduce PROJECTION

35. The mathematical details of the following discussion are based on presentations in Sedgwick 1980; Hagen 1986, chaps. 2–5; and Willats 1997, chaps. 2, 5.

36. The viewpoint is often called the "station point"; here I opt for the more intuitively vivid term. Informally, I will also refer to it as the "vantage point." Note that the viewpoint is not the same thing as a picture's "vanishing point." The viewpoint is a position in the space of the scene, outside the picture plane, introduced to define a perspective projection. A vanishing point is any position in the space of the picture plane to which two lines depicting parallel lines in the scene will converge. There may be as many vanishing points in a picture as there are such pairs of lines. The distinction between "one-point," "two-point," and "three-point" perspective describes vanishing points, not viewpoints.

LINES that trace straight paths between select points on the cube and the viewpoint.[37]

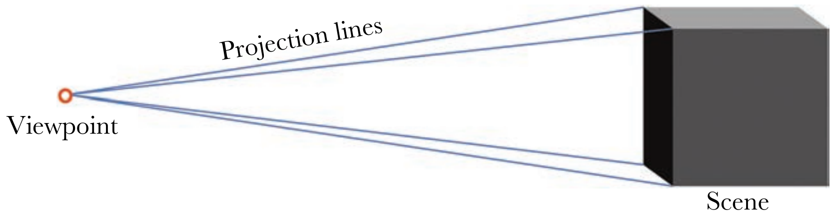

Crucially, we need not draw a line from *every* point on the cube to the viewpoint. Which points in the scene are relevant will depend on the style of perspective representation under consideration. For present purposes, let us consider only those points that lie along an edge of the cube; of these, let us select only those that can be connected to the viewpoint by a straight line, without passing through a surface of the cube. The resulting projection lines can be thought of as "lines of sight," since, like sight, they link the viewpoint to only those surfaces in the scene that are not occluded by another surface. For visual clarity in the diagrams above and below, I have indicated only those projection lines that connect the viewpoint to the accessible corners of the cube.

Into the spray of projection lines emanating from the viewpoint, we now introduce a PICTURE PLANE—analogous to Alberti's window. The picture plane can be positioned at any angle, but customarily, as suggested in the illustration below, it lies perpendicular to the line that connects the center of the scene to the viewpoint. At every point at which a projection line intersects the picture plane, a point is inscribed on the picture plane. When such points form a continuous line, a corresponding line is inscribed on the picture plane. If the picture plane is now seen face on, it reveals a side view of the cube.[38]

37. The diagrams of this essay are all drawn according to a highly imprecise system of depiction. The information they convey is impressionistic, not exact.

38. To be clear, talk of inscription here is purely metaphorical; to "inscribe" a point or line on the picture plane is just to define the geometry of the picture plane in a certain way. Also, in the account given here, edges in the scene are correlated with lines on the picture plane. But other styles of linear perspective can be produced by varying this choice. For example, edges themselves can be defined in a number of ways (DeCarlo et al. 2003). Beyond line drawing, there are other possibilities: in color drawings, colored surfaces in the scene are matched with regions of the same color on the picture plane; in cross-hatch

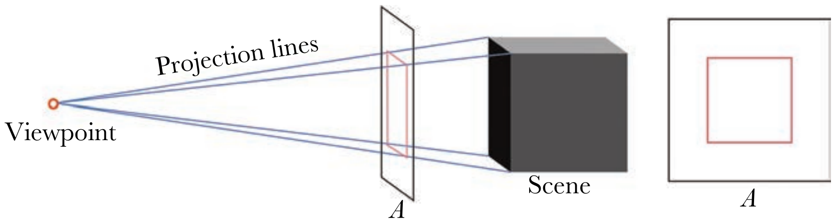

We can further manipulate the relative positions of the picture plane and viewpoint, with predictable consequences for the resulting image. The following figure illustrates the results of altering the location and orientation of the picture plane. Shifting the plane closer to or farther from the viewpoint (*B*) has the effect of altering the scale of the resulting image. Shifting the picture plane vertically (*C*) causes the projection of the scene to drift from the center. When the picture plane is tilted (*D*), the resulting image records exactly the same aspects of the cube as (*A*) but introduces the characteristic "railroad-track effect" of perspective projection, where edges of the cube that are in fact parallel are now represented by converging lines.

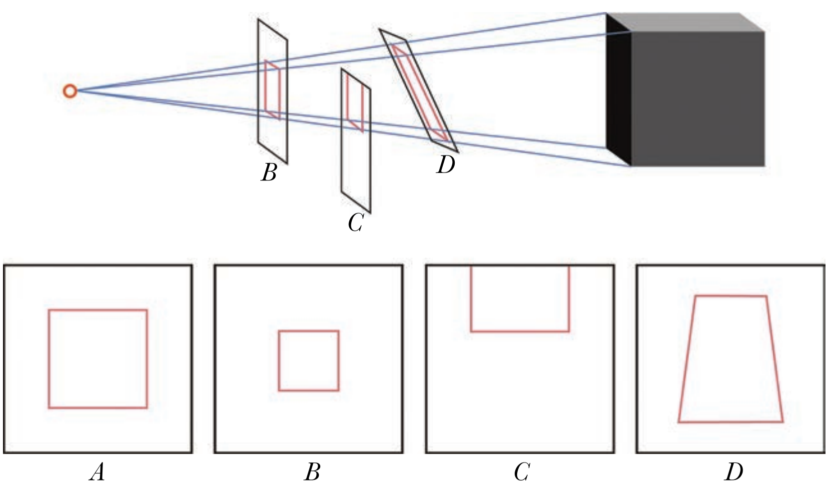

drawings, shaded surfaces in the scene correspond to regions in the drawing with a certain line density; and so on.

Each of the alterations just described changes the *way* the scene is represented, but none reveals additional information about the scene not already reflected in (A). By contrast, when the viewpoint is shifted, new features of the scene are revealed. For example, let us move the viewpoint above and to the side of the cube, while simultaneously sliding the picture plane to intercept the projection lines. The principal effect of repositioning the viewpoint in this way is that it can now "see" two additional faces of the cube, represented below on the picture plane (E).

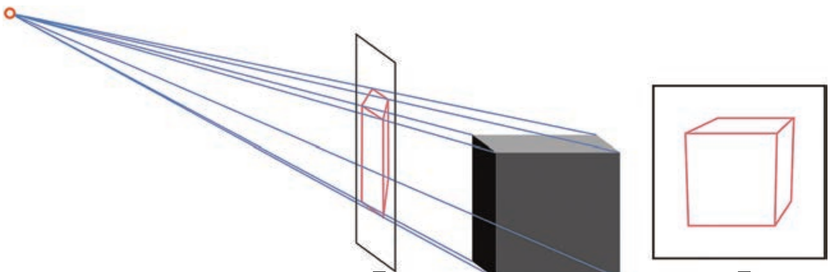

Images (A)–(E) are all produced by the same general method of projection; their differences are due to different selections of picture plane orientation and viewpoint. In general, there is no such thing as "the perspective projection" of a scene, independent of such parameters. As the two diagrams above suggest the image inscribed on the picture plane depends on three factors: (i) the shape of the scene itself; (ii) the position of the viewpoint relative to the scene; and (iii) the position, orientation, and size of the picture plane relative to the scene and the viewpoint. Thus *relative* to a choice of positions for picture plane and viewpoint, the method of perspective projection delivers a unique projection of any scene. Here it is convenient to collect picture plane and viewpoint position together into a single PROJECTIVE INDEX. Then we may say that, relative to a projective index, perspective projection determines a unique projected image of any scene.

Since the mechanism of picture generation is entirely explicit and determinate, it can be represented by a mathematical function, what I will call a PROJECTION FUNCTION. Such a function takes as inputs the spatial properties of the scene and a projective index, and outputs the inscribed picture plane. The mathematical details of this equation do not concern us here; it will suit our purposes best to reduce the formula to its most basic, structural elements.[39] In any given case, we may label the scene S and projective index j. Finally, naming the projection function $\mathsf{lin}(\cdot)$, we shall say that lin applied to S and j returns an inscribed picture plane P, that is, $\mathsf{lin}(S, j) = P$.

2.4.3. Linear Perspective as a System of Depiction

Linear perspective *projection* defines an algorithm for deriving pictures from scenes. The *system* of linear perspective, on the other hand, determines a mapping from pictures to content, and thus a standard of pictorial accuracy. The two are directly related: to a first approximation, a picture accurately depicts a scene in the system of linear perspective just in case there is a linear perspective projection from the scene to the picture. Accuracy in linear perspective requires linear projection.

To focus this proposal, let us confine our attention to L, the system of monochrome line drawing in linear perspective that corresponds to the projective technique described above. Now suppose two artists set out to draw a cube S from the vantage point used to produce image A above, relative to system L. The first picture, (4), is created in such a way that it conforms with linear perspective projection. Intuitively, (4) is an accurate pictorial representation of S in the system of linear perspective. But the second picture (5), even though it deviates only slightly from the intended rule, is such that that there is no index, much less the intended one, relative to which (5) could be derived by linear perspective projection. Intuitively (5) is an *in*accurate pictorial representation of S in the system of linear perspective. Thus accuracy in L and linear projection appear to march in lockstep.

(4)

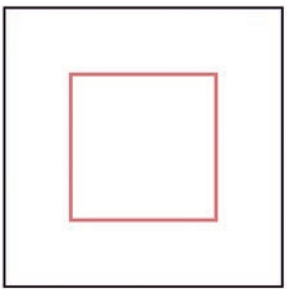

(5)

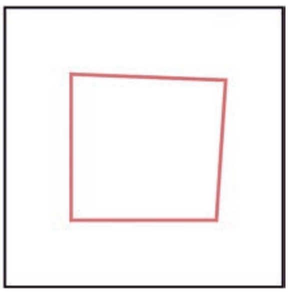

39. An explicit definition is supplied in Greenberg 2012.

This suggests that we can state the accuracy conditions for L quite simply in terms of the projective formula $\mathsf{lin}(S, j) = P$:[40]

(6) P accurately depicts$_L$ S iff $P = \mathsf{lin}(S, j)$.

The problem with this definition is that it makes the projective index j part of the system itself. This would mean that all accurate images in linear projection would have to be drawn from the same viewpoint. This is clearly false: both (8) and (9) below are accurate depictions of S in linear perspective, but only from different viewpoints. The solution is to allow some variability in j. One way to achieve this would be to existentially quantify over the projective index on the right-hand side of the equation.

(7) P accurately depicts$_L$ S iff there is some j such that $P = \mathsf{lin}(S, j)$.

On this analysis, two distinct projections of a scene, such as the two views of the cube below, are both accurate representations of that scene, since both can be projected from that scene according to *some* index.

(8)

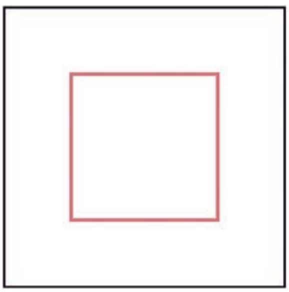

(9)

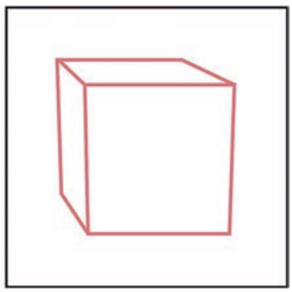

I concede that there is a strand in the concept of accurate depiction that conforms to this analysis, but my interest here is in a more specific notion, with equal currency in everyday usage. In this alternative, depiction is always relative to a projective index, and artists always intend that their pictures be interpreted relative to a particular projective index. If the picture is not a perspective projection of the scene from that index, then it is not accurate—*even if* it is a correct projection of the scene from some other index. Thus, if two artists create (8) and (9) but both intend to

40. To facilitate the statement of this definition, we will also assume that pictures are picture planes—geometrical instead of physical objects. Alternative views may be accommodated without difficulty.

represent the scene from the same vantage point, only one succeeds in accurately depicting the scene. To accommodate this idea, we now analyze depiction as a relation between a picture and scene relative to a system of depiction *and* a projective index:

(10) P accurately depicts$_L$ S relative to j iff $P = \mathsf{lin}(S, j)$.

