The Semiotic Spectrum

by Gabriel J. Greenberg

A dissertation written under the direction of
Jeff King (*co-chair*)
Jason Stanley (*co-chair*)
Peter Kivy
Matthew Stone
Dom Lopes (*external*)

Abstract

Because humans cannot know one another's minds directly, every form of communication is a solution to the same basic problem: how can privately held information be made publicly accessible through manipulations of the physical environment? Language is by far the best studied response to this challenge. But there are a diversity of *non*-linguistic strategies for representation with external signs as well, from facial expressions and fog horns to chronological graphs and architectural renderings. The general thesis of this dissertation is that there is an impressively wide spectrum of conventional systems of representation, corresponding to the many ways that the problem of communication can be solved, and that these systems can be described and explained using the tools of contemporary mathematical semantics.

As a partial corrective to the countervailing norm, this work concentrates on the class of systems arguably most different from language— those governing the interpretation of *pictorial images*. Such representations dominate practical communication: witness the proliferation of maps, road signs, newspaper photographs, scientific illustrations, television shows, engineering drawings, and even the fleeting imagery of manual gesture. I argue that *systems of depiction* and *languages* embody a parallel technologies of communication. Both are based on *semantics*: systematic and conventional mappings from signs to representational content. But I also provide evidence that these semantics are profoundly divergent. Whereas the semantics of languages are based on arbitrary associations of signs and denotations, the semantics of systems of depiction are based on rules of *geometrical* transformation. Drawing on recent research in computer graphics and computational vision, I go on to develop a precise theory of pictorial semantics. This in turn facilitates a detailed comparison of iconic, image-based representation, and symbolic, language-based representation.

A consequence of these conclusions is that the traditional, language-centric conception of semantics must be overhauled to allow for a more general semantic theory, one which countenances the wide variety of interpretive mechanisms actually at work in human communication.

Acknowledgements

The present work is the natural expression of twin life-long passions for visual representation and for philosophy. The first I owe to my parents: to my mother, Linda Rubinstein, who introduced me to the joys of drawing and story-telling; and to my father, Chip Greenberg, who gifted me his infectious curiosity about the mechanics of representation and the mysteries of codes. My first philosophical guides were Sam Amidon, Ed Bartlett, Rick Campman, Zeke Hecker, and Susan Klein, each of whom passed many patient hours with me teasing out the elusive nexus of mysticism, philosophy, and art. In college, I was expertly led into the illumination of analytic thought by Kati Balog, Troy Cross, Brian Scholl, and Michael Della Rocca. And my most lasting drawing lessons were those of Robert Reed and Joe Scanlan.

As a graduate student at Rutgers, Jason Stanley brought me into the thriving semantics and philosophy of language community, led in different capacities by Ernie Lepore, Jason, and Jeff King. This group of faculty and students has cultivated a heady blend of concerns which range freely among the fields philosophy, linguistics, computer science, and cognitive science. The methodology and sensibility of this dissertation is the direct inheritance of this group.

Of the philosophical morals I have learned over the past years, the most fundamental I have learned from my peers. They have included Tendayi Achiume, Luvell Anderson, Josh Armstrong, Dawn Chan, Sam Cumming, Pavel Davydov, Janelle Derstine, Richard Dub, Edinah Gnang, Rohit Gupta, Alex Jackson, Michael Johnson, Ben Levinstein, Karen Lewis, Kelby Mason, Ricardo Mena, Zachary Miller, Lisa Miracchi, Alex Morgan, Sarah Murray, Jenny Nado, Carlotta Pavese, Jessica Rett, Kevin Sanik, Adam Sennett, and Will Starr. A cadre of inspiring professors in the departments of Philosophy, Linguistics, and Computer Science at Rutgers have led by the example of their excellence as scholars and teachers. They have included Maria Bittner, Doug DeCarlo, Ernie Lepore, Jerry Fodor, Thony Gillies, John Hawthorne, Jeff King, Peter Kivy, Barry Loewer, Tim Maudlin, Roger Schwarzschild, Chung-chieh Shan, Ted Sider, Steve Stich, Jason Stanley, and Matthew Stone.

The germ of the present project was born in Roger Schwarzschild's semantics seminar, who

Contents

New York City, 2011. *Advisory graphic for a subway emergency intercom system.*

Chapter 0

General Semantics

The theme of this dissertation, familiar enough to students of semiotics, occurred to me one warm summer night in 2008, while I was staying at a friend's cabin in rural Vermont. Turning off the election-season radio broadcast, I stepped onto the back porch and found myself instantly engulfed by the darkness of the forest. The scene was peaceful, but hardly still. The arresting smells of cut grass and lilac bushes in bloom were heavy in the air. The shrill chirping of frogs and crickets rose and fell in swells of sound. Fireflies, with alternating flashes of bioluminescence, swirled over a field. Visible in the further distance were the more regular flashes of lights at the top of a satellite tower. A jet crossed the sky, displaying tiny red and green beacons on its wing-tips. And my cell phone buzzed quietly in my pocket, my reception in the woods just strong enough to receive text messages. As these signals drifted into perception, their deep commonalities became suddenly vivid. In varying degrees, each was an instance of systematic communication. Each was a solution to the same basic problem that faces all physical organisms whose survival depends on social exchange.

For physically distinct creatures, isolated from one another in space as well as time, information can only be exchanged by making it publicly accessible through manipulations of the physical environment. By modulating light and sound frequencies (or properties salient to other sensory modalities), and interpreting these modulations, living organisms coordinate and promote survival. But in order for such exchanges to be fast and reliable, they must be governed by shared codes— codes which determine how practicable information can be encoded in and then extracted from the altered environment. Many such codes, like those of the lilac, frog, or firefly, are biologically determined. Others are the product of human convention, including not only language, but also symbolic, diagrammatic, and pictorial signaling. The thesis of this dissertation is that there is an impressively wide spectrum of conventional systems of representation, corresponding to the many ways that the problem of communication can be solved, and that these systems can be de-

scribed and explained using the tools of contemporary mathematical semantics.

The remainder of this chapter is devoted to contextualizing and elaborating this claim, and anticipating the ways in which the remaining chapters provide evidence for it. In §1, I locate the subject of the dissertation within the broader philosophical category of representation. In §2, I describe the project of "general semantics", which aims to study the full range of conventional systems of representation, within the broad framework of natural language semantics. The present work is cast a contribution to this study. Finally, in §3 I discuss the conjecture that there is a "semiotic spectrum" of representational kinds, extending from purely symbolic to purely iconic representation. This dissertation is a study of the spectrum's polar limits.

1 Philosophical Setting

In the broadest terms, the subject of this dissertation is REPRESENTATION. Representation, as I shall use the term, is a very general asymmetric relation that holds between at least two entities—that which represents, and that which is represented. I will call that which represents a SIGN, and that which is represented its CONTENT. For A to represent B is for A to be in some sense *of*, or *about B*, and in many cases for A to *make ascriptions* about B. Characteristically, representation can hold between entities that are spatially and causally quite distant from one another, and lack any substantive commonality.

Representation inherits its philosophical importance from the central explanatory role it plays in common-sense psychology, contemporary cognitive science, and linguistics. We regularly explain why people and animals behave in the way that they do by appeal to their beliefs and desires. Cognitive scientists explain why organisms have the neural anatomy that they do in terms of the computational operations such anatomy can perform on biologically adaptive representations. And linguists explain why communicative exchanges take the form that they do in terms of the representational content carried by the signals that make up such exchanges. In all cases, explanations cast in terms of representation allow us to describe the relationship between agents, signs, and the world at a uniquely illuminating level of abstraction.

As is often noted, representation cannot be defined simply as the transmission of physical information. Smoke carries physical information about fire, and waves carry physical information about the direction of the wind. But these are not instances of representation in the sense intended here. Whereas phenomena such as these merely indicate their causes, representations are objects, events, or states which are specifically exploited by thinking agents to guide attitudes and actions.

Through such use, representations are given the function of ascribing properties to states of the world. And when this function is realized, representations are true, accurate, or correct. The capacity to be correct or incorrect characteristically distinguishes representations from ordinary objects like smoke or waves.[1] There remains the open and difficult question of exactly what grounds the difference between information transmission and genuine representation. But in this dissertation, I'll have little to say about this metaphysical puzzle.

A different breed of open question concerns the taxonomy and structure of representational kinds: what are the various forms of representation? how they related? what are their respective architectures? *These* are the questions taken up here. In what follows I'll try to specify the limited range of phenomena under consideration in this dissertation, in a way that is sensitive to the tremendous variety of representation found in nature.

Mental and public representation

Mental states have the ability to represent states of affairs outside the mind. Thus my concept of Nelson Mandela represents the person Nelson Mandela; my mental image of Mandela revisiting Robben Island in 1994 represents that historical situation; my current perceptual experience represents the desk at which I am writing this document. Mental states are arguably the fundamental forms of naturally occurring representation.[2] But mental states are not the only entities with representational powers. The word 'Nelson Mandela' *refers* to Mandela, and in the relevant general sense represents him. The following sentence *expresses*, in English, a certain proposition, and in the same general sense it represents that proposition.

Nelson Mandela revisited Robben Island in 1994.

And the photograph below represents a particular moment during Mandela's return.

---

[1]This conception of the relationship between content and functional norms is derived from Burge (2010).
[2]See Burge (2010) for extensive discussion.

Broadly speaking, there are two kinds of representation. MENTAL REPRESENTATION is that which relates cognitive states and processes— such as concepts, thoughts, mental images, perceptual states, and the like— to their content. PUBLIC REPRESENTATION associates content with objects and states which are physically instantiated outside of the organisms that produce them, and perceptually accessible to others.[3] They include firefly flashes, cricket chirps, and products of human communication like sentences, pictures, diagrams, and gestures. What is the relationship between mental and public representation? The dominant view, which I accept here, is that public representation is always derived, one way or another, from mental representation.[4] Yet the topic of my dissertation is the derivative thing: public representation. Though the theory of mental representation informs the present investigation, it is not its subject matter.

There are two main justifications for this choice of focus. First, though public representations may be metaphysically derivative, they differ in basic ways from mental representations. The dominant naturalistic theories of mental representation ground it in universal features of human biology and our causal interactions with the physical world.[5] By contrast, public representation is at least partly a product of human social convention, as evinced by the dramatic cross-cultural diversity of linguistic, diagrammatic, and pictorial systems of representation.[6] In addition, it is public representation that plays the primary role in communication: while solipsistic thought requires only mental representation, communication requires that we generate sounds, gestures, or marks on paper which are endowed with representational properties. Thus public representation's social roots bring with it a unique set of problems and questions. It constitute's a legitimate subject of study in its own right.

Second, public representation is by nature easier to study than mental representation, and therefore serves as one important window into the obscure mechanisms of the mind. On one hand, public representations, unlike mental representations, are themselves perceptually accessible, so it is relatively easy to measure and model their structure. On the other hand, because public rep-

---

[3]This taxonomy is not intended to be exhaustive. In particular, abstract objects such as propositions may have representational powers. (King 2007, pp. 129-32) Though King himself would likely classify propositions as artifacts, one need not: Frege (1918) conceived of "thoughts" as timeless abstracta, independent of human affairs, but representational nonetheless. Another possible locus of representations are the internal states of man-made computers. Note also that not all public objects which figure in communication are necessarily representations. Thus Kivy (1984) argues that instrumental music lacks representational properties.

[4]Three classic proposals of this form are Grice (1957), Lewis (1969; 1975), and Fodor (1975). See Speaks (2010, §3.1) for a survey of such proposals.

[5]Theories of this form are proposed by Dretske (1981), Millikan (1984), andFodor (1992). See Loewer (1997) for a survey of such proposals.

[6]This point, obvious for languages, is still not widely recognized with respect to pictorial systems; this issue is taken up in Chapter 1. On cultural and historical variation in diagrammatic systems, see Tufte and Schmieg (1985) and Tversky (2000).

resentations play a central role in social exchange, sophisticated communicators such as humans have developed ways to talk and think directly about representational properties such as truth, accuracy, meaning, and entailment. And because we are able to form robust and stable judgements about these properties, it is possible to develop detailed mathematical theories which model their behavior. It is presumably for these kinds of reasons that the formal analysis of public language has played such a key role in the study of symbolic thought.[7] My hope is that by developing parallel formalizations of other kinds of public representation, we will be able to make new strides in understanding their mental counterparts. In particular, the formal analysis of pictorial representation may serve as a model for the philosophical understanding of picture-like representations in the mind, such as sensory perception and mental imagery.

Improvised and systematic representation

Even among public representations there are basic differences in kind. Suppose you have lost your favorite writing pen. In order to help you regain it, I ask you what it looks like. You hold up another pen in response to my question. Your act of holding up a pen communicates specific information— to wit, the missing pen looked like the held pen. Thus you used the held pen to *represent* the lost pen. But there was no fully systematic rule that associates the act of holding up a pen with the information that some other item looked like it. Instead, I infer the significance of your improvised gesture by reasoning about the probable motivations behind it. This case of representation was devised "on the fly"; it was an example of IMPROVISED REPRESENTATION.[8]

Improvised representation is an effective means of communicating so long as the stakes are sufficiently low and the information which must be communicated is relatively unspecific. When the stakes are high, some guarantee must be secured that a given signal will be interpreted in a predictable and efficient manner. The solution is for the signal sender and receiver to converge on a *code*, or scheme of SYSTEMATIC REPRESENTATION. Systematic representation allows communicators to bypass the unpredictable complexity endemic to improvised signaling, relying instead on established rules of representation.

Consider an example championed by Lewis (1969). On April 18, 1775 Paul Revere and his accomplice Robert Newman faced the following communication problem: it was necessary for Newman, on one side of the Charles River, to communicate the path of the British invasion to Revere, located on the other side of the river. One can imagine that the men might have left the problem

---

[7]Fodor (1975) is a classic source of such developments.

[8]Stanley (2000) marks this distinction between means of conveying content which are improvised and pragmatic, and those that are systematic and semantic.

alone, to be resolved by improvisation. But the stakes were too high; improvised representation is typically highly ambiguous, and in this case, the reliable exchange of information was crucial. A pre-arranged, systematic solution would be necessary. Thus the famous code: Newman would hold up one lantern if the British were coming by land, and two if by sea. This is a paradigm case of systematic representation.

My dissertation is concerned exclusively with systematic representation. Improvised representation is governed by all-purpose reasoning and general world-knowledge. A theory of improvised representation depends largely on an inclusive account of the psychology of the organisms engaged in communication. Systematic representation, by contrast, is always governed by a small number of interpretive rules. A theory of systematic representation attempts to articulate the rules underlying particular systems of representation. Though systematic representation is, by definition, constrained, a theme of this dissertation is the variety rule-architectures at work in human communication. The differences among such systems turns out to be illuminating.

Conventional and biological representation

Finally, forms of representation which are both public and systematic further divide into those which are regulated largely or entirely by biology, and those for which social convention plays a significant role. For example, Von Frisch (1967) first described a system of public representation used by honey bees to communicate the location of sources of water, pollen, and nectar. By flying along a diagonal axis in relation to the position of the sun, these bees are able to convey precise spatial information to their hive mates. The code by which such information is communicated is surprisingly precise and consistent. Presumably it is also fixed by genetics: it is a case of BIOLOGICAL REPRESENTATION. It is very likely that many forms of human signaling, such as smiling on the part of infants, have similar biological determiners. (Wolff 1963)

By contrast, signaling systems which are established by explicit agreement, such as Paul Revere's lantern code, are obvious examples of CONVENTIONAL REPRESENTATION. Following Lewis' (1969) landmark analysis of convention, I will also count natural languages as conventional systems of representation. This is so even though the structure of human languages is partly constrained by biology, and even though languages are not, in general, established by explicit agreement. As I will argue in Chapter 1, systems of pictorial representation are conventional in the same way. Lewis (1969) observed that such representational conventions come about as solutions to coordination problems of communication. In this way, conventional systems of representation are cast as social tools, created and sustained by humans to overcome recurrent obstacles to information

exchange. The variety of systems of representation corresponds to the variety of way such tools may be designed and deployed.

We have now arrived, finally, at the general subject of this dissertation: those forms of representation which are public, systematic, and conventional. Following King and Stanley (2005), I propose that these types of representation are exactly those which have a *semantics*. Thus my concern in this dissertation is not with representation generally, but with those kinds of representation that are governed by semantics. This usage of "semantics" is admittedly irregular, since the term is typically applied to languages alone. Yet, as I will now discuss, my usage corresponds to a natural broadening of the traditional purvey of natural language semantics.

2   General Semantics

I have identified my subject as those forms of representation which are at once public, systematic and conventional. It should be evident that this class includes more than the natural languages. For example, maritime signaling with flags, traffic lights, and trail-blazes all rely on interpretive rules for signals which are obviously systematic and fixed by social convention. And we can imagine a *general* semantics, a study much like that of contemporary natural language semantics, which explores this broader class of systems.[9] Indeed, the kinds of examples mentioned above all seem to rely on simple correlations of signs and content very much akin to the lexicons of natural language. An extension of natural language semantics to include these other, essentially symbolic systems, would be straightforward.

But a general semantics should not stop there. In Chapter 1, I will argue that *pictorial* representation, like linguistic representation, is both systematic and conventional. If this conclusion is correct, then a general semantics must account not only for languages and language-like systems, but also for those systems which are paradigmatically *unlike* language. (Some have attempted to assimilate pictorial representation to the familiar architecture of linguistic representation; but I will argue that this approach is misguided.) Others have argued persuasively that systems of diagrammatic representation are governed by rigorous logics, and even offered detailed semantic analyses of these systems.[10] A general semantics would study *all* such systems of representation— symbolic codes, languages, diagrams, pictures, three-dimensional models, audio representations, and so on. It would attempt, for each kind of system, to articulate the rules governing the interpretation signs

---

[9]The phrase "general semantics" is a play on Lewis' (1970) article by the same name. Lewis envisioned his proposal providing a framework for the semantics of languages, not systems of representation in general.

[10]For example, Shin (1994) was the first to provide a model theory and accompanying proof of soundness and completeness for Venn diagrams.

in that system.

So conceived, general semantics subsumes natural language semantics, and thereby gains extra philosophical edge. While linguistic representation does appear to constitute a natural kind, worthy of independent study, it does *not* appear to be the most general natural category in the nearby taxonomic vicinity. Philosophical inquiry, whose special charge is inquiry into general and abstract regularities, should be especially interested in the broader categories into which linguistic representation fits. It is general semantics, not natural language semantics, which seems to me of primary philosophical interest.

The project of general semantics is more or less that of the discipline of *semiotics* or *semiology*, the "science of signs." From its inception, semioticians have envisioned an ecumenical project, far broader than contemporary linguistic analysis. Consider, for example, Saussure's early statement of purpose:

> A language is a system of signs expressing ideas, and hence comparable to writing, the deaf-and-dumb alphabet, symbolic rites, forms of politeness, military signals, and so on.... It is therefore possible to conceive of a science *which studies the role of signs as part of social life.* It would form part of social psychology, and hence of general psychology. We shall call it *semiology* (from the Greek *sēmeîon*, 'sign'). It would investigate the nature of signs and the laws governing them. (Saussure et al. 1922, §33)

The difference between the agenda envisioned here and that of semiotics is largely one of history and methodology, not subject matter. The progenitors of semiotics were Saussure et al. (1922) and Peirce (1906); it was later developed by Morris (1946), Barthes (1968), Sebeok (1977), Eco (1979), and others. The discipline came of age in the era of structural linguistics and anthropology, before the influences of Montague and Chomsky were fully ramified, and later iterations were influenced by phenomenology and post-structuralism. Even at its most scientific, semiotics has been an informal discipline. By contrast, the course taken here is the inheritance of contemporary analytic philosophy of language, formal linguistics, and computer science. It aspires towards mathematically precise models of actual behavior, even when this means abstracting away from important empirical variations. While freely drawing from the insights of semiotics, I hope that this approach can shed light where semiotics has fallen short.

I conclude this section by clarifying what I take to be the three key elements of any theory within general semantics: signs, content, and systems of representation. SIGNS are the bearers of content; they are abstract types whose tokens are physical signals. Words and sentences are examples of signs; so are gestures, diagrams, and pictures; in the right setting, actions, events, or states may also instantiate signs. Linguistic theory posits only simple signs with atomic structure

and complex signs with syntactic tree structure, but a general semantics need not be limited by this assumption. Interpretive rules take sign structure as input, whatever that may be, including graph-theoretic, topological, or geometrical structure, or some hybrid of these.[11] For example, in Chapter 3, I'll develop a semantics for pictures where the signs have geometrical structure.

CONTENT, in one of its primary instances, is that information which is systematically associated with signs in the course of communication. The nature of content is a philosophically fraught subject, and I wish to avoid substantive commitments here. Typically, theorists countenance both referential and representational content. Referential content, normally associated with individual linguistic terms, includes objects, sets of objects, and properties. Representational content is normally associated with clausal linguistic expressions, but here with complete diagrams and pictures as well. It construes the world as being a certain way, and as a consequence, it either correctly or incorrectly represents the world. Thus many theorists have either modeled content as, or identified content with correctness-, accuracy-, or truth-conditions. Linguists have recently posited still other forms of content, including *dynamic* content— whose function is to transform a conversational state, rather than merely describe the world. Further complications have to do with various modes of content for questions, commands, and declaratives. For the purposes of this dissertation, I shall typically model the content of complete expressions (linguistic or pictorial) as truth- or accuracy-conditions, modeled in turn as sets of possible worlds. But I flag here that this is only a working assumption, useful for the purposes of the analytical tasks at hand and familiar to a range theoretical projects. I trust the presentation here can be translated into other terms if necessary.