This formula is nearly right, but somewhat too demanding. Linear perspective is often invoked in such a way that, even when the entire projective index is explicitly fixed, the size of the resulting picture is immaterial to its accuracy. Again, suppose two artists intend to draw the cube from the vantage point of image A, and produce the following pair of pictures:

(11)
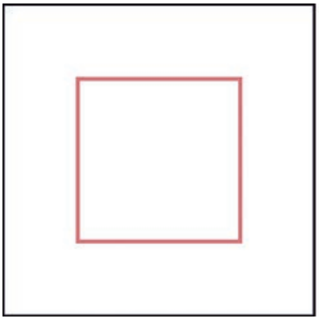

(12)
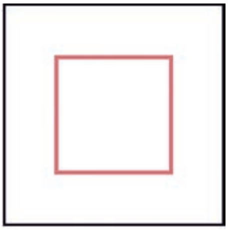

Intuitively, both (11) and (12) are perfectly accurate representations of S, from the vantage point specified in A, despite their difference in size. But this flexibility is not allowed by (10), for $\mathsf{lin}(S, j)$ yields a picture plane of fixed size, and any deviation from this leads to a failure of the biconditional. There are a variety of ways of handling this complication; a simple one is as follows.[41] Let us use the term SHAPE to describe what we normally mean by the intrinsic shape of an object *excluding* its size. In this sense, a large and a small equilateral triangle have the same shape. Then we can revise the definition of L as follows:

(13) P accurately depicts$_L$ S relative to j iff the shape of P = the shape of $\mathsf{lin}(S, j)$.

One final emendation is required. According to (13), any arrangement of shapes could depict a given cube S, relative to an index, so long as it

41. A more general and precise treatment of the same issue is presented in Greenberg 2012.

conformed to the geometric requirements of lin. These arrangements could include lines left by waves in the sand, floor tiling, or an infinity of abstract shapes. Yet arguably, these entities should not all count as accurate depictions of S.[42] It is not that they inaccurately represent S, but rather that some neither accurately nor inaccurately represent it; they are not representations of S at all. The problem appears to be the same as that dramatized by Goodman's objection to resemblance theory. The solution, it seems, should also be the same. What is required is a condition that is sufficient to guarantee pictorial reference without generally determining whether the representation in question is accurate or not. This was just the mandate of the representation condition REF(P, S) introduced in response to Goodman's objection, and I shall rely on it here as well.

> LP *The System of Linear Perspective*
> for any picture P, scene S, index j:
> P accurately depicts$_L$ S relative to j iff
> the shape of $P =$ the shape of lin(S, j) & REF(P, S).

We can illustrate this analysis with a moderate expansion of the essay's diagrammatic idiom. The top half of the diagram below illustrates the condition that a picture P accurately depicts$_L$ a scene S relative to j. The bottom half illustrates the condition that the shape of $P =$ the shape of lin(S,j). For convenience, I hereafter omit illustration of the reference condition.

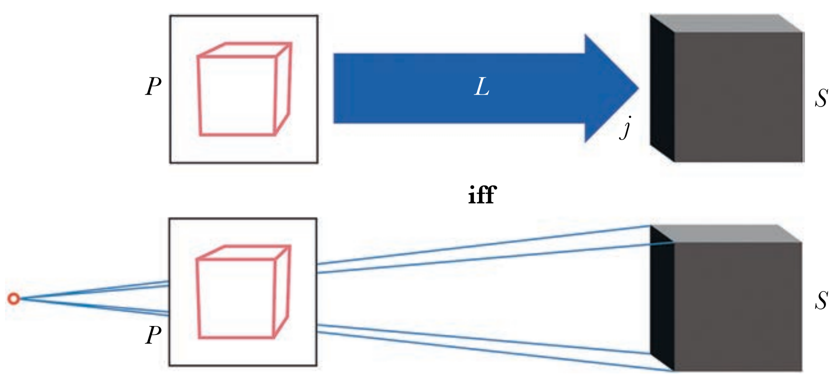

The shape of P = the shape of lin(S, j)

42. This problem was raised by Putnam (1981, 1). The solution adopted here is compatible with Putnam's own.

Proposition LP is a powerful tool for the theorist of depiction. By describing the exact relationship between a scene and a picture required by linear perspective, its sets a precise gold standard for any theory of depiction. The challenge for resemblance theorists is to reformulate LP with strict equivalence in terms that conform to the resemblance schema R3.

2.4.4. Resemblance Theories of Linear Perspective

Given that linear perspective cannot be defined in terms of intrinsic similarity, a natural thought is that it should instead be analyzed in terms of similarity *under* perspective projection. According to this suggestion, a picture depicts a scene just in case the picture is similar in shape to a given projection of that scene. That is, for any picture *P*, scene *S*, and index *j*:

(14) *P* accurately depicts$_L$ *S* relative to *j* iff *P* is similar to $\mathsf{lin}(S, j)$ with respect to shape & REF(*P*, *S*).

According to this proposal, illustrated below, accurate depiction is determined by first taking a linear projection of *S* relative to *j* (on the right side of the figure below) and then requiring that the result be perfectly like *P* with respect to shape.[43]

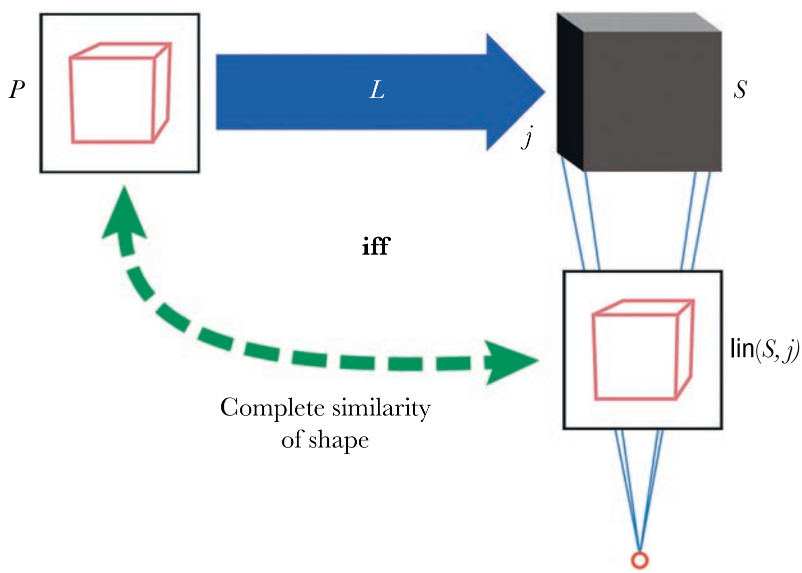

43. A proposal of this sort is defended by Hyman (2000; 2006, chap. 5).

Proposition (14) has the virtue of being equivalent to LP, which we already established as providing the correct conditions for accuracy in L:

LP P accurately depicts$_L$ S relative to j iff the shape of P = the shape of $\mathsf{lin}(S, j)$ & $\textsc{ref}(P, S)$.

According to (14), the condition on accurate depiction is that P and $\mathsf{lin}(S, j)$ are similar with respect to shape; that is, they share all the same shape properties. But this is equivalent to the claim that the shape of P = the shape of $\mathsf{lin}(S, j)$, which is exactly the condition imposed by LP. Unfortunately, despite the fact that (14) succeeds as an equivalent reformulation of LP, it is not a resemblance theory of depiction. For although (14) requires similarity between two things, they are not the picture and the scene—instead, they are the picture and a *projection* of the scene. This is similarity between picture and scene in name but not in substance. For the relation defined here is neither reflexive nor symmetrical, and it cannot be made to comply with R3. In general, similarity under a transformation of only one relatum—such as that posited by (14)—is not similarity between the relata; for real similarity, both relata must be transformed simultaneously.

The point here is not merely one of logical typology. Rather, it is this: if (14) counts as a resemblance theory of depiction, then almost nothing is meant by the claim that depiction is grounded in resemblance. Consider, for example, the relation *being the biological mother of*. This relation, which patently is *not* similarity based in any straightforward sense, can be equivalently reformulated as a similarity relation under the transformation of one relatum. Let the function m map individuals to their biological mothers, and let X be the set of properties of *being M* for every mother M. Then for any individuals x and y, x is similar to $m(y)$ with respect to X if and only if x is the mother of y.[44] Thus similarity under the transformation of one relatum has little to do with genuine similarity. The condition on depiction specified in (14) cannot count as a resemblance theory on pain of undermining the interest of the theory.

But if (14) is not a genuine similarity analysis, how else can the resemblance theorist hope to capture the structure of linear projection encoded in LP? To understand the answer that resemblance theorists have themselves elected, let us return to the original intuition that motivated the resemblance account. We noted that the drawing of Obama and

44. At the present time, I do not know what the limitations of this technique are—whether there are relations that cannot be reformulated in this way.

Obama himself seem to *look similar*. That is, if one were able to view the picture and the person simultaneously, and from the right vantage point, there would emerge a similarity at least in shape between the visual appearance of Obama and the visual appearance of the drawing of Obama. In general, accurate depiction seems to covary with how a picture and the scene it depicts each appear in visual experience. This suggestion is illustrated below.

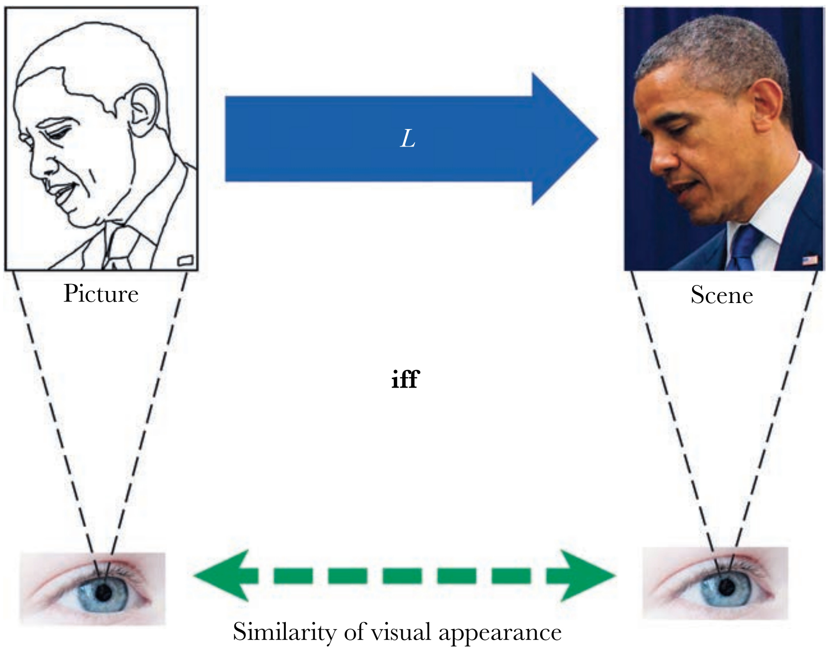

This observation suggests a framework for realizing a genuine resemblance analysis of linear perspective, since we are asked to compare common properties of both relata—the visual appearance of the picture and the visual appearance of the scene. But how can we make this intuitive proposal sufficiently precise? A natural thought is this: following Renaissance scholars, let us treat linear perspective projection as a *model* of visual appearance. Then similarity between picture and scene with respect to visual appearance may be modeled as similarity between picture and scene with respect to the linear projection of each. The idea, illustrated below, is that a picture *P* depicts a scene *S* just in case the linear projection of *S* is similar in shape to the linear projection of *P*.[45]