Finally, SYSTEMS OF REPRESENTATION are the interpretive rules which map signs to content. They can be implemented in the cognitive mechanisms of individuals, and shared by interpretive communities. A semantics for such a system describes the rules of the system in such a way as to abstract from the computational details of different possible implementations. General semantics, as I have defined it here, is the study of the semantics of those systems of representation which have been established by social convention. The aim of this dissertation is to make a contribution to this general study, and to demonstrate its viability.

3   The Semiotic Spectrum

My investigation of semantic representation is guided by Pierce's (1906) distinction between SYMBOLIC and ICONIC representation, which I take, at least pre-theoretically, to demarcate a natu-

---

[11]It is worth noting that all such structures, including syntactic trees, can be modeled as graphs, and it may be that graph theory provides the most general framework for describing the various possible inputs to semantic interpretation. I won't pursue this speculation here.

ral distinction among types of semantic system. The paradigm case of symbolic representation is linguistic meaning, but the class of of symbols also includes stop lights, smoke signals, maritime signal flags, gang signs, and conventional gestures like "thumbs up". The paradigm case of iconic representation is pictorial depiction, exemplified by drawings, paintings, photographs, and maps, but the class of icons also includes audio recordings, three-dimensional models, and depictive gestures. To a first approximation, symbolic representations are those which bear a merely arbitrary relation to their content, while iconic representations are those which bear more intimate or natural relation to theirs— a relation often characterized in terms of "resemblance" or "likeness".

Between the poles of iconic and symbolic representation seem to lie a vast swath of intermediate representational kinds. These are systems of representation which appear in various ratios to be both iconic and symbolic. Proximal to pure pictorial representation are images which represent non-visual modalities visually, such as spectrograms (visual representations of sound) and thermograms (pictorial representations of heat distribution). Somewhat more symbolic are stylized images, such as cartoons and stick figure drawings. The vast majority of images created before the advent of photography exhibit some form of stylization. Even more symbolic and abstract representations include directional arrows, temporal graphs, and Venn diagrams. Further down the scale towards language are pictographic writing systems such as Egyptian hieroglyphics or the historical precursors of contemporary Chinese *Han* characters. Finally, phenomena such as quotation, onomatopoea, and conventional sound effects (e.g. "knock, knock!") all suggest the presence of picture-like expressions in natural language itself. Any empirically adequate account of human communication must allow for this incredible range of representational systems.

A natural hypothesis is that these representational kinds constitute a more or less smooth spectrum extending from the purely symbolic to the purely iconic.[12] I term this the SEMIOTIC SPECTRUM. But since so little is known about the mechanisms of non-linguistic systems, it remains an open question whether the observed variation does in fact reflect a continuum of possibilities. This dissertation begins to investigate this question by charting out the poles of the hypothesized spectrum— pictorial representation, at one extreme, symbolic representation at the other. I will not address the rich variety of intermediate representational kinds; happily, there is much more work to be done.

In the 20^{th} century, the semantic analysis of representational systems has has focused almost exclusively on language. When philosophers and linguists *have* undertaken semantic analysis of non-linguistic representation, the general strategy has been to extend linguistic models of meaning

---

[12]Similar claims are defended by Eco (1979) and McCloud (1993).

to the non-linguistic realm.[13] But our rich understanding of language has come at the expense of an equally detailed and rigorous understanding of non-linguistic sign systems.

My dissertation is intended as a partial corrective. I begin with the class of systems arguably most different from language: those governing the interpretation of *pictorial images*. I argue that the kinds of imagery which figure in practical communication— maps, architectural drawings, road signs, gestures, and the like— are governed by systematic and conventional semantic rules, broadly akin to those of natural language. But whereas the semantics of languages are based on arbitrary associations of sign and denotation, conventional "systems of depiction," as I term them, are based on rules of *geometrical* transformation. Drawing on recent research in computer graphics and computational vision, I go on to develop a geometrical semantics in precise formal detail, for a simple system of perspective line drawing. One upshot is that the traditional, language-based conception of semantics must be overhauled to allow for the wide variety of interpretive mechanisms actually at work in human communication.

In Chapter 1, I present and defend the Pictorial Semantics Hypothesis, the claim that pictures are associated with their content by systematic and conventional interpretive rules. Chapter 2 examines the most orthodox theory about what kind of semantics pictures actually have. According to this proposal, pictorial representation is always grounded in resemblance. I argue that the very structure of this theory renders it empirically inadequate: there are systems of depiction which cannot be modeled in terms of resemblance alone, but must instead be based on a more general conception of geometrical transformation. In Chapter 3, I therefore introduce the Geometrical Projection Theory, an alternative, and largely novel account of the semantics for pictures, according to which pictorial representation is grounded in geometrical projection. Finally, in Chapter 4 I turn to symbolic representation. Here I survey extant attempts to define symbolic representation, and offer a new account which is aimed at teasing apart the crucial difference between symbolic and iconic representation.

Ultimately, I conclude that both pictorial and linguistic representation have semantics; both are governed by conventions which systematically map signs to content; both are types of social tool created and sustained to solve coordination problems of communication. Yet, as I will argue, they embody deeply divergent strategies to achieve this mandate; they are very different kinds of tools, suitable for very different kinds of communication. Whereas linguistic representation is based on arbitrary associations and compositional rules, pictorial representation is based on methods of geo-

---

[13]See, for example, Goodman (1968), Shin (1994), Lascarides and Stone (2009*a*). There are also exceptions: Malinas (1991) and Blumson (2009*b*).

metrical projection. This distinction reflects an even more basic one. Purely iconic representation is based on *general* rules for mapping signs to content, while purely symbolic representation is based on *particular* rules for mapping individual signs to content. This deceptively subtle contrast explains many of the obvious differences between the two types of representation, and suggests why each has found such a central role in human communication.

Chapter 1

Pictorial Semantics

Because humans cannot know one another's minds directly, every form of communication is a solution to the same basic problem: how can privately held information be made publicly accessible through manipulations of the physical environment? Language is by far the best studied response to this challenge. But there is a diversity of non-linguistic strategies for representation with external signs, from facial expressions and fog horns to chronological graphs. Among these alternatives, the class of *pictorial* representations dominate practical communication— witness the proliferation of maps, road signs, newspaper photographs, scientific illustrations, television shows, architectural drawings, and even the fleeting imagery of manual gesture.[1] In such cases, what bridges the gap between the physical media of pictures and the information they manage to convey? Over the course of the next three chapters, I offer a partial answer to this question.

The thesis of this chapter is that pictures are associated with their informational content by

---

[1]See McCloud (1993), Tversky (2000), and Lascarides and Stone (2009*a*; 2009*b*) for a wide range of examples.

13

systematic rules selected by social convention— not, as one might have thought, by a purely bio-logical reflex of vision, nor by open-ended rational speculation. In this respect, the interpretation of pictures is importantly similar to the interpretation of sentences in a language. I will call this the *Pictorial Semantics Hypothesis*.[2] In Chapter 2, I will argue against the most prominent extant proposal about what kind of semantics pictures actually have; according to this view the semantics of pictorial representation are always grounded in resemblance. The failure of this view sets the stage for Chapter 3, where I defend a specific and largely novel theory of pictorial semantics. According to this idea, for a picture to accurately depict a scene is for there to be a projective mapping from the three-dimensional features of the scene to the two-dimensional features of the picture plane. Recovering a picture's content from its surface features involves working out the kind of scene from which the picture could have been projected. In this respect, the strategy used to en-code information in pictures is decidedly unlike that employed by languages. I will call this the *Geometrical Projection Theory*.

If correct, these conclusions have import for the study of a variety representational kinds be-yond the pictorial. By providing a precise formal framework in which geometrical structures are as-sociated with correspondingly geometrical content, they show that rigorous semantic analysis can assume forms quite unlike those most familiar today. Languages constitute just one class among a broad spectrum of representational systems, all of which deserve to be studied under a general science of semantics. In addition, this framework may serve as a model for the philosophical analy-sis of visual perception. Just as the study of public languages has informed our understanding of private cognition, the present account of public images may help to guide the philosophic theory of perceptual representation generally.[3]

The argument of the chapter is laid out in four parts. In §1, I articulate the Pictorial Seman-tics Hypothesis, the claim that pictures are associated with their contents by rules which are both systematic and conventional. Sections §2-4 defend the key components of this hypothesis. In §2, I argue that pictures are the bearers of genuine representational content. In §3, I provide evidence that the mapping between pictures and their content is systematic. In §4, I argue that these system-atic mappings are established by social convention. In §5, I provide a formal, accuracy-conditional

---

[2]I shall simply assume that natural languages have semantics, but even this is still contested. Stanley (2000) and King and Stanley (2005) survey such anti-semantic accounts of language and argue against them. The present paper takes the same basic considerations that motivate King and Stanley and applies them to the case of pictures.

[3]There are other potential applications as well. A precise semantics of pictures would be illuminating for the study of imagistic mental representation— whether or not the suggestive analogies between these domains are taken literally. (Block 1982, pp. 1-16; Pylyshyn 2006, ch. 7) In the philosophy of science, debate has arisen recently about the nature of scientific models and scientific representation in general; here pictures are frequently held up as a paradigm case (see, e.g. French 2003; Giere 2004, §4; Cohen and Callender 2006, p. 11). Even in metaphysics a pictorial semantics might elucidate the metaphysical account of possible worlds known as "pictorial ersatzism." (Lewis 1986, §3.3)

framework for any theory of pictorial semantics. Finally §6 defines the scope of inquiry for the chapters that follow.

1 The Pictorial Semantics Hypothesis

The use of languages in human society is universal. Languages, manifested variously in speech, writing, and bodily signs, enable the fast and reliable communication of complex information between individuals. It is widely believed that languages are based on complex rules which regulate the production, structure, and meaning of linguistic signals. The study of SEMANTICS, in particular, hypothesizes the existence of tacitly understood rules which govern meaning— rules which map linguistic signs to representational content. Such theories contribute to an explanation of linguistic communication by describing mechanisms through which language users employ shared conventions to encode and decode information in linguistic signs. Since the 1960's and 70's the mathematical study of semantics has flourished, producing detailed predictions about an ever-widening range of linguistic data. Now one of the central pillars of modern linguistics, semantic theory has proven to be a progressive and exact research program.

Of course, there are other, non-linguistic forms of representation as well, most notably, representation by pictorial images. The use of pictorial representations as tools of public expression is ancient in human history, and thrives without special training or tools in the form of iconic gesture. In modern industrial society, pictures 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. And the roots of pictorial representation run deep: the spatial organization of the human visual cortex strongly suggests that picture-like representation is one the basic strategies for information management implemented by the brain, particularly in low-level perceptual processing.[4]

The thesis of this chapter is that we should posit a semantics for pictures, just as we have posited a semantics for natural language. The proposal is not, implausibly, that pictorial and linguistic representation are governed by the same *kind* of semantics— only that both are governed by semantics.[5] More explicitly:

(1) ***Pictorial Semantics Hypothesis***

Pictures are associated with their content by systematic and conventional interpretive rules.

---

[4]See Pylyshyn (2006) for extended discussion. It remains a controversial question whether there are also picture-like representations in higher cognition.

[5]On the general conception of semantics see Tarski (1944, p. 345) and Haugeland (1991, p. 62).

This Hypothesis is constituted by three key claims, each of which will be defended separately over the course of this chapter: (i) that pictures have representational content; (ii) that the association between pictures and their content is systematic; (iii) that the systematic association between pictures and their content is conventional.

The Pictorial Semantics Hypothesis has few allies. Both philosophers of language *and* philosophers of depiction have doubted that non-linguistic representation is amenable to the same kind of semantic analysis that characterizes language. As a consequence, representation by pictorial images is widely dismissed as a legitimate subject of systematic study. Those who are moved by such skepticism must admit that pictures often play a pivotal role in successful communication, but they insist that pictorial communication is wholly accounted for by some admixture of facts about human visual perception, general principles of rationality, and heterogeneous world knowledge. In short, they believe there is a *pragmatics* of pictorial representation, but not a semantics. Thus Walton (1990, p. 351) writes that language "is to be defined in semantic and/or syntactic terms," but "depiction is a pragmatic notion." And Stanley (2000, p. 396) concludes, "we should therefore be suspicious of attempts to forge philosophically significant analogies between the different processes underlying the interpretation of linguistic and non-linguistic acts."[6]

What accounts for the widespread skepticism about pictorial semantics— given the widespread *acceptance* of linguistic semantics? One reason, I believe, is that the evidence for pictorial semantics has not been fully appreciated; one of the central tasks of this chapter is to review this evidence. There are also a number of substantive conceptual and empirical worries one might have about a pictorial semantics; another central task of this chapter is to address these concerns. But there are other sources of doubt which stem from basic misunderstandings about the project of pictorial semantics. In the remainder of this section, I address three of these issues, with an eye towards clarifying the core aims of the present investigation.

First and foremost is the simple conflation of *pictures* and *visual art*. The interpretation of works of art is by nature unbounded, in part because artistic meaning is specifically tuned to expressive and metaphorical significance beyond literal content, in part because such meaning is often a product of intentionally violating the very interpretive norms that govern ordinary communication. This makes artistic meaning an admittedly unlikely candidate for systematic semantic analysis. But just as linguists do not take poetry as their primary source of data, there is no reason that stu-

---

[6]And Borg (2006, p. 261) announces "fundamental differences between communicative acts in general and linguistic acts in particular. Linguistic acts, uniquely in this area, have a crystallized component to their meaning, an element which they carry with them across all contexts and which may be accessed by a competent language user even if she has no access at all to the speaker's original intentions... Yet these genuinely code-like qualities seem very different to the properties of other communicative acts, which depend on context in a far more constitutive way."

dents of pictorial representation should take visual art works as their point of departure. Instead, we should only expect robust semantic regularities to emerge from the use of pictures deployed for practical information exchange, under conditions in which efficiency and fidelity are at premium. Examples include road signs, maps, architectural and engineering drawings, textbook illustrations, police sketches, and so on. In this paper I will confine my attention exclusively to pictures which 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.)

A second reason for doubt is the assumption that, as a matter of definition, semantics has the structure of linguistic semantics. Yet it is correctly observed that the mechanisms underlying pictorial representation are not those of linguistic representation. 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 contexts. 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 or photograph of a scene and the scene itself is one of intimate geometrical correspondence, nothing like the stipulative association between a word and its denotation. Since semantics is defined in terms of the the atomic and arbitrary architecture of language, and since pictorial representation resists such definition, semantics is not for pictures.

But this argument assumes an obtuse vision of pictorial semantics. Insofar as the idea of a pictorial semantics is plausible, it is not the idea of a semantics whose architecture is just like that of language. Rather, by abstracting away from the atomic and arbitrary basis of linguistic representation, we can isolate a more general concept of semantics. The Pictorial Semantics Hypothesis posits a semantics for pictures only in this ecumenical sense— a systematic and conventional mapping from signs to content.

A third and final concern is that none of the extant theories of depiction could plausibly form the basis of a systematic pictorial semantics. In many cases this is quite deliberate. For example, the *Symbol Theory* of depiction introduced by Nelson Goodman in his 1968 *Languages of Art* was articulated with groundbreaking rigor, but it rejected the view that depiction was in any substantive sense systematic. Goodman held the implausible view that pictures are associated with their content only by arbitrary stipulations, much like the relation between the word "tree" and the property of being a tree. A much more promising tack is pursued by *Perceptual Theories*, which define the content of a picture in terms of the cognitive and perceptual effects it has on viewers. Yet in

general, such theories purposefully eschew the systematicity of semantic analysis, holding instead that an exact definition of depiction ultimately rests in part on the gnarly mechanics of biological vision. The only theory of depiction which immediately lends itself to a semantics is the orthodox *Resemblance Theory*, according to which a picture represents a scene in virtue of being *similar* to it in certain relevant respects. But such accounts are generally products of philosophical dialectic, not quantitative empirical inquiry. And they have been the subject of numerous and vigorous objections over the last fifty years. In Chapter 2, I add my voice to this dissenting chorus, arguing that the the very structure of such similarity-based views forces them into false predictions for a broad range of cases. But if none of the extant philosophical theories of depiction could form a suitable basis for a pictorial semantics, how could the Pictorial Semantics Hypothesis possibly be right?

The argument of this chapter is not intended as an endorsement, even implicit, of any of the extant philosophical theories of depiction. If the Pictorial Semantics Hypothesis is correct, then we will need to identify some other, perhaps novel account of depiction to answer to the requirements of a pictorial semantics. And it is this mandate which Chapter 3 aims to fulfill. Inspired by the study of images in projective geometry and recent work in computational perceptual science, the Geometrical Projection Theory which I'll defend there diverges from its predecessors both in the substance of its analysis and the exactness of its expression. According to this theory, a picture depicts a scene just in case the picture can be derived from the scene by general but precise rules of geometrical projection. Thus pictorial representation is grounded neither in arbitrary meaning associations, subtle resemblances, nor the perceptual responses of viewers. The result is an account that boasts a rigor, simplicity, and empirical validity which has eluded other proposals; it is an appropriate foundation for a semantics of pictorial representation.

I hope to have clarified the rudimentary commitments and ambitions of the Pictorial Semantics Hypothesis. In the remainder of the chapter I will elaborate and defend each of its three key components: (i) the presupposition that pictures have representation content; (ii) the posit of systematic mappings from pictures to their content; and (iii) the claim that these systematic mappings are established in human communication by social convention.

Finally, a terminological note: a PICTURE, as I use the term, is a *type* of pictorial sign, not a concrete token. I intend to include typical cases of photography and so-called "representational" (as opposed to abstract) painting and drawing. But I will not consider composite images, such as comic books, or images which integrate symbolic elements like textual labels or color-coding. For purposes of brevity, I shall also ignore such abstract representations as cartesian graphs, pie charts,

and Venn diagrams— as well as alternative media such as cartoons, film, relief, scale models, music, and audio recording.[7]

2 Pictorial content

The Pictorial Semantics Hypothesis presupposes that pictures have REPRESENTATIONAL CONTENT. In this section I will elaborate and defend this position.[8] At minimum, by claiming that pictures have representational content, I mean that pictures are representations which specify a way reality could be, with respect to objects and properties that can be physically remote from the picture itself. The content of a picture is identical to the way reality could be, as specified by the picture.[9] The opposing view holds that pictures are ordinary, non-representational objects; their communicative power stems only from their causal effects on the minds of viewers.

The idea of pictorial content at work here is familiar enough. Of Frank Lloyd Wright and Jack Howe's 1935 architectural drawing with which we began, it is natural to say that it depicts a house, of a certain shape and color, in a certain relationship to its environment, perched over a waterfall, and so on. In this way, we commonly associate pictures— flat, marked surfaces— with specifications of reality as having specific, spatially extended features. That is, we naturally associate pictures with representational content. (Such content need not specify how reality *actually* is; the drawing here was made before it was certain that the house depicted would be built.)

Content in this sense is not an aesthetic or impressionistic property of pictures. Like sentences, images can reliably carry precise and useful informational content. For example, the exact content carried by the Frank Lloyd Wright drawing, about a hypothetical building plan, played a crucial part in the Kauffman family's eventual decision to build the house, at considerable cost.[10] In general, the use of pictures to communicate specific information in high-stakes contexts, such as those involving engineers, architects, and pilots, demonstrates the efficacy of pictorial content, and should quell the suspicion that the content of pictures is any way less exact or stable than that of sentences.

Pictorial content must be distinguished from other semantic properties of pictures. First, the content of a picture is not its REFERENT. Whereas the content of a picture is the depictive information it carries, its referent is the scene which the picture is supposed to carry information *about*.

---

[7]The discussion here can be recapitulated for nearly any of these media. See Kivy (1984, ch. 1-2) for a detailed parallel discussion of musical content.

[8]The general thrust of this section— that pictures have content much like sentences, and that this content can be defined in terms of correctness conditions— is clearly prefigured in Wittgenstein (1921/1997, §2.21-224).

[9]Since this characterization does not put direct constraints on the nature of content, it is compatible with the many divergent theories of content on offer in contemporary debates. See Siegel (2009, §3) and King (2007, ch. 1) for discussion.

[10]See Toker (2003, pp. 195, 250-7).

While pictorial reference and content both answer to colloquial uses of the phrase "what is depicted", they come apart strikingly in cases of misrepresentation.[11]

Let's say that I attended the Obama press conference on April 27, 2006 at the National Press Club in Washington, D.C. The next day I decide to draw what I saw there at a particular moment during the day, from the vantage point of my front row seat. I produce the first image at right. A week later I set out to draw the same scene, but this time poor memory and lack of common sense conspire against me, resulting in the second image.[12] 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. In this sense the two images have the same *referent*. On the other hand, there is a clear sense in which they do *not* depict the same thing. The first picture depicts Obama with short hair; the second depicts him with longer, spiky hair. In this sense, the two images have different *content*.