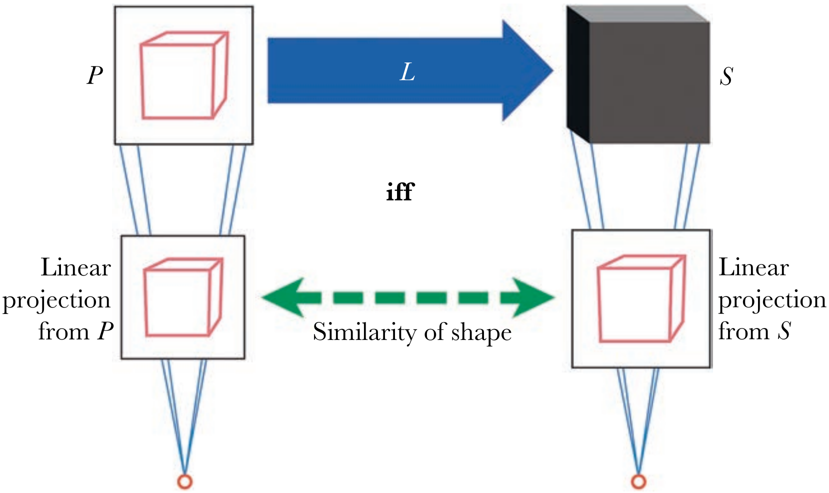

As the diagram makes clear, this proposal relies on the interesting suggestion that we need not treat a picture merely as the result of a projective process; we may also treat it as a scene that is itself projected. The current hypothesis is that, by taking the linear projection of *both* the scene *and* the picture and requiring that the resulting images be alike with respect to shape, we may be able to recapture LP. I shall term this analytic approach the method of SIMULTANEOUS PROJECTION. Unlike the asymmetrical analysis (14) that defines depiction in terms of similarity under selective projection of just one relatum, the method of simultaneous projection requires similarity under the projection of *both* relata.

Note that the method of simultaneous projection requires *complete* similarity of shape between the two derived projections. If it did not—if, for example, it only required approximate similarity of shape—then the conditions for accuracy would be satisfied even if the original picture were only approximately a projection of the scene. The resulting account would predict that a slightly skewed image of a cube could accurately

45. Proposals along these lines are defended by Peacocke (1987) and Hopkins (1998).

depict a cube. But since perfect accuracy makes no such allowances, the similarity in shape between the simultaneously derived picture planes must be total. In what follows, by "similarity of shape," I mean complete similarity of shape.

As we move to make simultaneous projection analysis explicit, we must grapple with a detail that may at first appear to be a technical triviality but in fact has important repercussions. We have seen that perspective projection is always defined relative to a projective index. In the method of simultaneous projection, it is natural to use this index to determine projection from the scene (as on the right side of the above figure), just as we did in the asymmetrical analysis we rejected earlier. But the current proposal also requires that we take another projection of the picture itself. The question is, projection relative to what index?

A first answer is that we simply borrow the projective index used for the scene and apply it to the picture. Then we can define depiction in L as follows:

(15) P accurately depicts$_L$ S relative to j iff $\mathsf{lin}(P, j)$ is similar to $\mathsf{lin}(S, j)$ with respect to shape & REF(P, S).

The problem with (15) is that it makes no sense. A projective index specifies the spatial relationships between a scene, the plane of projection, and the viewpoint. But in general, the spatially extended object S and the flat picture P will have completely different spatial characteristics. As a consequence, there is no general and motivated way of using a single projective index to simultaneously determine the relationship between S and a picture plane, and P and another picture plane. Each requires its own specifications.

The best alternative is to allow that we must specify an independent projective index for the picture. Then we may simply stipulate that depiction is relative to a *pair* of projective indices, where the first is associated with the picture and the second with the scene. Thus we may say, for any picture P, scene S, and pair of indices i and j:

(16) P accurately depicts$_L$ S relative to i, j iff $\mathsf{lin}(P,i)$ is similar to $\mathsf{lin}(S, j)$ with respect to shape & REF(P, S).

According to (16), we determine whether P depicts S in three steps. First we take the perspective projection of S, $\mathsf{lin}(S, j)$. Second we take the perspective projection of P, $\mathsf{lin}(P,i)$. Finally we compare the shapes of the two resulting projections. This proposal raises two questions. First, does it correctly characterize linear perspective? My answer is that it does, but only when certain constraints are imposed on the projective index. Second, is it a genuine resemblance theory, or does it merely have the trappings of resemblance? I will argue that while (16) does not conform with R3, it does conform with a plausible revision of R3. I discuss each point in turn.

To begin, let us see how (16) may be equivalent to LP. First, if $\mathsf{lin}(P,i)$ and $\mathsf{lin}(S,j)$ are completely similar with respect to shape, then the shape of $\mathsf{lin}(P,i) =$ the shape of $\mathsf{lin}(S,j)$. Thus (16) is equivalent to:

> (17) P accurately depicts$_L$ S relative to i, j iff the shape of $\mathsf{lin}(P,i) =$ the shape of $\mathsf{lin}(S,j)$ & REF(P, S).

If we can further show that the shape of $\mathsf{lin}(P,i) =$ the shape of P, as suggested in the diagram above then by simple substitution, (16) is equivalent to LP. By now it should be clear that whether this proposition succeeds depends heavily on the selection of the projective index i. Perspective projection is not an "innocent" transformation; it systematically transforms scenes in a variety of ways. Indeed, having fixed the scene's projective index j, there are infinitely many ways of selecting the projective index i for the picture that demonstrably invalidate (16).[46] For example, suppose that the projective index specifies that the plane onto which P is projected is tilted forward toward P, as in the following illustration.

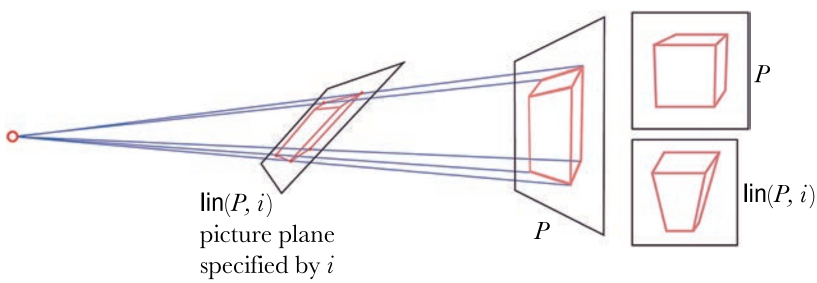

In this case, the result of projecting the scene S relative to j and the picture P relative to i are *not* completely similar with respect to shape, as the diagram below clearly demonstrates. Thus accurate depiction and

46. Lopes (1996, 24) makes this point regarding theory of depiction in Peacocke 1987, a view in many ways similar to (16).

similarity under simultaneous projection may come apart given a poor choice of the projective indices.

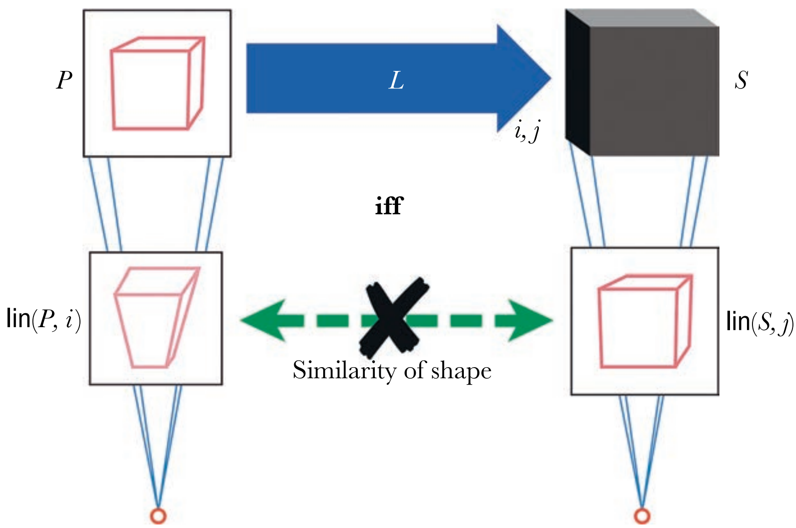

Yet it turns out that we can salvage (16) if we impose special restrictions on the projective index. Recall the analogy to comparisons of visual appearance. Obama and the picture of Obama look similar. What was implicit in this observation was that the picture was viewed *face on*. If we had tilted the picture obliquely away from our eye, it would no longer have been the case that the picture and the scene looked similar. This suggests that, so long as we constrain the conditions under which the picture is viewed to rule out such anomalous angles, then similarity in appearance may be sustained.

This idea has a precise correlate in projective geometry. Simply put: for any given scene, whenever there is a flat surface in the scene, if the picture plane is *parallel* to that surface, then the shape of that surface is preserved perfectly under linear perspective projection—with the caveat that the shape on the picture plane will be smaller in size than the original shape in the scene. This fact is illustrated below. Suppose our scene S consists of two parallel lines on a flat plane. We then take two projections of S. In the first, (A), S is projected onto a picture plane that is parallel to it. In the second, (B), S is projected onto a picture plane that is perpendicular to it.

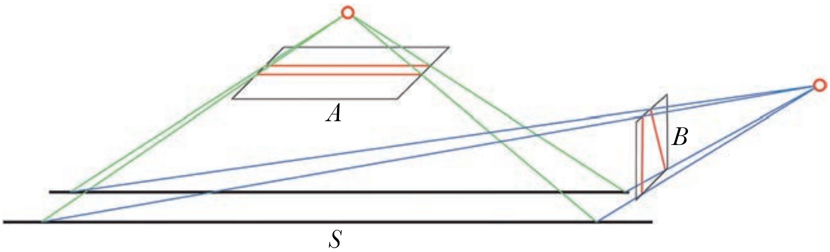

The results of each projection are shown below. Image A, produced from the parallel picture plane is a characteristic "bird's-eye view" of S. Note how the lines in A are parallel, just as they are in S. Meanwhile, image B, produced from the perpendicular picture plane reveals the "train tracks effect" of parallel lines converging. A preserves the shape, though not the size of S. B preserves neither.

(18)

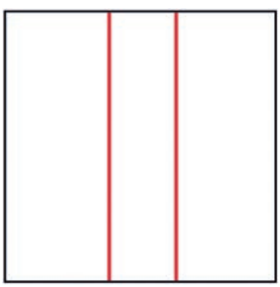

A

(19)

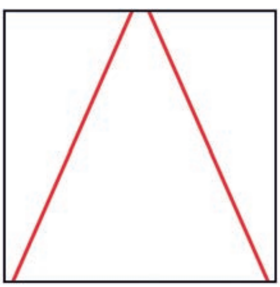

B

We may draw the following conclusion: so long as we constrain the projective index such that P is always projected onto a plane parallel to it, then we are guaranteed that P and $\mathsf{lin}(P, i)$ will have the same shape. Let us call this the PARALLEL PLANES CONSTRAINT.[47] With the addition of this constraint, (16) successfully becomes equivalent to LP.

47. Some authors, such as Hopkins (1998), define similarity by comparing the visual angles subtended by each object relative to the viewpoint, rather than their projected shape on a picture plane. But angle subtended exhibits greater variability than projected shape. I am not currently sure how to state the parallel planes constraint within this framework. In any case, talk of subtended angles, rather than projected shape, renders the operative notion of similarity in shape considerably less intuitive and direct.

We now turn to the question of whether (16) is a genuine resemblance-based analysis. On one hand, (16) is clearly an improvement over its asymmetrical predecessor. For instead of comparing *P* to *S*-under-transformation, (16) compares *P*-under-transformation to *S*-under-transformation. By applying the projection function uniformly, it in effect states that *P* and *S* are similar with respect to shape-under-transformation.

On the other hand, the projective index and the scene index are necessarily distinct; whereas the scene index locates the viewpoint and picture plane at whatever position is specified by the artist, the projective index must always locate the picture plane parallel to *P*. Thus, strictly speaking, different operations are performed on each of *P* and *S* before the comparison of shape. This threatens the claim that (16) is a genuine resemblance analysis. A defensive reply is that depiction may be defined in terms of near similarity—similarity of *P* and *S* under only modestly different transformations. The violation of strict resemblance is slight enough to overlook. This is an option, but I think resemblance theorists have a more elegant solution at their disposal.

Thus far we have assumed that depiction is a relation between a picture and a scene. But suppose instead we let depiction be a relation between a picture-index pair, and a scene-index pair. We may call the first a CENTERED PICTURE and the second a CENTERED SCENE.[48] For a picture *P* and index i, the centered picture is written P_i, and the corresponding centered scene is S_j. Thinking of the relata of depiction in this way is a departure from our original position, but it is not implausible. On one hand, LP taught us that depiction obtains only relative to a selection of index. It is cogent to think that this index is part of our notion of a scene. What a picture depicts is not merely a piece of reality—but a piece of reality from a certain point of view. Centered scenes answer to this description. On the other hand, it cannot be denied that artists create images with intended viewing conditions in mind. The illusory powers of a picture are most active when they are viewed from a particular distance and angle.[49] It is plausible that these conditions are built into our conception

48. The terminology of centering is borrowed from Lewis (1979); the idea of a centered scene as a way of characterizing image content is described by Blumson (2009b).

49. This fact is vividly displayed by cases of anamorphosis, in which an apparently distorted image appears “normal” when viewed from an unusually oblique angle. The present framework elegantly handles anamorphosis by allowing the intended viewpoint for the picture to be realized in the projective index of the centered picture.

of a picture. What depicts a scene is not merely a picture, but a picture from a certain view. Centered pictures in turn answer to this description. Thus the move to centering is harmonious with our intuitive conception of depiction.

In order to complete the reformulation of (16) in terms of similarity between a centered picture and centered scene, we must finally introduce the corresponding notion of PROJECTION SHAPE: projection shape is just the shape of a centered object under a projection. For example, for S_j to have the property of lin-squareness is just for the lin projection of S_j to be square.[50] Using these definitions, we may now equivalently reformulate (16) as follows:

(20) P_i accurately depicts$_L$ S_j iff P_i is similar to S_j with respect to lin-shape & REF(P, S).

And illustrated thus:

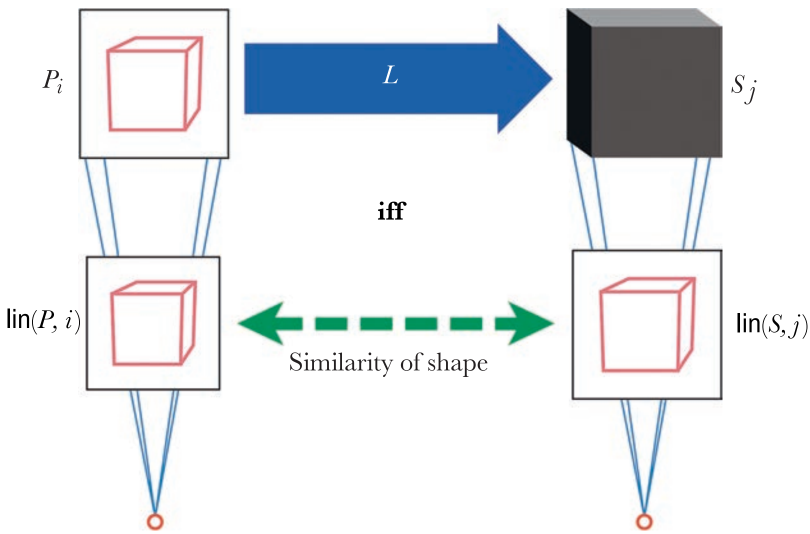

Proposition (20) clearly qualifies as a resemblance theory of L. But since R3 presupposes that depiction is a relation between a picture and scene *simpliciter*, R3 must be revised to accommodate (20), as follows.

50. More formally, where H is a shape predicate, and X_y is a variable ranging over centered pictures or scenes, the corresponding lin-shape property may be defined using lambda abstraction: $\lambda X_y.H(\text{lin}(X,y))$.

(Note that $\mathcal{F}_I$, the set of relevant properties determined by a system I, should no longer contain first-order properties of pictures and scenes but only properties under projection—but this requires no formal amendment.)

R4 *Variable Centered Resemblance with Reference Condition*
for any system of depiction I, centered picture P_i, centered scene S_j:
P_i accurately depicts$_I$ S_j iff P_i is similar to S_j with respect to $\mathcal{F}_I$
& REF(P, S).

The ability of the schema expressed by R4 to accommodate the projective definition of linear perspective constitutes a significant victory for resemblance theory. For an especially common system of depiction, resemblance theory is able to provide a precise and descriptively adequate theory of accurate depiction. Yet even as it succeeds in the case of linear perspective, I shall argue in the following section that resemblance theory cannot be sustained for all systems of depiction in general.

3. The Case against Resemblance

Resemblance theory embodies a very general position about the structure of depiction: it holds, not just of one particular system of depiction, but of *any* system of depiction, that accuracy in that system is grounded in resemblance. In this section I will contest this claim. I argue instead that there are commonly used systems of depiction whose conditions on accuracy *cannot* be characterized in terms of similarity; these systems require for accuracy that a picture *differs* from its subject according to specific rules of geometric transformation. Thus, my complaint with resemblance theory is not that there are no systems of depiction that can be grounded in resemblance, but rather, that there are some systems of depiction that cannot.[51] My ultimate diagnosis is that resemblance theory has mischaracterized the basic architecture of accurate depiction. Rather than resemblance, accurate depiction in general is grounded in the more inclusive phenomena of geometrical transformation.

51. The problems I shall raise have to do exclusively with resemblance-based accounts of pictorial shape; I adopt this narrow focus because some resemblance theorists have held the surprising view that shape is the only dimension of similarity required for depiction (Peacocke 1987; Hopkins 1998). I intend to meet these authors on their own terms. Challenges for resemblance-based accounts of pictorial color, along with suggested responses, are discussed at length by Lopes (1999), Hyman (2000), Dilworth (2005, 68–69), and Newall (2006).

I begin in section 3.1 by first describing the system of curvilinear perspective and its associated method of projection; I go on to show that the techniques that were used to analyze linear perspective in terms of similarity in the last section cannot be applied to the curvilinear case. Then in section 3.2 I present the core contribution of this essay: a general argument to the effect that no kind of similarity can ground accurate depiction in curvilinear perspective. Thus curvilinear perspective is an insurmountable counterexample to resemblance theory. I conclude by describing the general class of such counterexamples, of which there are many, and addressing potential objections.

3.1. Curvilinear Perspective

Pictures in *linear* perspective are constructed by projecting a scene from a point onto a flat picture plane. By contrast, pictures in *curvilinear* perspective are constructed by first projecting the scene from a point onto a *curved* surface, and then flattening this surface by a standard technique to yield the final image (see Loverde and Weisstein, 2012). Such pictures can be drawn by hand but are characteristically the products of "fish-eye" cameras. These cameras produce photographs that contain noticeably warped lines, especially around the periphery of the image.

Curvilinear and linear perspective have much in common. For example, in both systems, as objects increase in distance from the viewpoint, they are represented by regions of diminishing size on the picture plane. As a consequence, in both systems, parallel lines in the scene traveling away from the viewpoint appear to converge on the picture plane. The major difference between linear and curvilinear perspective lies in how they represent straight lines. For while linear perspective *does* transform shape, it also preserves line straightness. If a line is straight in the scene, it will be straight in any linear perspective projection of that scene. This is not the case for curvilinear perspective: straight lines in the scene become curved under curvilinear perspective projection (with the exception of those that pass through the center point of the picture space).

Recent studies of visual perception strongly suggest that curvilinear perspective more closely approximates natural perspective than linear perspective, though the correct explanation for this phenomenon remains a matter of dispute.[52] This fact should permanently quell any instinct to denounce curvilinear perspective as "invalid" and declare linear perspective "valid." They are simply two systems of perspective, each with its own interest and utility.

3.1.1. Curvilinear Perspective Projection

Curvilinear perspective projection proceeds in two distinct steps. The first step is exactly like linear projection, save that the surface that the scene is projected onto is curved; here I will assume that this surface is always a hemisphere. Since this surface is not the final picture plane, I shall call it the PROJECTION SURFACE.

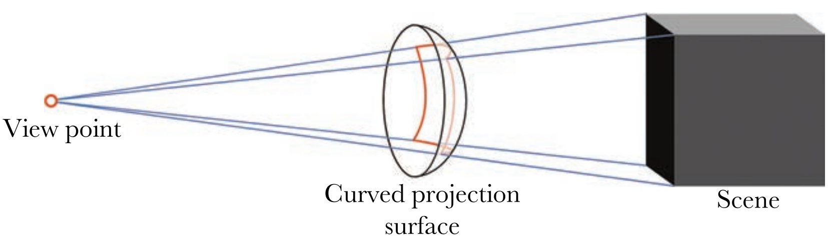

52. Among the defenders of the view that natural perspective is curvilinear are none other than the nineteenth-century perceptual scientist and physicist Hermann von Helmholtz ([1867] 1962). The view has been defended by philosophers and artists (Hansen 1973, Arnheim 1974, and Hansen and Ward 1977) and more recently in empirical work by vision scientists (Rogers and Rogers 2009).

The second step translates the lines inscribed on the curved projection surface onto the picture plane. This process is entirely determinate, with no variable parameters. The curved projection surface is oriented so that it directly faces the picture plane, as indicated below. Then, for every point inscribed on the projection surface, a line is extended perpendicularly from the picture plane to meet it. Where these lines intersect the picture plane, a correlate point is inscribed there. At the end of the process, the image on the projection surface has been transposed in flattened form onto the picture plane. The resulting image is shown at right in the figure below.

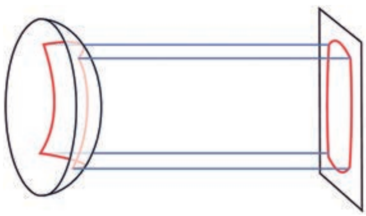

Curved projection surface   Picture plane

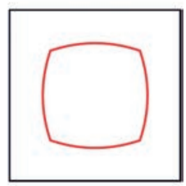

Picture plane

Just as we were able to describe the method of linear perspective projection with a simplified formula, we may do the same for curvilinear perspective. Let curv be the projection function for curvilinear perspective, and k a projective index. Then we may say that the projection function curv applied to a scene S and index k yields a picture P: $\mathsf{curv}(S,k) = P$.

3.1.2. Curvilinear Perspective as a System of Depiction

Once again, we find that the recipe for curvilinear perspective *projection* is directly related to the *system* of curvilinear perspective. To a first approximation, a picture accurately depicts a scene in the system of curvilinear perspective just in case the picture can be derived from the scene by curvilinear perspective projection. We review the evidence for this claim below.