In addition, not all content associated with pictures is genuinely pictorial. In the 17th century painting by Philippe de Champaigne at right, the rendering of the tulips represents life, the painting of the hour-glass represents the passing of time, and that of the skull, inevitable death.[13] Yet none of these symbolic elements figure in the image's pictorial content, in the sense intended here. Instead, the painting depicts a flower, a skull, and an hour-glass, in a certain arrangement, at a certain time, and under certain lighting conditions, but nothing more.[14]

Philippe de Champaigne, "Vanitas" c. 1671

Finally, the content of a picture is distinct from what it happens to *communicate* to a particular viewer on a particular occasion; rather the picture's content is what it *depicts* independent of what

---

[11]The distinction between pictorial reference and pictorial content is widely observed in the Philosophy of Art; see for example Beardsley (1958, pp. 270-2), Goodman (1968, pp. 27-31), and Lopes (1996, pp. 151-2). Various ways of prising reference from content, including cases of misrepresentation, are discussed by Knight (1930, pp. 75-6), D. Kaplan (1968, pp. 198-9), and Lopes (1996, pp. 94-8). Cummins (1996, pp. 5-22) observes a parallel division with respect to mental representation. In general, the distinction between pictorial reference and pictorial content is closely aligned with Kripke's (1977) distinction between speaker reference and semantic reference. A more thorough analysis might reveal that these are species of the same phenomenon. See Burge (2010, pp. 30-46) for an extended discussion along these lines.

[12]Both drawings are based on the photograph by Mannie Garcia, AP.

[13]See Lubbock (2006).

[14]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, §2) and Peacocke (1987, p. 383). Goodman (1968, p. 5), for one, appears to reject the distinction altogether.

is communicated.[15]

For example, if you write to me asking how my vacation in Nova Scotia is going, I might respond with nothing but the drawing at right, thereby communicating to you is that I am passing the summer in happy and idle recreation. But this is not what is depicted: what is depicted is me sailing a boat. However misleading, it is also compatible with the picture on its own that I have had a stressful, miserable experience (shortly after the moment depicted, the boat capsized and my dissertation sank to the bottom of the ocean).

In general, what a picture communicates on a particular occasion is independent of what it depicts. What is communicated can exceed the content of picture— as in the case above; or it can fall short— as when a picture is only partially viewed; or the two may be disjoint— when a picture is improperly viewed or interpreted. Of course, if such behavior accompanies *every* act of interpretation, then the content associated with pictures may change or dissolve. The associations of pictures with content inevitably depends on the interpretive practices of the population of artists and viewers in which the picture has currency; I discuss this matter later in the chapter. But content is never held hostage to the responses of specific individuals on specific occasions. It is this individual-independent content which I propose the semantic rules associate with pictorial signs.

I have pinpointed the concept of representational content presupposed by the Pictorial Semantics Hypothesis. But should we believe that pictures have such content? The alternative is the null hypothesis that pictures are ordinary, non-representational objects. According to this eliminativist view, pictures carry physical information in the same way that any physical object carries physical information about its causal sources. Thus waves carry information about the direction of the wind and smoke carries information about the location of fire. The eliminativst alleges that this the only sense in which pictures indicate the state of the world; by contrast, I maintain that they carry genuine representational content.

Unlike ordinary objects, representations— objects with representational content— construe the world as being a certain way. As such, they can be evaluated according to whether or not they are correct in the way they construe the world to be. For example, the fact that sentences have content

---

[15]The distinction between *what is communicated* and *what is depicted* by a picture is borrowed directly from Grice's (1975, pp. 43-4) distinction between *what is communicated* and *what is said* by a sentence.

is reflected by the fact that we can evaluate sentence for truth and falsity; truth and falsity are measures of the correctness of a sentence's representational content. By contrast, unadulterated sticks and stones cannot be evaluated for truth or falsity, though they may carry floral and geological information, because they are not bearers of representational content. It is in exactly this sense that pictures are like sentences, and unlike sticks and stones.

While it is admittedly awkward to speak of a picture being "true" or "false", the counterpart of truth-like correspondence with reality for pictures is *accuracy*.[16] It is perfectly natural to evaluate pictures for their relative accuracy and inaccuracy. This reflects the fact, I suggest, that pictures have content. It is in virtue of carrying content that pictures represent the world accurately or inaccurately. The concept of accuracy invoked here plays a central regulatory role in information exchange with pictures. For example, accuracy is the standard of pictorial fidelity which 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 which artists characteristically strive for, and viewers expect them to achieve. When the stakes are lowered, it is permissible for artists to infringe upon this standard proportionately. Accurate depiction is simply representation without error. Accuracy in this sense is not the exclusive denotation of the word "accuracy" in colloquial English, but it is one of them.[17]

It may help to clarify what accuracy is not. By accuracy, I do not mean mere *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.[18] 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.[19] Nor does accuracy imply *closeness to reality*, in any straightforward sense; a full-scale animated model of Obama is arguably more similar to

---

[16]Fodor (1975, pp. 180-1) argues that, even in principle, pictures cannot have truth-values.

[17]More or less the same concept of accuracy is widely invoked in discussions of pictorial realism (see, e.g. Lopes 1995; Abell 2007; and Hyman 2006, pp. 194-7). Yet such accounts typically aim to analyze the richer concept of *realism*, taking some form of (non-factive) accuracy as a merely necessary but insufficient condition. Accuracy also corresponds to what Walton (1990, §3.2) calls the relation of "matching".

[18]At any rate, this is *one* sense of the term "realism" among many. See Jakobson (1921/1987), Lopes (1995), and Hyman (2006, Ch. 9) for discussion.

[19]See Block (1983, pp. 655-8).

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, a point vividly demonstrated by Frank Lloyd Wright's architectural drawing of a *proposed* building plan with which we began.

For any given state of the world, a picture's content determines whether the picture represents that state accurately or inaccurately. In other words, it determines the conditions under which a picture is accurate or inaccurate. Such *accuracy-conditions* can in turn be used to characterize the content of a picture, just as *truth-conditions* are used to characterize the content of sentences. In §5, I'll provide a formal exposition of this same idea; and I'll use this analysis to define pictorial semantics, formally, as a mapping from pictures to accuracy-conditions.[20]

A final concern allows that pictures have representational content, but objects that it is not the sort of content required by a semantics. Instead, semantics essentially involves the definition of *truth*, and counterpart notions such as accuracy fall short of this criteria.[21] While accepting that pictures are neither true nor false, I reject this rigid truth-theoretic conception of semantics. Contemporary semantic theories are keyed to a variety of alternative standards of *correctness*.[22] I simply add accuracy to this roster. I conclude that pictures have representational content, and that this content is a suitable basis for a semantics of pictures.

3 Systematicity

I have argued that pictures are bearers of representational content. The Pictorial Semantics Hypothesis further holds that pictures are assigned their content by systematic rules.[23] The thesis that there are systematic mappings from pictures to content is widely doubted. The opposing view is that pictures themselves simply do not make systematic contributions to the information they happen to convey. Like winks or kicks under the table, pictures are thought to impose at most loose and defeasible constraints on interpretation.[24] The rest is left to rational inference and quirks

---

[20]Truth and accuracy appear to come apart in two ways. First, pictures and descriptions are accurate and inaccurate, while only sentences are true and false. Second, accuracy comes in degrees while truth is binary.

[21]The assumption, widespread among semanticists, is typified by Lewis' (1970, p. 18) statement: "Semantics with no treatment of truth conditions is not semantics." That said, it is clear that Lewis, if not others, had a quite ecumenical notion of truth. For example, in Lewis (1969), he defines the satisfaction conditions for imperatives in terms of truth. Clearly it is not natural to speak of imperatives as being "true" or "false"; but equally clearly, their content should be characterized in terms of a constraint on how the world must be.

[22]See Hamblin (1958), D. Kaplan (1999), and Starr (2010) for discussion of a variety of non-truth-conditional approaches.

[23]Here and throughout I use the term "systematicity" to denote the property of being systematic, as applied to interpretive schemata. This is distinct from the use of the same term by linguists in discussions of compositionality to denote a specific feature of language in which expressions of the same type can be inter-substituted while preserving interpretability. (Szabo 2010)

[24]Describing the parallel anti-semantic view of language, Stanley (2000, p. 396) uses the examples of the meaning conveyed by kicking someone under the table or tapping them on the shoulder; Borg (2006, p. 261) uses the example of raising one's eyebrow.

of psychological response; there is no *code* for interpreting pictures.

Tellingly, for the greater part of the twentieth-century, the same kind of anti-semantic attitude dominated philosophical views of language. Ordinary language was thought to be essentially unamenable to systematic analysis. Yet such skepticism ultimately gave way. Following the pioneering work of Chomsky (1965), Montague (1970*a*, 1970*b*) and others, theorists of language increasingly warmed to the idea that natural languages are basically rule-governed phenomena. They have even largely converged on the particular proposal that linguistic communication relies on *semantics*— systematic and conventional codes which arbitrarily assign basic meanings to simple linguistic signs, and then derive complex semantic content from strings of signs, by recursive rules of combination and context sensitivity.[25] This is an explanatory posit, not an article of faith. Humans are able to interpret sentences they have never encountered before, and they are able to do so with extraordinary speed and reliability. Such feats would not be possible, it is thought, if linguistic interpretation were governed solely by open-ended guess work. Instead, these phenomena are best explained by the hypothesis that humans have tacit knowledge of interpretive rules, and they interpret novel sentences by automatically applying these rules.[26] The plausibility of this hypothesis has increased as semantic theories for language have grown ever more detailed and precise.

I propose that essentially the same argument can be made about the interpretation of pictures.[27] The first key observation here is that humans are consistently able to extract content from pictures they have never seen before. For example, a cognitively normal consumer of contemporary global media would have no trouble interpreting the architectural drawing presented at the beginning of this chapter, even if she had never seen it before. She would not necessarily be able to identify the particular building depicted, but she would easily recognize the basic three-dimensional spatial and chromatic configuration the picture represents. The fact that we can all interpret unfamiliar images in this way demonstrates that picture interpretation is not based on a finite list of memorized correspondences. Instead, it stems from some general capacity which may be applied to new particular cases. But what is the nature of this general capacity?

The second key observation is that the interpretation of unfamiliar images is *fast* and *reliable*, despite the fact that both the stimuli and the recovered content are significantly complex. Pre-

---

[25]See King and Stanley (2005) for a canonical statement of this view.

[26]The phenomena described here is commonly termed "productivity." Linguists also argue for the same conclusion on the basis of "systematicity". (See footnote 23.) In either case, the details of the explanation are by no means straightforward. See Lewis (1975), Soames (1989), and Dummett (1996) for three very different perspectives on the matter.

[27]Here I will focus on considerations of productivity for pictorial interpretation. I do believe, however, that it is possible to mount arguments for pictorial semantics from phenomena that correspond roughly to the linguistic concept of systematicity. (See footnote 23.) For example, if a picture is interpretable, so is its mirror image; and certain parts of images can be shifted, rotated, and replaced while maintaining interpretability. I believe the conclusions suggested by this evidence are much the same as those argued for in the main text.

sented with our initial architectural drawing, normal media consumers would attribute to the image approximately the same spatial interpretation, and each would be able to do so nearly instantaneously. These facts are not easily explained by supposing that interpretation is based merely on open-ended speculation. That is the kind of reasoning we engage when choosing a chess move, or guessing the intentions behind an ambiguous nod. It is characteristically slow and variable, particularly when applied to complex stimuli— nothing like the automatic and consistent responses exhibited in picture interpretation. A better explanation is simply that the mechanisms at work here are governed by general interpretive rules. The speed and reliability with which humans associate unfamiliar complex images with content is explained by the systematic application of these rules. Viewers need not explicitly represent such rules themselves; it is sufficient that they apply certain systematic capacities in regular ways.

It is undeniable that pictures are used in communication; any theory of communication owes an explanation of how this transpires. The considerations reviewed here indicate that at least part of the explanation will be systematic. Positing systematic interpretive mechanisms is the best available explanation for our observations of human interaction with pictures. Of course, this hypothesis is open to empirical falsification. But since this argument is essentially the same as the argument offered for the existence of a semantics of language, the conclusion is just as robust. We should expect to find a semantics of pictures.

The extent and limits of this claim must be understood. First, the claim of systematicity should not be construed too broadly. The preceding argument shows only that a significant portion of pictorial communication is governed by semantic regularities. I do not deny that pragmatic principles play an important role as well. For example, I discussed earlier how an image of a sailboat, sent by mail, would reasonably elicit from its intended audience the judgement that the passengers depicted were especially happy and carefree. But this judgement cannot be warranted by the picture itself, since the passengers themselves are scarcely visible in the image; instead, it is partially justified by the audience's knowledge of the artist's likely intentions, informed by further knowledge of the social norms governing the communicative exchange. That is, the judgement is the result of unconstrained pragmatic reasoning; such imaginative engagement with pictures is ubiquitous. But note that, as is often the case, the pragmatic reasoning on display here depends crucially on first securing that content which is made available semantically. One cannot deduce that the passengers depicted are sailing happily, unless one first appreciates that the image depicts a sailboat, and even more basically, a certain boat-shaped arrangement of edges in spaces. The extraction of this last dimension of content is strictly determined by the structure of the image itself and the rules of in-

25

terpretation, and is not the result of pragmatic speculation. It is a substantive and difficult question exactly how much content typically associated with pictures should be counted as purely semantic. My contention here is only that some significant portion is, and in Chapter 3 I will make this idea precise.

The claim of systematicity must not be understood too rigidly either. An instrumental part of modern natural language semantics was brought on by the development of the theory of CONTEXT-SENSITIVITY. Context-sensitivity arises when the content of a linguistic expression depends on the context in which it is uttered, but in ways that are strictly constrained by the underlying grammar.[28] Perhaps the simplest example in English is the indexical "I", the referent of which varies with speaker, and thus context, but in a manner that is obviously rule-governed. (The grammar requires that "I" refer to its utterer, whoever that may be.) Cases like this illustrate an important moral. The mere fact that the content of a linguistic expression depends in part on context does not imply that such content is not determined by a semantics. Instead, the semantics may associate expressions with meanings which systematically regulate the contribution of context. The same lesson applies to pictorial representation. Given the considerations provided above, we should expect the mapping of pictorial signs to content to be systematic. But as we shall presently see, such content depends in part on context. This does not imply that pictorial representation has no semantics; it suggests instead that the semantics is context-sensitive.

Consider a series of drawings, each produced under different conditions, yet all of which coincidently turn out to be qualitatively identical to the accurate depiction of Obama presented above. One depicts Obama; one depicts a wax sculpture; a third depicts a chance configuration of lines in the sand. In so far as the concepts *Obama*, *wax-sculpture*, and *lines in the sand* are part of— or play an essential role in specifying— the content of each image, such content must be context-sensitive. Since the structure of the image is the same in each case, only the variation in their context of expression can account for the variation in their content. Yet, as the context-sensitivity of language demonstrates, these facts do not imply that the content of each image is *entirely* context-dependent, hence unsystematic.[29] Indeed, it is quite obvious that each scene depicted shares some basic spatial configuration, specified by the structure of the image itself. It is this core component which, I propose, the semantics of pictures directly associates with image structure; the rest is contributed by context.

Following precedents in linguistic theory and recent work in the philosophy of depiction, I

---

[28]The classical source of this view is David Kaplan (1989). See Stanley and Gendler-Szabo (2000) for more recent elaboration.

[29]Balog (2009, §2) presents such an argument against pictorial semantics.

propose that a pictorial semantics associates each picture with two levels of content: a primary, EXPLICIT level of content is context-invariant and sparsely geometrical; a secondary, EXTENDED level of content is context-dependent and richly referential.[30] The explicit content of an image is that minimal spatial and chromatic profile of visible surfaces which can be derived from the structure of the image alone. (In Chapter 3, I develop a precise theory of explicit pictorial content.) The extended content of an image elaborates the explicit content with reference to individuals, properties, and richer spatial features. It is extended content that is most naturally grasped when we talk about, say, the picture depicting *Obama* as having a certain kind of *hair cut*.

But extended content always depends constitutively on explicit content. For example, it is part of the extended content of the Obama picture that the man's hair has a very specific outline shape, one which we struggle to articulate in words. It is this kind of purely structural feature which is supplied by the explicit content of the image. Explicit content also functions to constrain the way that context may enrich a picture's representational properties. A single image type might accurately depict Obama, or lines in the sand. But the same image could not accurately depict a cube or a plane, no matter the circumstances of its creation. The image's explicit content places hard constraints what the image may be used to accurately depict. It constitutes the core of pictorial representation.

In this and subsequent chapters, I directly defend the systematicity of pictorial semantics only with respect to explicit pictorial content. That is, I defend the claim that there is a context-invariant component to pictorial content which is associated, by systematic rules, with image structure. Any theory of extended content, which takes into account the influence of context, necessary depends on a prior theory of explicit content. This dissertation is restricted to the more basic task.

Despite the evidence that pictures are systematically mapped to their content, and the appropriate qualifications of this claim, the very possibility of such a mapping is widely doubted. Language, it is thought, is the paradigmatic site of systematic semantics. It is clear that language provides a poor model for pictorial representation, but this fact is taken to impugn the viability of a pictorial semantics. On this view, there can be no semantics outside of language. Three kinds of concern motivate this outlook.

---

[30]This idea has two sources. The suggestion that pictures have two levels of content— a thin, geometrical level, and a rich, referential level— is due to Haugeland (1991, pp. 73-5), who distinguishes between "bare-bones" and "fleshed-out" pictorial content; this idea is developed extensively by Kulvicki (2006*a*, p. 59). (Siegel (2006, p. 2) discusses analogues for perceptual content.) The suggestion that a sign's context-sensitive content can be derived from a thin level of context-invariant meaning, together with context is due to David Kaplan (1989). Kaplan theorizes that the grammar associates each lexical entry with a "character" (also called a "standing meaning"), which, together with context, determines "content". Explicit content in my sense is the pictorial analogue of Kaplan's character; extended content in my sense is the analogue of Kaplan's content. Thanks to Matthew Stone for suggesting the term "explicit content".

(i) First, it does not seem that pictures have syntax in the manner of linguistic expressions. Semantic rules are typically understood as systematic mappings from syntactic structures to contents. But without syntactic structures to take as inputs, it seems that a semantics of pictures could never gain firm footing. Yet this objection misunderstands the general character of semantics. In the most general sense, it is *structure*, not logical syntax, that is the point of departure for semantics. Thus I grant that pictures do not have a *syntactic* structure like that of linguistic expressions, and they may not be governed by substantive rules of syntactic well-formedness.[31] Yet in Chapter 3 I will develop a view according to which pictures have *geometrical* structure instead. The semantics of pictorial representation associates geometrical structures with representational content using systematic techniques of geometrical projection.

(ii) A second concern has to do with compositionality. Compositionality is often understood as the recursive derivation of the meaning of a complex expression from the meanings of each of a finite number of atomic parts, along with the way these parts are put together. But pictures do not seem to be composed from a finite set of atomic parts. Hence their content cannot be derived compositionally; and without compositionality, it is wagered, there is no semantics.[32] But the worry is unfounded. Compositionality, in its most general construal, is the derivation of the content of a complex expression from its internal structure— it need not require that this structure is decomposable into finitely many atomic parts. In the formalization presented at the conclusion of Chapter 3, I demonstrate that the proposed geometrical semantics for picture is fully compositional in this general sense.

(iii) The last objection is purely inductive. The only semantic theory which is well-understood is that of language. There have been no other correspondingly plausible candidates for a systematic account of depiction. And it is fundamentally unclear what such a system of pictorial representation would be. If there were a semantics of pictures, it is thought, such an account would have been established by now. The preceding discussion has amounted to a defense of the very idea of a systematic pictorial semantics. But I have not proven that any such semantics exists. Until a systematic theory of depiction is actually articulated, this last form of skepticism will reasonably persist. It is this concern which Chapters 2 and 3 are intended to assuage. The remainder of the present chapter continues to establish the necessary groundwork for subsequent development.

---

[31] For example, it might be that a picture belongs to the system of black and white line drawing just in case it consists of (a) black lines and (b) white regions. This amounts to a rule of well-formedness, but not a substantive one.

[32] Thanks to Jeff King and Mark Baker for pressing this concern.

4 Conventionality

I have argued that pictures are assigned representational content by systematic rules. The Pictorial Semantics Hypothesis further holds that these rules are conventional. There are a diversity of interpretive systems for pictures, varying across history and culture, which are selected for use by social convention, just as languages are. The opposing view is that the ability to interpret pictures is simply a manifestation of the biologically fixed abilities which make up human visual perception. If there were a theory of pictures, it would be a byproduct of the science of vision, not a topic of independent interest.[33] On this view, conventional knowledge plays no role in the interpretation of pictures.

Despite its naïve plausibility, this deflationary hypothesis is now widely discredited. As early as 1921, the linguist Roman Jakobson observed, "it is necessary to learn the conventional language of painting in order to "see" a picture, just as it is impossible to understanding what is said without knowing the language." (1921, p. 21) Since then, theorists of pictorial representation have documented a diversity of pictorial conventions, or SYSTEMS OF DEPICTION.[34] Such systems crucially mediate between pictures and the informational content they carry.

Consider, by way of illustration, the pair of images at right. The first contains spatial as well as chromatic information about the scene featuring my plant— e.g. that the *cone-shaped* object is a particular shade of *brown*. Part of what determines its chromatic content is of course the color of the pigments in the picture itself. If these colors had been different, the picture would have had a different content. By contrast, the second image is a black and white line drawing of the same scene. It has the same spatial content as the first drawing, but contains no chromatic information about the scene it depicts. If it *were* a color drawing, it would depict a plant, pot, and desk which were all paper-white; plainly this is not the case. So while both pictures have content, the relationship between the surface properties of the picture and its content is different in the two cases. The explanation for this is that different systems of depiction, like line drawing versus color drawing, determine different ways of mapping pictures

---

[33]This need not imply that seeing a picture of a scene is the same as seeing the scene itself. After all, black and white drawings look nothing like the scenes they depict. Instead, it is held that the mechanisms which ground picture interpretation are fully determined by (but not necessarily the same as) those which ground normal visual perception.