Here we will be concerned with U, the system of monochrome line drawing in curvilinear perspective that corresponds to the projective technique described above. As before, suppose that two artists set out to draw a cube S from the vantage point used to produce the image above, relative to system U. The first picture, (21), is created in such a way that it conforms with curvilinear perspective projection. And (21) is also an accurate pictorial representation of S in the system of curvilinear perspective. But the second picture (22), due to the incompetence of the artist, deviates slightly from the intended rule in such a way that there is no projective index relative to which (22) could be derived by curvilinear perspective projection. In parallel, (22) is an *in*accurate pictorial representation of S in the system of linear perspective.

(21)

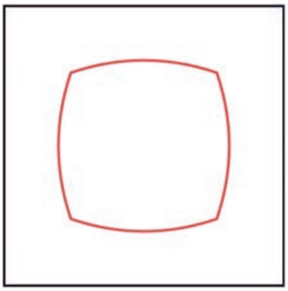

(22)

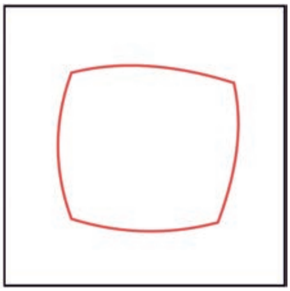

Clearly accuracy in curvilinear perspective and curvilinear projection covary closely. These examples suggest we can even state the accuracy conditions for U quite simply in terms of the projective formula $\mathsf{curv}(S,k) = P$. From here, we proceed through the same chain of reasoning that led us to adopt LP as the definition of linear perspective. I will not recapitulate every step; the same considerations of index relativity, insensitivity of accuracy judgments to size, and the need for a representation condition all apply. Together, these observations lead to the following definition of the system of curvilinear perspective U:

CP *The System of Curvilinear Perspective*
for any picture P, scene S, index k:
P accurately depicts$_U$ S relative to k iff
the shape of $P =$ the shape of $\mathsf{curv}(S, k)$ & REF(P, S).

Like its linear counterpart, CP provides a correct and exact definition of accuracy under curvilinear perspective.

3.1.3. Resemblance Theories of Curvilinear Perspective?

The crucial question for the resemblance theorist is how CP may be equivalently reformulated in terms of similarity. I'll now show that the techniques used to successfully analyze linear perspective in this way cannot be applied to curvilinear perspective.

In the last section, we saw how resemblance theorists used the method of simultaneous projection to successfully define accurate depiction in linear perspective. This proposal was inspired by the intuition that originally motivated the resemblance theory: namely, that a picture and the scene it depicts are similar in appearance. Yet this same intuitive idea does not naturally apply to the case of curvilinear perspective. Accurate curvilinear depictions do *not* resemble their subjects in appearance, at least not to the same degree exemplified by linear perspective. The "trippy" oddity of the curvilinear photograph on p. 264 makes this point obvious. This fact may be illustrated diagrammatically with respect to a curvilinear projection of a cube:

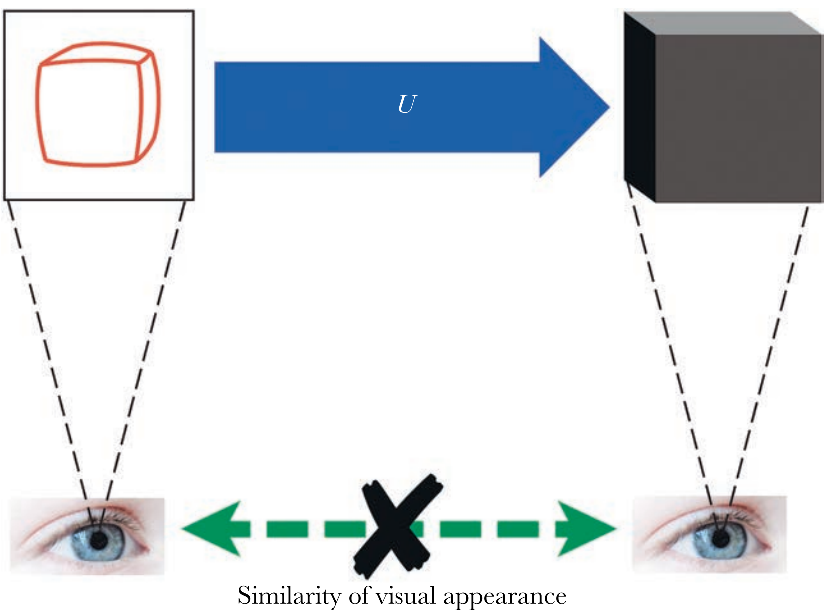

In the case of linear perspective, we went on to use the method of simultaneous projection to analyze depiction in linear perspective in terms of similarity with respect to shape-under-perspective-projection. But again, this tack fails for curvilinear perspective, as I will now explain.

Below I reproduce the simultaneous projection analysis of linear perspective and present its natural counterpart formula for U.

(20) P_i accurately depicts$_L$ S_j iff P_i is similar to S_j with respect to lin–shape & REF(P, S).

(23) P_i accurately depicts$_U$ S_j iff P_i is similar to S_j with respect to curv–shape & REF(P, S).

We saw before that (20) is not valid if the projective index is left unconstrained. Certain variations in the index such as the tilting of the picture plane had the effect of undermining resemblance between the projection of S and the projection of P. The same concern applies in the curvilinear case.[53] However, it was also shown that when the parallel planes constraint was imposed on the projective index, (20) was thereby validated. Is there a complementary restriction that can be imposed on the projective index of (23) that will render it valid? Unfortunately, for the resemblance theorist, there is not.

The crucial feature of (20) that enabled its eventual success was the existence of pictorial indices relative to which a linear perspective projection of a picture P would preserve P's shape. But this is precisely what we do not find in the curvilinear case. In the system of curvilinear perspective, straight lines are always projected onto curved lines, and curved lines are projected onto *more curved* lines—this is so no matter how the plane of projection is oriented in relation to the scene and the viewpoint. Lines in P will always become more curved under projection. As a consequence, there is no projective index such that a curvilinear projection of a picture P preserves P's shape.

To see this point more vividly, consider again the grey cube, S, and an accurate depiction of it in curvilinear perspective, P. The diagram below illustrates the simultaneous projections of S and P—where the projective indices are configured in the most natural way. As the figure makes clear, curv(P, i) and curv(S, j) do *not* have the same shape. For while the lines in curv(S, j) are curved projections from the straight edges of the cube, the lines in curv(P, i) are the even *more* curved projections from the already curved lines in P. Thus, in this case, similarity of shape under simultaneous projection does not covary with accurate depiction. So (23) is incorrect.

53. Lopes (1996, 24) notes this problem for the theory in Peacocke 1987, a cousin of (20).

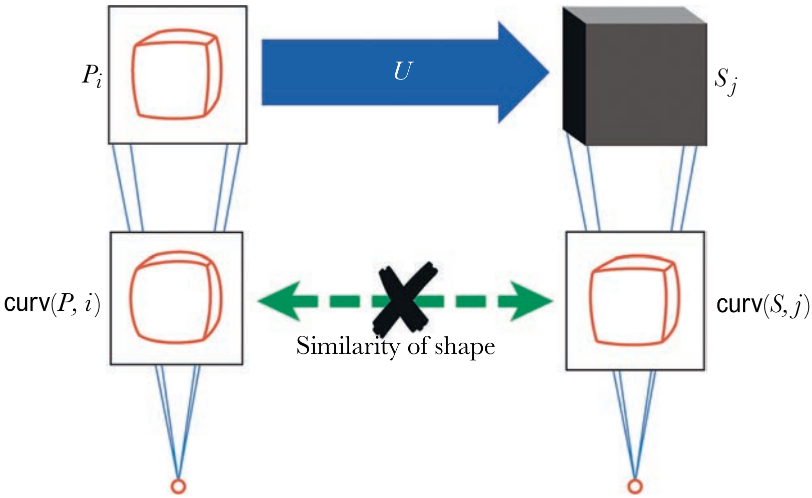

Importantly, the fact that curv(S, i) and curv(P, j) are *somewhat* similar in shape is not enough to save the proposed analysis. As we saw above, accuracy in curvilinear perspective is highly sensitive to subtle variations in pictorial shape. A slightly misdrawn curve can render an otherwise acceptable drawing inaccurate. To capture this shape-sensitivity, the conditions for depiction in U must be stringent; flexible conditions of similarity would allow more variation than U permits. If depiction in U can be grounded in similarity, it must be *complete* similarity of shape. But this is precisely what fails in the case above.

Dominic Lopes (pers. comm.) has suggested an ingenious attempt to rescue the resemblance theorist, by replacing the parallel planes constraint with a more sophisticated requirement. Perspective depiction is always defined relative to at least one parameter, the viewpoint. Why not think that, in the case of curvilinear perspective, the *degree of curvature* of the projection surface is another, contextually fixed parameter? We have so far considered only projection surfaces with a fixed, positive degree of curvature. Lopes suggests that the parallel planes constraint be augmented with an additional requirement that the degree of curvature for the projection of P always be zero—that is, perfectly flat. This solution, while technically viable, comes at too high a philosophical cost. It effectively requires that the comparison of similarity occurs between a *linear* projection of the picture and a *curvilinear* projection of the scene. The result is an instance of resemblance theory in name only; in fact, it requires asymmetrical transformations of its relata. While it is true that the original parallel planes constraint also imposes asymmetrical requirements on the picture and scene, it is at least possible for both objects to satisfy the constraint at the same time (if the scene was also a plane). This is not the case of the zero-curvature constraint. It is impossible to apply a zero-curvature projection to both relata while computing real curvilinear projection—instead, the result is always an instance of linear projection. The proposed solution is *necessarily* asymmetrical, thus a radical departure from the guiding principle of the resemblance analysis.

In conclusion, the strategy of analyzing depiction in terms of similarity of simultaneously projected shape, which worked so well for linear perspective, cannot succeed for curvilinear perspective. The key feature of curvilinear perspective that undermines the analysis is that there is no way of preserving the shape properties of a flat surface under curvilinear projection. In U, a picture of a picture can never have the same shape as the original. Depiction in U is inevitably transformative. But perhaps there is some other strategy for analyzing curvilinear perspective in terms of similarity? There is not. In the next section, I generalize the negative results of the preceding discussion: there is *no* acceptable resemblance analysis that provides sufficient conditions for accurate depiction in curvilinear perspective.

3.2. Against Resemblance Theories of Curvilinear Perspective

In this section, I will argue that the system of curvilinear perspective gives rise to counterexamples for any resemblance theory of depiction. These are examples in which there is resemblance of the required sort *and* pictorial reference but not accurate depiction. Hence they show that the conditions described by the best resemblance theory are insufficient for depiction.

3.2.1. Argument Outline

Let us begin by reviewing the definition of R4:

> R4 *Variable Centered Resemblance with Reference Condition*
> for any system of depiction I, centered picture P_i, centered scene S_j:
> P_i accurately depicts$_I$ S_j iff P_i is similar to S_j with respect to $\mathcal{F}_I$
> & REF (P, S).

I will argue that within U, the system of curvilinear perspective, there is a centered picture P_i and scene S_j such that:

Premise 1. It is not the case that P_i accurately depicts$_U$ S_j.
Premise 2. The reference condition holds: REF(P, S).
Premise 3. P_i is similar to S_j with respect to $\mathcal{F}_U$.

Thus there is a case in which there is no accurate depiction (premise 1), yet the conditions sufficient for accurate depiction according to R4 are met (premises 2 and 3). So R4 is wrong. Diagrammatically, the argument proceeds as follows (suppressing projective indices for readability). The accuracy conditions for U, according to R4, are shown at left. The counterexample is described at right.

According to R4, for any P, S:

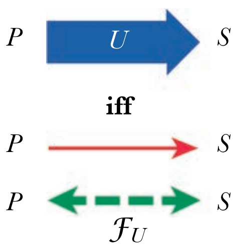

Yet there is a P, S such that:

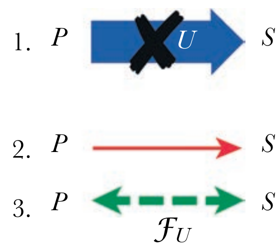

The premises shown above are the central premises of my argument. Together they contradict R4. I'll argue for each in turn.

3.2.2. The Case

The argument revolves around an unusual but certainly possible episode of drawing. Suppose that, having mastered the art of hand drawing in curvilinear perspective, I have now taken up the task of learning to hand draw in curvilinear perspective while *blindfolded*. For practice, I select various objects from my studio, gaze at them intently for several minutes, then don the blindfold and begin drawing. Unfortunately, my skill at this task is still nascent, and most of my attempts are highly inaccurate. On one occasion, I select as my subject an old drawing S. I look carefully at S and select a vantage point such that the imagined projection surface is centered directly above S. (Let j be the projective index so specified for the scene, and i the projective index implicitly specified for the new picture, such that $i = j$.) Then I apply the blindfold and begin to draw. When I am done, I discover to my surprise that the drawing I have produced, P, is qualitatively indistinguishable from S:

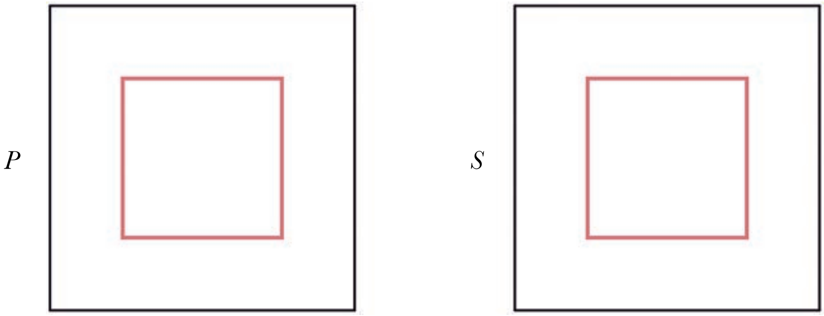

Of course, it is highly unlikely that, after donning a blindfold, I could attempt to draw S and produce P. A scribble seems a more probable result of this process than a regular shape, much less one qualitatively indistinguishable from the subject. But if a scribble might result, then so might the regular shape. The unlikeliness of the scenario described is no mark against its possibility. With that worry at bay, I will now argue that the centered picture P_i does not accurately depict$_U$ the centered scene S_j (premise 1); that P depictively refers to S (premise 2); and further, that P_i resembles S_j in any respect that the resemblance theorist might reasonably invoke (premise 3). I conclude that the resemblance theory of depiction is false.

3.2.3. Premise 1

Our first task is to show that P_i is not an accurate depiction of S_j in the system of curvilinear perspective. To begin, recall that curvilinear projection works by first projecting the scene onto a curved projection surface and then transposing this curved surface onto the flat picture plane. Let us consider the result of this process when the scene in question is the flat surface S. First S is projected onto a curved projection surface:

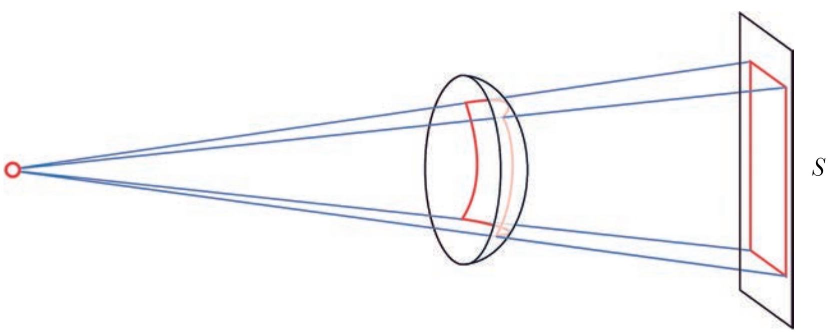

This surface is then transposed onto the picture plane. The result, call it R, is the product of applying the curvilinear projection function curv to S relative to viewpoint j. That is, $\mathsf{curv}(S, j) = R$.

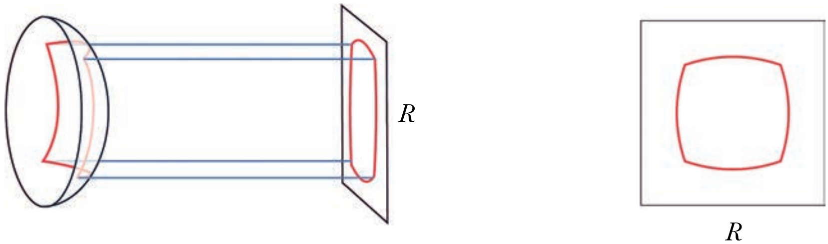

Since R is the result of taking the curvilinear projection of S relative to a certain projective index, by CP it follows that R accurately depicts$_U$ S relative to the same index:

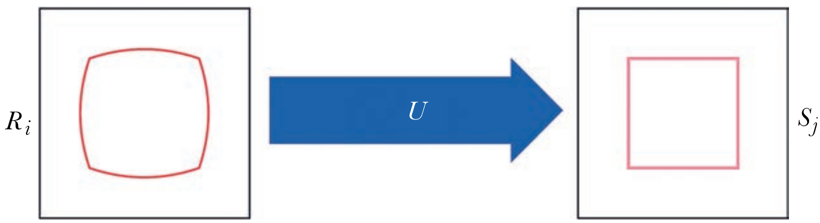

But now recall how sensitive CP is to picture shape. Minor changes in curvature to the lines in R will result in inaccuracy. So long as we keep the projective index fixed, then R (or a scale copy) is the *only* picture that accurately depicts$_U$ S. But R and P have very different shapes, as evidenced below:

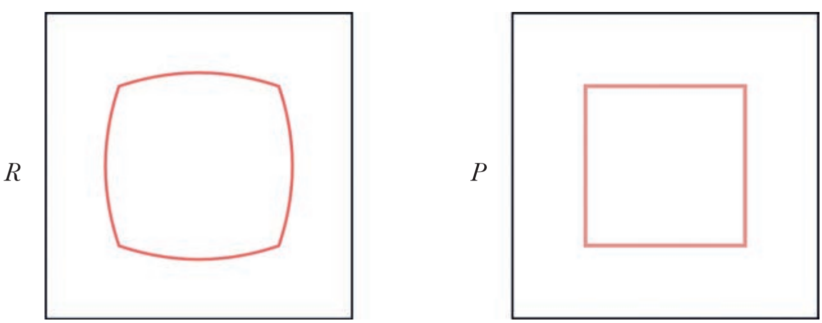

Since the projective index that was used to generate R is the same one used in the evaluation of P, it follows immediately that P_i does not accurately depict$_U$ S_j. The point is further dramatized if we consider the flat surface that P_i *would* accurately depict in curvilinear perspective:

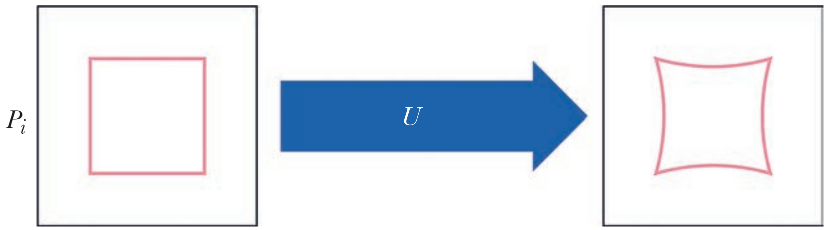

Quite obviously, the flat surface that P would accurately depict differs dramatically from S; so P does not accurately depict S.

In general, a signature feature of curvilinear projection is that it maps straight lines in the scene onto curved lines in the picture plane. Thus, no matter the projective index, the straight lines in S will be mapped to curved lines in the picture plane. Yet the lines in P are not curved. Thus it cannot be a curvilinear projection of S for *any* index. It follows that P_i does not accurately depict$_U$ S_j.

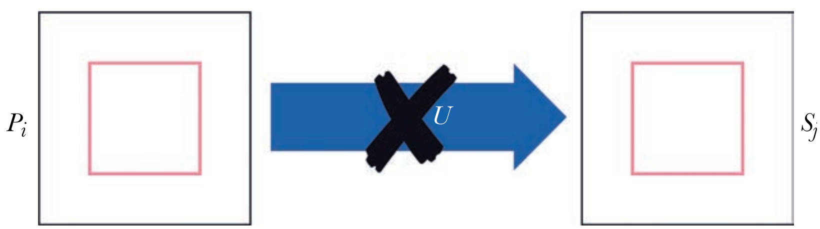

3.2.4. Premise 2

Next, we wish to establish that the reference condition is satisfied with respect to P and S—that P bears the relation of pictorial reference to S. In our earlier discussion of pictorial reference, we used the scribbled drawing of Obama made while blindfolded as a paradigmatic example of pictorial reference without accurate depiction. If we were correct in identifying that picture as a case of pictorial reference to Obama, then surely P also bears the relation of pictorial reference to S. Whether or not P accurately depicts S, it was drawn with the honest intention of representing S, and there is no reason to think that the intended reference failed. P is a picture *of* S, that is, a pictorial reference to S, thus the reference conditions are satisfied with respect to P and S.

P  S

3.2.5. Premise 3

Finally, we will now show that P_i and S_j are similar with respect to $\mathcal{F}_U$. By stipulation, of course, P and S are qualitatively indiscernible. So it is very natural to conclude that they are similar with respect to $\mathcal{F}_U$. My task here is to substantiate this suggestion without controversial assumptions about how a resemblance theorist might wish to specify $\mathcal{F}_U$.

I make just one, quite plausible assumption about the resemblance condition of R4: it is sensitive only to the qualitative features of the scene and picture, rather than any singular properties of each.[54] That is, the resemblance condition is sensitive only to features of the scene and picture such as size, shape, color, and so on, but not to features like being one of Granny's favorite objects, or containing a particular particle, or being a particular picture. While this *is* a substantive constraint on R4, it is entirely of a piece with the spirit of resemblance theory, and I know of no actual resemblance theory that violates it.[55]

The intuitive idea that the resemblance condition is sensitive only to qualitative properties of objects may be expressed more formally using a substitution principle. Objects that are qualitatively indistinguishable

54. I am not sure how to define qualitative as opposed to singular properties, but I am confident there is a real distinction to be drawn. Qualitative properties are properties such as size and shape; singular properties are properties such as *being that desk*. See Adams 1979 for discussion.

55. Many resemblance theories make explicit commitments that seem to entail this assumption. For example, Files (1996, 403) stipulates that "the resemblance theory of pictorial content adverts, obviously enough, to a sharing of *appearance* properties." If "appearance" properties and qualitative properties are not the same, my argument is undiminished by substituting the former for the latter.

may be freely substituted into the resemblance condition with no change in the satisfaction of the condition. This suggestion in turn may be formulated explicitly as a principle of "qualitative indifference":

> For any system of depiction I, and any three objects A, B, and C :
> if A and B are qualitatively indistinguishable,
> and A and C are similar with respect to $\mathcal{F}_I$,
> then B and C are similar with respect to $\mathcal{F}_I$.