[34]Three excellent books on the topic are *Art and Illusion* (Gombrich 1960), *Varieties of Realism* (Hagen 1986), and *Art and Representation* (Willats 1997). The term "system of depiction", or a close cousin, is widely used (see, e.g. Manns 1971, p. 286; Newall 2003, p. 384; Abell 2009, p. 222).

to content.

Differences between systems of depiction are not limited to variation in the treatment of color; they differ in their handling of line, shading, and texture as well. Importantly, they also determine divergent treatments of pictorial *geometry*. Consider the two drawings below; each depicts a cube, but in a different way. In the drawing on the left, *converging* lines in the picture are used to depict the parallel edges of the cube. In the drawing on the right, *parallel* lines are used to depict the parallel edges of the cube. The first is rendered in a PERSPECTIVE system, the latter in a PARALLEL system.[35] If the drawing at right *were* in perspective it would depict an irregularly shaped solid— not a cube at all. Thus, like line and color drawing, these two systems determine different relations between features of the picture surface and its content.

PERSPECTIVE   PARALLEL

Far from artificial constructs, these systems of depiction and a vast diversity of others have been widely used across history and culture. Classical Chinese painting, for example, employed a type of parallel depiction; another dominated ancient Egyptian imagery, with very little variation for over 2,000 years; meanwhile, Western post-Renaissance depiction has typically been presented in perspective.[36] Notably, parallel depiction systems like those used by classical cultures are the norm in contemporary engineering and architectural drawing. All such systems exhibit a characteristic stability, whereby artists and viewers may continue to exploit the same system across time and in varying contexts.

These facts cannot be explained by appeal to the normal workings of the perceptual system alone. The activity of the perceptual system is widely recognized to be cognitively "impenetrable". That is, its internal processes are in an important sense automatic; they cannot be modified through the interference of rational inference or conscious decision making. The visual system is also believed have essentially the same structure in all human populations. By contrast, the use of systems of depiction varies widely across human populations, and is depends heavily on contingencies of

---

[35]The class of parallel systems includes the more familiar and specific systems of oblique, isometric, and orthogonal projection. Orthogonal projections include plan and elevation depictions common in architecture and engineering. The drawing above is an isometric projection.

[36]On classical Chinese depiction, see Willats (1997, p. 55). On classical Egyptian depiction, see Hagen (1986, pp. 168-72).

social context. Thus the interpretation of pictures cannot be explained by appeal to biological vision alone.

At the same, it is evident that those systems of depiction used by human societies depend in crucial ways on the workings of the human visual systems. The interpretive computations required by systems of depiction are broadly of a kind with those performed by the human visual system. Overlapping cognitive resources are called upon in normal vision and the interpretation of pictures. This idea is made concrete in Chapter 3. Because of their perceptual origins, the space of humanly realizable systems of depiction is indeed constrained by universal features of human vision. Not anything can become a human system of depiction.

On the other hand, the conventionality of pictorial representation frees it from the strictures and ideosyncracies of vision. Systems of depiction introduce expressive devices never found in perceptual experience. These include outlines, alternative color schemes (like black and white, or negative color), color coding, and alternative projection systems (like parallel projections or map projections), to name just a few. Such departures from normal vision demand interpretive familiarity. Systems of depiction also determine objective standards of accuracy that do not depend on the visual impairments or special abilities of individuals. This facilitates reliable communication. In this way, systems of depiction are both conventional crystalizations of visual perception and extensions of it.[37]

Thus systems of depiction are the pictorial analogues of languages: culturally specific interpretive rules for encoding content in signs. Are systems of depiction therefore conventional in the same way that languages are? They are. Although systems of depiction have a very different internal architecture from languages, the use of systems of depiction in human communication is conventional. More carefully: there are infinitely many possible systems of depiction, considered as abstract interpretive rules; among these, a select few are selected by social convention for use in human communication, just as languages are.

Following Lewis' (1969) influential account, conventions of interpretation, such as languages, arise when a population wishes to communicate, but more than one methods of communication is available for use. In order to successfully communicate, the population must collectively coordinate on one such method for repeated use, thereby establishing a convention. Systems of depiction are conventional in exactly this way. Consider a concrete example: the engineers in a construction firm must share drawings amongst themselves while designing and building a bridge. It is important to

---

[37]See Gombrich (1960, pp. 3-30), Kennedy (1974, pp. 24-7), Lopes (2009, pp. 18-9), Durand (2002, §4), and DeCarlo and Stone (2010, p. 5).

everyone that all the drawings produced within the firm be accurate and be accurately interpreted, since inaccuracy in production or interpretation could have disastrous effects. The only reliable and efficient way to achieve this goal is to establish a standard system by which drawings are produced and read. That is, the engineers must collectively establish a convention of producing and interpreting drawings according to a given system of depiction.

The same phenomena is repeated, at a larger scale, in contemporary media. For example, the use of perspective line drawing is a global convention. Here, the coordination among artists and viewers is tacit and disparate. Nevertheless, the basic considerations are the same: all seek to efficiently communicate approximately accurate information; yet there are always multiple ways this might be achieved; communicative success is ultimately realized by coordinating on a stable interpretive convention for pictures. Of course, contemporary media relies on a great variety of systems of depiction. Just as many human societies are multi-lingual, human societies can and do use a multitude of systems of depiction. Those which are used recurrently become conventions.

In spite of the considerations reviewed here, many have thought that systems of depiction *could not* be conventional, even in principle. Paralleling the debate about systematicity, objections here tend to take languages as the paradigm case of conventionality, and argue that, because of their disanalogies with language, systems of depiction could not be conventional in the same way. Dispelling such concerns will help elucidate the central claim of this section. I focus on three objections in particular.

(i) If systems of depiction were conventional, it is thought, they would be arbitrary in the way that languages are. Yet while words bear a merely arbitrary relationship to their content, pictures bear a much more intimate and direct relationship to theirs. Thus systems of depiction cannot be conventions. This objection trades on a simple but widespread confusion between the sort of arbitrariness associated specifically with language, and that which is a feature of all conventional behavior.

It will be helpful here to recall a key feature of David Lewis' influential analysis of convention. Lewis observed that for a regularity of behavior within a population to count as a convention, there must be some alternative regularity with which the population *could* have conformed, while still realizing their dominant needs. Part of what makes a regularity a convention, then, is that it reflects a *choice* (not necessarily a conscious one) on the part of the group, between mutually acceptable options. For example: driving on the left is a convention in some nations, in part because drivers *could* have collectively driven on the right instead, while still realizing their dominant need to pass each other while driving.

32

The sense in which all conventions are *arbitrary* is simply this: a population which conforms with a given convention could have conformed with a different convention, while still realizing their dominant goals. The use of, say, perspective line drawing by some population is a convention, in part because that population could have used parallel projection line drawing for the same purposes instead; in that sense the use of perspective line drawing by that population is arbitrary. In the same way, the use of English by some population is a convention, in part because the population could have used Shona for the same purposes; in that sense the use of English is arbitrary. But this kind of arbitrariness is not unique to systems of symbolic representation.[38]

The other kind of arbitrariness *is* unique to systems of symbolic representation. This is the kind displayed by the arbitrary link between the word "tree" in English and the property it denotes of being a tree; there is no intrinsic connection between the sign and its content. The same kind of arbitrariness is *not* displayed by the link between a picture of a tree and the tree it depicts. Since there are non-symbolic systems of representation, such arbitrariness is not a necessary feature of conventional systems of representation. And it is not even sufficient. As Fodor (1975, p. 178) has observed, if there are mental symbols, they bear an arbitrary relation to their content, but the relation is not determined by convention. I discuss this kind of arbitrariness at greater length in Chapter 4, in my analysis of symbolic representation.

Thus it is essential to distinguish one kind of arbitrariness, *in a group's choice to use a particular system of representation*, from arbitrariness *in the relation between a sign and its content within a system*. The former is a necessary condition on conventionality; the latter is specifically a feature of language, compatible with, but not necessary for conventionality. "Arbitrary" here has simply acquired two related but independent meanings. Public systems of depiction and public languages are both selected for use, by populations, conventionally. But of the two, only public languages rely on arbitrary pairings of signs and contents. A simplified example may help to make this point more vivid.

Suppose you and I must communicate our credit card numbers in a crowded room. To guard our privacy, we must come up with some kind of code. Here are four codes we might use:[39]

---

[38]Eco (1979, pp. 189-92) is one of the few authors to make this point explicitly.

[39]Here I put the message expressed in quotation marks, to distinguish it from the numeral thereby indicated. Strictly speaking, it is probably most apt to think of both relata of the code as numerals, not numbers.

+1 Code	+2 Code	A-List Code	B-List Code
"1" means 0	"2" means 0	"9" means 0	"2" means 0
"2" means 1	"3" means 1	"3" means 1	"5" means 1
"3" means 2	"4" means 2	"4" means 2	"4" means 2
and so on...	and so on...	"7" means 3	"8" means 3
		"5" means 4	"9" means 4
		"2" means 5	"6" means 5
		"8" means 6	"1" means 6
		"1" means 7	"3" means 7
		"0" means 8	"7" means 8
		"6" means 9	"0" means 9

The +1 and +2 codes exploit simple computations to disguise their message; other such codes could employ other once-place mathematical functions like multiplication or exponentiation by a constant. The A-List and B-List codes employ distinct, arbitrary pairings between numerals. There are as many such possible codes as their are possible distinct sets of pairings. Let us suppose that all four codes described here are sufficiently secretive for our purposes.

If the problem of communicating with codes is a recurrent one, it is convenient to establish a convention; the convention dictates which of the infinitely many possible codes (including the four above) we are to select for use on subsequent occasions. In the first sense, the selection would be arbitrary. Because all four codes suit our purposes equally well, the choice to use the +1 code instead of the +2 code or the A-List code is an arbitrary choice. But now consider that the A-List and B-List codes are *arbitrary* in a way that the +1 and +2 codes are not. They rely on arbitrary pairings of input and output, whereas the others rely on mathematical laws. Clearly, then, arbitrariness in the selection of a code for use is quite a different thing from arbitrariness in the structure of the code itself.

The lexicons of natural language are the analogues of the A-List and B-List codes. They are conventions of interpretation which are based on arbitrary pairings of signs and contents. Systems of depiction are the analogues of the +1 and +2 codes. They are possible conventions of interpretation which are based on law-like associations between signs and contents. Although I have said little thus far about the internal structure of systems of depiction, in Chapter 3 I develop a theory according to which they are based on rules of geometrical projection. For present purposes, such projective rules are merely complex cousins of the +1 and +2 codes. The general moral is this: conventionality requires that the agents exploiting the conventional rule have made an arbitrary selection of that rule over another. But conventionality makes no claim about the nature of this rule. It may be arbitrary, like a lexicon, or law-like, like a geometrical projection. Systems of depiction, as it happens, are often deployed in communication by convention.