To bring this principle to bear on our argument, it remains to extend it to *centered* pictures and scenes. There are a variety of ways to achieve this; here I offer one straightforward approach. Once again, the intuitive idea to capture is that the resemblance condition is sensitive only to qualitative properties of its relata. But we should allow that differences in projective index could result in differences that are relevant to the resemblance condition. On the other hand, so long as the projective indices of two relata are identical, then the resemblance condition should be determined entirely by the qualitative properties of the uncentered relata.

For example, given two centered objects A_i and B_j, if $i = j$, then whether or not A_i is similar in the relevant respects to B_j should depend entirely on the qualitative properties of A and B respectively. In terms of substitution, if A and B are qualitatively indistinguishable *and* $i = j$, then we should be able to freely substitute A_i for B_j into the resemblance condition with no change in satisfaction. Explicitly then:

> ***Qualitative Indifference*** for any system of depiction I, and any three centered objects A_i, B_j, and C_k:
> if $i = j$ and A and B are qualitatively indistinguishable and A_i and C_k are similar with respect to $\mathcal{F}_U$, then B_j and C_k are similar with respect to $\mathcal{F}_U$

We will now wield this principle to achieve our conclusion. To begin, recall that similarity with respect to $\mathcal{F}_U$ is a similarity relation, therefore reflexive. In the case of the centered scene S_j, it follows that:

(24) S_j is similar to S_j with respect to $\mathcal{F}_U$.

By stipulation, P is qualitatively indistinguishable from S. Also by stipulation, the projective indices for P and S are identical because P was drawn from S using the same vantage point that P itself was intended to be viewed from. This does not reflect a sophisticated choice of vantage point; it is just the very natural projective index that situates the viewpoint perpendicular to the center of the flat surface S and perpendicular to the center of the picture P.[56] Thus, so far as *Qualitative Indifference* is concerned, P_i and S_j are indiscernible, so intersubstitutable. Then by *Qualitative Indifference* and (24), premise 3 follows:

(25) P_i is similar to S_j with respect to $\mathcal{F}_U$.

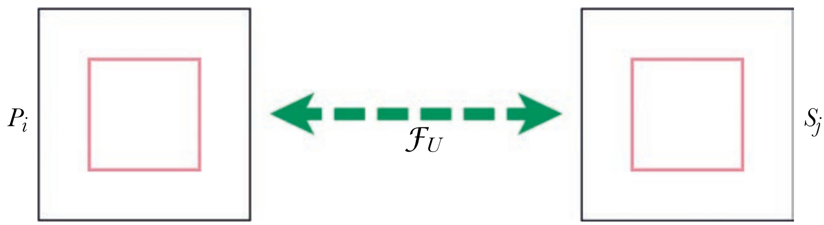

Thus, so long as theories of resemblance conform with *Qualitative Indifference*, we may derive our final premise with no further assumptions about the particular account of similarity at work.

3.2.6. Conclusion

By combining premises 2 and 3, we see that P_i and S_j satisfy both the resemblance and reference conditions of R4. By R4, it follows that P_i accurately depicts$_U$ S_j. But careful consideration of the system of curvilinear perspective justified premise 1, that in fact, P_i does *not* accurately depict$_U$ S_j.

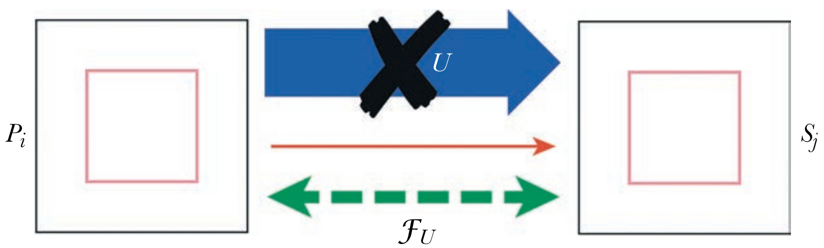

56. I can see no reason why such a stipulation could not be sustained. Except this: on some formal implementations of U, it might be impossible that a single projective index could be applied to distinct locations in space. But the identity condition in *Qualitative Indifference* could be easily and minimally modified to accommodate this requirement. For example, it could be replaced by requiring relational similarity between indices.

Thus, contrary to R4, no degree of similarity is sufficient to guarantee accurate depiction in curvilinear perspective. As I have shown, this result holds for any credible definition of similarity. So R4 is false.

What does this argument reveal about R4's shortcomings? In the case above, the essential problem for R4 was that *P* and *S* were *too* similar. As the discussion of premise 1 made clear, *P* could accurately depict only scenes that differed from it in certain ways; and *S* could be accurately depicted only by pictures that differed in opposite ways. Yet *P* and *S* were qualitatively indiscernible; this meant that they satisfied the similarity requirement but were insufficiently different to be related by curvilinear projection.

In general, accurate depiction in the system of curvilinear perspective depends on a certain degree of *difference*; the point is not merely that curvilinear perspective *tolerates* a certain amount of difference but rather that it *requires* it. Yet it is the nature of a similarity condition that it puts no lower bound on similarity; it makes no allowance for such conditions of nonsimilarity. So even as accuracy in curvilinear perspective requires difference, the similarity condition in R4 requires none. These requirements are obviously not equivalent. The case above was simply constructed to reveal how they may come apart.

The argument just advanced trades on the insufficiency of any degree of resemblance to ground accurate depiction. But this, recall, was also the moral of Goodman's famous objection to resemblance theories. Yet the similarities between these conclusions are similarities in form alone.[57] Goodman's argument reveals that some pragmatic preconditions *beyond* a resemblance-based criteria of "fit" are required to ground resemblance. But Goodman never explicitly considers whether resemblance is a suitable theory of such "fit." By contrast, I have taken care to focus on cases where pictorial reference is already guaranteed—so Goodman's objection is stayed—showing instead that resemblance has no further role to play in grounding accurate depiction. The lessons of Goodman's discussion have been largely incorporated by contemporary resemblance theorists. The conclusions of the present argument, I contend, cannot be.[58]

57. Thanks to an anonymous reviewer for pressing this issue.

58. An additional note on the logic of my argument: recall that I began the discussion of curvilinear perspective by showing that it could not be analyzed by the method of simultaneous projection. Specifically, I showed that similarity in shape under simultaneous projection could not be a necessary condition on depiction. Now I have argued

3.3. Against Resemblance

Resemblance theory is an account of the general structure of depiction, for it holds that *any* system of depiction is grounded in resemblance. We have seen that this claim fails; even at its most sophisticated, resemblance theory cannot model the accuracy conditions of inherently differential systems of depiction like curvilinear perspective. Thus resemblance theory does not describe the universal architecture of a theory of depiction.

But some may feel this conclusion is too hasty. Might it still be the case that resemblance theory is *approximately* correct? Perhaps resemblance theory correctly describes some "core" class of systems of depiction, while deviations from resemblance correspond to peripheral cases of pictorial representation.[59] According to this view, linear perspective, which can be analyzed in terms of resemblance, is an example of a core system of depiction, while curvilinear perspective is not.

Yet it is difficult to see what could justify this ordering of the systems of depiction. The fact that linear perspective is more common than curvilinear perspective seems to have more to do with practical considerations—the ease of creating such images with a pencil and straightedge, for example—than any fundamental features of pictorial representation. Further, recent empirical research appears to confirm the long-held view that the structure of human visual perception corresponds more nearly to curvilinear than linear projection (Rogers and Rogers 2009). (The research shows that curved lines, viewed at the periphery of vision, appear straight.) If this is so, then curvilinear perspective should arguably be classified as a more central case of pictorial representation than linear perspective. That there is room for empirical debate on this question

---

that similarity of any kind is not sufficient. What is the relationship between these two conclusions? It should first of all be emphasized that the insufficiency argument constructed here applies straightforwardly to the simultaneous projection analysis (though its domain of application is also more general). So the simultaneous projection analysis gives conditions that are both unnecessary and insufficient. In general, the conclusions secured here are best interpreted as showing, generally, that the conditions that ground depiction cannot be conditions of resemblance. Resemblance conditions that are plausibly sufficient tend to turn out not to be necessary; and none turn out, in the end, to be sufficient either.

59. This is the position of Kulvicki (2006). While Kulvicki does not explicitly allow that his is a resemblance theory, his only substantive requirement on depiction is that it preserve projective invariants, hence that pictures and scenes be similar with respect to projective invariants.

indicates the poverty of this line of defense as a philosophical justification for resemblance theory.

But the problem with the objection is even deeper. I have chosen curvilinear perspective as my counterexample to resemblance theory because the system is easily defined, relatively common, and involves vivid failures of resemblance in shape. But it is by no means the only system that resists analysis in terms of resemblance. Using the rubric of the argument presented here, readers may go on to identify other counterexamples for themselves: cases in which an accurate picture of a scene necessarily differs from its subject.[60] Examples include the conventions of scale that accompany the average road map, as well as systems of contour drawing, and the variety of map projections. In the domain of color and tone representation, many systems are recalcitrant to similarity analysis. Such systems involve: overall shifts in lightness, darkness, and contrast; false color and negative color; and the amplification of hue—as in technicolor—or suppression of it—as in the brown-tinted color of nineteenth-century European painting. Of the last, Gombrich (1960, 46–47) aptly wrote, "It is a transposition, not a copy." In general, a relation that requires difference in some dimension cannot be analyzed in terms of one that allows an arbitrarily high degree of similarity. Yet such a similarity condition lies at the heart of any resemblance theory. Thus resemblance theory inevitably fails.

So while the skeptic alleges that I have identified an odd outlier that does not submit to analysis in terms of resemblance, nearly the opposite is true: only a small and delicately delineated set of systems of depiction can be analyzed in terms of resemblance. Minor changes to the rules governing these systems—for example, in terms of the linearity of perspective—render these systems incompatible with a resemblance-based analysis.

A different kind of objection holds that I have misconstrued the ambitions of resemblance theory. On this view, resemblance theory never seriously aspired (or never should have aspired) to identify both necessary and sufficient conditions on depiction. Its only aim was to identify necessary conditions: for any picture, scene, and system of depiction, *if* the picture accurately depicts the scene in that system, *then* the picture is similar to the scene in certain respects. Such a theory, it must be admitted, gives up the central aspiration of supplying constitutive conditions on

60. Dilworth (2005, 68–69), for one, makes precisely this objection against resemblance treatments of pictorial color, albeit rather briefly.

depiction, of explaining what depiction fundamentally *is*. Instead, this attenuated version of resemblance theory is merely descriptive.

Yet there is a danger that the claim of necessary similarity is trivial. The problem is that similarity is cheap. Necessarily, any two objects are similar to each other in infinitely many respects (Goodman 1972). For example, the number three and I are similar with respect to not being the number two, the number four, and so on. For their claim to be substantive, resemblance theorists must specify the respects in question. But it turns out to be difficult, perhaps prohibitively difficult, to specify exactly what these similarities are. We have already seen that similarity with respect to projection shape cannot be a necessary condition on accurate depiction. Nor even similarity with respect to such benign features as betweenness relations: the projective counterparts of points that lie along a line in the picture plane do not necessarily lie along a line in the scene. It is probably impossible to prove that there are no substantive necessary resemblance conditions on accurate depiction because it is probably impossible to rigorously define a "substantive condition." But I have yet to encounter a proposed similarity condition that adequately reflects the impressive variety of systems of depiction.

A distinct and final objection holds that my definition of resemblance, in terms of the philosophical concept of similarity, is too narrow. There are other, more elastic notions of resemblance, such as the mathematical concept of *isomorphism*, which may be impervious to the arguments introduced above. But this objection misunderstands the breadth and flexibility of my chosen concept of similarity, defined only as sharing of properties. The properties in question need not be first-order, qualitative properties; they may be properties cast at any level of abstraction. Indeed, isomorphism is just a species of abstract similarity—similarity with respect to relational structure. It is therefore high time we looked beyond resemblance.

4. Beyond Resemblance

I have argued that resemblance theory fails as a general account of depiction because there are actual systems of depiction that require of accurate images that they differ from their subjects in systematic ways. At the same time, it is clear that certain dimensions of resemblance *do* matter for accurate depiction in many systems. In normal color photography, for example, similarity of apparent surface color is a necessary condition on accuracy. Such considerations suggest that, in general, accurate depiction cannot be characterized wholly in terms of difference, or wholly in terms of similarity, but should instead be defined by the broader notion of *transformation* that incorporates both.

Still, resemblance theory and a transformational view of depiction, though distinct, are closely aligned. Both hold that depiction is grounded in a certain kind of "fit" between the structural features of a picture and its subject; both reject the view that depiction is based on arbitrary associations between pictures and scenes, or on the spontaneous perceptual judgments of individual viewers. Further, if accurate depiction is grounded not in resemblance but in systematic transformation, those resemblances that have seemed so powerful to so many theorists are not ruled out. Instead they are understood as transformational invariants, properties of a scene that survive projection onto the picture plane. The transformational account is a reorientation and extension of resemblance theory, affording greater flexibility and freedom from a limiting interpretation of resemblance theory's motivating intuitions. It recognizes that there is more to accurate depiction than invariants.

Once we appreciate how systems of depiction are precisely defined by geometrical algorithms such as those of perspective projection, the view that depiction is grounded in projective transformation becomes the natural one. Rather than trying to forcibly recast these transformational rules in terms of resemblance, we should instead take them at face value. The correct definition of linear perspective is simply the one given by linear perspective projection. The many systems of projection correspond to the many systems of depiction. And different systems of depiction are grounded in different kinds of geometrical projection. (Of course, these are projections of more than merely spatial geometry; many define complex manipulations of color.) These conclusions should not be surprising. Despite the apparent efficiency of resemblance-based interpretation, interpretive rules based on geometrical projection make more natural use of the innate computational machinery involved in visual, perspectival perception.[61] To recover information from a picture, an audience member must work backward from picture to scene, inferring what sort of scene the image must have been projected from, much as the visual system recovers information from light projected to the eye.

In this essay, I have purposefully confined my discussion to the analysis of representation as it arises in publicly available pictures. Still,

61. I develop these conclusions in detail in Greenberg 2012.

I believe the same basic considerations extend to other domains where resemblance or isomorphism views of representation have taken hold, especially in the theory of scientific representation and in cognitive science. Insofar as researchers in these areas take pictures as a paradigm of representation and the resemblance theory of depiction as the correct analysis, I hope at least to have raised doubts, and indicated a different way forward. Detailed application of the present considerations to these areas must await future development.

References

Abell, C. 2007. "Pictorial Realism." *Australasian Journal of Philosophy* 85, no. 1: 1–17.

———. 2009. "Canny Resemblance." *Philosophical Review* 118, no. 2: 183–223.

Adams, R. M. 1979. "Primitive Thisness and Primitive Identity." *Journal of Philosophy* 76, no. 1: 5–26.

Alberti, L. B. 1991 [1435]. *On Painting*, ed. C. Grayson. London: Penguin Classics.

Arnheim, R. 1974. "Hansen's Curvilinear Perspective." *Journal of Aesthetics and Art Criticism* 32, no. 3: 424–44.

Beardsley, M. C. 1958. *Aesthetics: Problems in the Philosophy of Criticism*. San Diego: Harcourt, Brace, and Company.

Block, N. 1982. *Imagery*. Cambridge, MA: MIT Press.

Blumson, B. 2007. "Resemblance and Representation." PhD diss., Australian National University.

———. 2009a. "Defining Depiction." *British Journal of Aesthetics* 49, no. 2: 143–57.

———. 2009b. "Pictures, Perspective, and Possibility." *Philosophical Studies* 149, no. 2: 135–51.

Burge, T. 2010. *Origins of Objectivity*. Oxford: Oxford University Press.

Cohen, J., and C. Callender. 2006. "There Is No Special Problem about Scientific Representation." *Theoria* 21, no. 55: 67–85.

Costello, D., and D. M. Phillips. 2009. "Automatism, Causality and Realism: Foundational Problems in the Philosophy of Photography." *Philosophy Compass* 4, no. 1: 1–21.

Cummins, R. 1996. *Representations, Targets, and Attitudes*. Cambridge, MA: MIT Press.

Da Vinci, L. 1970 [c. 1500]. *The Notebooks of Leonardo da Vinci*, ed. and trans. J. P. Richter. New York: Dover Publications.

DeCarlo, D., et al. 2003. "Suggestive Contours for Conveying Shape." *ACM Transactions on Graphics* 22, no. 3: 848–855 (TOG).

Descartes, R. 2001 [1637]. *Discourse on Method, Optics, Geometry, and Meteorology*, trans. P. J. Olscamp. Indianapolis: Hackett.

Dilworth, J. 2005. "Resemblance, Restriction, and Content-Bearing Features." *Journal of Aesthetics and Art Criticism* 63, no. 1: 67–70.

Eco, U. 1979. *A Theory of Semiotics*. Bloomington: Indiana University Press.

Files, C. 1996. "Goodman's Rejection of Resemblance." *British Journal of Aesthetics* 36, no. 4: 398–412.

French, S. 2003. "A Model-Theoretic Account of Representation (or, I Don't Know Much about Art—but I Know It Involves Isomorphism)." *Philosophy of Science* 70, no. 5: 1472–83.

Gibson, J. J. 1960. "Pictures, Perspective, and Perception." *Daedalus* 89, no. 1: 216–27.

———. 1971. "The Information Available in Pictures." *Leonardo* 4, no. 1: 27–35.

Giere, R. 2004. "How Models Are Used to Represent Reality." *Philosophy of Science* 71, no. 5: 742–52.

Gombrich, E. H. 1960. *Art and Illusion: A Study in the Psychology of Pictorial Representation*. London: Phaidon.

Goodman, N. 1968. *Languages of Art: An Approach to a Theory of Symbols*. Indianapolis: Bobbs-Merrill.

———. 1972. "Seven Strictures on Similarity." In *Problems and Projects*, 437–47. Indianapolis: Bobbs-Merrill.

Greenberg, G. 2012. "Pictorial Semantics." Unpublished manuscript. gjgreenberg.bol.ucla.edu/.

Hagen, M. A. 1986. *Varieties of Realism: Geometries of Representational Art*. Cambridge: Cambridge University Press.

Hansen, R. 1973. "This Curving World: Hyperbolic Linear Perspective." *Journal of Aesthetics and Art Criticism* 32, no. 2: 147–61.

Hansen, R., and J. L. Ward. 1977. "Perception of Pictorial Space in Perspective Pictures." *Leonardo* 10, no. 3: 261–62.

Helmholtz, H. 1962 [1867]. *Treatise on Physiological Optics*, ed. J. P. C. Southall. Vol. 3. Mineola, NY: Dover.

Hopkins, R. 1998. *Picture, Image and Experience*. Cambridge: Cambridge University Press.

Hyman, J. 2000. "Pictorial Art and Visual Experience." *British Journal of Aesthetics* 40, no. 1: 21–45.

———. 2006. *The Objective Eye: Color, Form, and Reality in the Theory of Art*. Chicago: University of Chicago Press.

Jakobson, R. 1987 [1921]. "On Realism in Art." In *Language in Literature*, 19–27. Cambridge, MA: Harvard University Press.

Kaplan, D. 1968. "Quantifying In." *Synthese* 19, no. 1: 178–214.

Knight, H. 1930. "Aesthetic Experience in Pictorial Art." *Monist* 40, no. 1: 74–83.

Kulvicki, J. 2006. *On Images: Their Structure and Content*. Oxford: Clarendon.

Lewis, D. 1979. "Attitudes *De Dicto* and *De Se*." *Philosophical Review* 88, no. 4: 513–43.

———. 1986. *On the Plurality of Worlds*. Oxford: Basil Blackwell.

Lopes, D. 1995. "Pictorial Realism." *Journal of Aesthetics and Art Criticism* 53, no. 3: 277–85.

———. 1996. *Understanding Pictures*. Oxford: Oxford University Press.

———. 1999. "Pictorial Color: Aesthetics and Cognitive Science." *Philosophical Psychology* 12, no. 4: 415–28.

Loverde, T., and E. W. Weisstein. n.d. "Fisheye Perspective." In *MathWorld–A Wolfram Web Resource*, mathworld.wolfram.com/FisheyePerspective.html (accessed November 4, 2012).

Lubbock, T. 2006. "Philippe de Champaigne: Vanité (1660s)." *Independent*, www.independent.co.uk/arts-entertainment/art/greatworks/de-champaigne-philippe-vanit-1660s-798027.html.

Malinas, G. 1991. "A Semantics for Pictures." *Canadian Journal of Philosophy* 21, no. 3: 275–98.

Manns, J. W. 1971. "Representation, Relativism and Resemblance." *British Journal of Aesthetics* 11, no. 3: 281–7.

Morris, C. W. 1946. *Signs, Language and Behavior*. Upper Saddle River, NJ: Prentice Hall.

Neander, K. 1987. "Pictorial Representation: A Matter of Resemblance." *British Journal of Aesthetics* 27, no. 3: 213–26.

Newall, M. 2003. "A Restriction for Pictures and Some Consequences for a Theory of Depiction." *Journal of Aesthetics and Art Criticism* 61, no. 4: 381–94.

———. 2006. "Pictures, Colour and Resemblance." *Philosophical Quarterly* 56, no. 225: 587–95.

Novitz, D. 1975. "Picturing." *Journal of Aesthetics and Art Criticism* 34, no. 2: 145–55.

Peacocke, C. 1987. "Depiction." *Philosophical Review* 96, no. 3: 383–410.

Peirce, C. S. 1906. "Prolegomena to an Apology for Pragmaticism." *Monist* 16, no. 4: 492–546.

Putnam, H. 1981. *Reason, Truth, and History*. Cambridge: Cambridge University Press.

Pylyshyn, Z. 2006. *Seeing and Visualizing: It's Not What You Think*. Cambridge, MA: MIT Press.

Rogers, B., and C. Rogers. 2009. "Visual Globes, Celestial Spheres, and the Perception of Straight and Parallel Lines." *Perception* 38, no. 9: 1295–312.

Sedgwick, H. A. 1980. "The Geometry of Spatial Layout in Pictorial Representation." In *The Perception of Pictures*, vol. 1, 33–90. London: Academic Press.

Suarez, M. 2003. "Scientific Representation: Against Similarity and Isomorphism." *International Studies in the Philosophy of Science* 17, no. 3: 225–44.

Van Fraassen, B. 2008. *Scientific Representation: Paradoxes of Perspective*. Oxford: Clarendon.

Walton, K. L. 1973. “Pictures and Make-Believe.” *Philosophical Review* 82, no. 3: 283–319.

———. 1990. *Mimesis as Make-Believe: On the Foundations of the Representational Arts*. Cambridge, MA: Harvard University Press.

Willats, J. 1997. *Art and Representation: New Principles in the Analysis of Pictures*. Princeton: Princeton University Press.

Wittgenstein, L. 1997 [1921]. *Tractatus Logico-philosophicus*, trans. D. Pears and B. McGuinness. London: Routledge.