(ii) A second objection holds that, for all we have just said, languages must be *more* conventional than systems of depiction. The evidence is this: pictures produced according to one system of depiction can typically be successfully interpreted by viewers familiar only with other, even very distant systems.[40] For example, a viewer familiar only with systems based on the geometry of *perspective* projection would be able to recover the spatial information from the picture at right, even though it is the product of a system based on *parallel* projection. By contrast, sentences from one language are not generally comprehensible to speakers of another language. Thus languages must be more a product of convention, and systems of depiction less so.

The argument relies on an implicit false premise: although viewers may manage to recover the *crude* content of pictures from unfamiliar systems, it is never possible to determine with certainty the *exact* content of a picture without acquaintance with the relevant system. Thus, without knowing whether the picture above belongs to a perspectival system, or a system with some alternative geometry, it is impossible to tell whether the object represented is in fact a cube, or a related but irregular solid. Conventional knowledge is required, after all, for correct interpretation.

What the objection *does* reveal is that the space of humanly implementable systems of depiction is narrower than the space of humanly implementable languages, restricted as it is by the universal structure of the human visual system. As a consequence, human viewers can generally make educated guesses about the interpretation of pictures in unfamiliar systems. Yet this difference between languages and systems of depiction is only one of degree. Genetically determined neural anatomy put severe constrains on the space of languages that can be used by humans, as well.[41] In both cases, restrictions are imposed by the biological structure of the interpretive anatomy.

Ultimately, it is immaterial to the conventionality of a given system of representation whether it belongs to a wide or narrow space of humanly realizable alternatives. The conventionality of a given system requires only that there exist *some* alternative. Driving on the left, for example, may be a convention, even when there is only *one* possible alternative which satisfies the goals of population— driving on the right. Drawing in perspective, like driving on the right, or speaking French, is a matter of social convention.

(iii) A third and final objection, due to Lopes (1992, §3-4; 1996, §6.6), contests the claim that systems of depiction are conventional, on the grounds that the choice to use one system of depiction over another is rarely arbitrary. For example: drawings from systems of black and white depiction

---

[40]See Carroll (1999, pp. 42-4).
[41]See Pinker (1995).

and those from systems of color depiction reveal different features of their subjects, and require different technologies to produce. The choice to use one kind of system or the other in a given communicative exchange is invariably motivated by the practical demands of the task at hand. We might specifically wish to communicate color information, or specifically wish to suppress it. But, it is argued, whenever the choice to use one system or another has a specific rationale that use is no longer based on convention.

The objection rests on the principle that conventions only arise when the population in question is truly indifferent between alternative regularities in behavior, so that the choice between them is entirely arbitrary. But the principle is false. It is true that, for a regularity in behavior to be a convention, the population must select that regularity from a space of alternatives. But the selection need not be unmotivated. By way of illustration, consider a population for whom, through a quirk of genetics, driving on the left side of the road is discomfiting, while driving on the right side of the road is relatively pleasurable. (And not driving at all is still more pleasurable.) It is a matter of common knowledge throughout the population that this is so. Nevertheless, if all members of the population drove on the right, and expected all others to do the same, it would still be a convention, despite the fact that it is the mutually preferred alternative. It would be a convention, in part, because all would have been prepared to drive on the left, if necessary, in order to pass each other on the road. Thus a regularity in behavior can be a convention, despite the fact that it would be the preferred behavior quite apart from convention; also a regularity can be a convention even if it would be dispreferred. Conventionality does not require indifference.

In his analysis of convention, Lewis (1969) is sensitive to these facts. Rather than requiring that a populations be indifferent between alternative regularities, Lewis imposes a subtler, and more plausible requirement. For a regularity to be a convention, it must be one among a set of alternatives which are such that each member of the population prefers conformity to that regularity *on the condition that everyone else conforms*. Thus, I prefer to speak English if everyone else speaks English; I prefer to speak Shona if everyone else speaks Shona. This condition can be met, even if the alternatives are preferentially ranked. All things being equal, I may prefer to drive on right; but this preference is trumped by the desire to drive in the manner of everyone else. This is sufficient to qualify the regularity for conventionality.

In the original case of depiction, the dominant goal of the typical communicative exchange is to communicate accurate information using a picture. Thus, while it may be that, all things being equal, I prefer to produce and view color drawings, I prefer most to use the same system of depiction that everyone else is using. Under such conditions, a system of black and white depiction

36

or one of color depiction could become the convention. Thus the conjecture that the use of any given systems of depiction is often or always motivated by practical concerns does not impugn its conventionality.

Together, the last three sections constitute my defense of the Pictorial Semantics Hypothesis. I have argued: (i) that pictures are bearers of representational content; (ii) that the mapping between pictures and their content is systematic; and (iii) that there are a diversity conventional systems according to which pictures are associated with their content. The Pictorial Semantics Hypothesis can now be more precisely construed as the hypothesis that every actually occurring system of depiction determines a systematic and conventional mapping from pictures to their content. The aim of the next section is to render this claim with sufficient formal precision that we may reliably evaluate particular theories of pictorial semantics.

5   Accuracy-conditional pictorial semantics

The aim of this section is to give formally explicit expression to the Pictorial Semantics Hypothesis, by defining the concept of pictorial content adumbrated above in terms of its accuracy conditions relative to a system of depiction. This accuracy-conditional framework will set the stage for the assessment and development of specific theories of pictorial semantics in the next two chapters.

We begin by developing a formal model for depictive content that abstracts away from the distinguishing features of different systems of depiction.[42] There are a number of ways to proceed here, but following Blumson (2009*b*), a natural course is to extend the popular possible-worlds theory of linguistic content to the case of pictures.[43] On this view, since a picture represents a certain way things might be, its content can be modeled by the set of possible worlds so represented.[44] But this is too crude: a world could be just as a picture represented it at one *time*, but not at another.

---

[42]There are two sorts of theory of content. One seeks to place purely *formal* constraints on the structure of content, but remains silent on the metaphysical nature of the objects realizing these constraints. This is the norm in linguistic theory. (See Lewis (1970, p. 32) for a brief but apt defense of this norm.) The heroic few (e.g. King 2007) take on both the formal and metaphysical questions. In this paper I am concerned only with the formal characterization of pictorial content.

[43]For a canonical statement of the possible-worlds account of content, see Stalnaker (1984). The application of this theory to pictorial content has been proposed by Howell (1974), Ross (1997, pp. 17-21), and Blumson (2009*b*). Alternatively, we coud define the content of a pictures as a complex property (Lopes 1996, p. 111) or as a structured situation (Malinas 1991, p. 276). Further proposals might be usefully adopted from the theory of linguistic content, perceptual content, or the content of mental imagery.

[44]Does the existence of "impossible" pictures like the Penrose triangle or the illustrations of Escher require that pictorial content also include *impossible* worlds? (See Huffman (1971) and Sorensen (2002) for examples.) It is known that all familiar impossible pictures can be divided into possible parts. (Kennedy 1974, pp. 149-50; Sorensen 2002, p. 346) On the strength of this observation, I deny that impossible pictures have content, while allowing that their possible parts are individually contentful, though mutually incompatible. Thus the pressure to countenance impossible worlds is diffused. Blumson (2009*b*, §7), for one, opts instead to include impossible worlds in pictorial content.

Similarly, it could be just as a picture represented it at one *location*, but not at another. It is therefore inappropriate to think of pictures as representing the way the world is across the board; instead we should think of them as representing worlds at particular times and locations.[45] (For now I choose to ignore the further matter of a picture's *angle of view*. I will return to this issue in Chapter 3 once we have developed a richer understanding of the mechanics of perspective.) We can think of these particular locations and times within worlds as *scenes*. Pictures do not, as adverted before, represent the entire world as being a certain way— they represent a more limited scene as being a certain way. Thus the content of a picture is modeled as a set of scenes. We may define a SCENE simply as a ⟨*world, location, time*⟩ triple, what is sometimes called a "centered world."[46]

Depictive content is modeled by *sets* of scenes, rather than individual scenes, because pictorial representation is always indeterminate.[47] In the system of black and white line drawing, for example, the picture of Obama does not represent Obama's suit as being any particular color. So the picture is compatible with worlds in which the suit is blue, as well as worlds in which it is grey or red. While color systems are not indeterminate in this particular way, nearly all systems are indeterminate with respect to *occluded surfaces*— that is, those surfaces obscured by some more nearby surface. Thus the picture of Obama does not represent Obama's *back* as being any particular way; it is compatible with worlds in which Obama has wings, occluded from view, and worlds in which he doesn't. In such cases the picture fails to specify a complete description of the scene before it, hence the characterization of its content as a set, rather than a single scene.

Though the content of a picture may always be modeled by a set of scenes, *which* set of scenes must depend on the operative system of depiction. For example, the set of scenes compatible with the content of a color drawing is more restricted chromatically than those compatible with its black and white counterpart. A quite natural suggestion is to define pictorial content in terms of DEPICTION— understood as a two-place relation between a picture and a scene, specified relative to a system of depiction.[48] The idea is that a single picture may depict a range of possible scenes— for example a single black and white line drawing may depict a chromatic variety of scenes— but the placement of lines in the picture ensures that every scene the picture depicts has the same basic configuration of objects in space. Taken together, the set of scenes which a picture depicts in a

---

[45]See Blumson (2009*b*, §2).

[46]The original idea is due to Lewis (1979). For our purposes, it doesn't matter much how we think about possible worlds, though conceptions of worlds as constructed from linguistic objects may introduce certain difficulties.

[47]See Lopes (1996, §6.1) and Ross (1997, p. 18).

[48]A different approach, pursued by Ross (1997) and Howell (1974) is to characterize pictorial content indirectly, by considering what it must be in order to satisfy true sentences about pictures. This kind of consideration has its place, but my primary interest here is in pictures, not language about pictures. Still another strategy is developed by Malinas (1991, pp. 288-9), who defines picture content in terms of what "queries" a picture answers with "yes", "no", or "neither". It is hard to evaluate this proposal, because such terms are so alien to the concepts we ordinarily use in thinking about pictures.

system defines the picture's content in that system.[49]

My only complaint here is that the term "depicts" can be used ambiguously to track either pictorial reference or pictorial content. What we need to specify is a depictive relation that holds between all and only the scenes which conform with a picture's content, independent of its intended referent. This can be isolated by focussing on the *accurate* or *successful* depiction of a scene. Thus the spiky-hair image of Obama *accurately* depicts all and only scenes with spiky-hair-shaped objects of the appropriate sort— though none of these scenes are the actual referent of the picture. The definition of pictorial content can then be expressed unambiguously, for any picture *P* and system of depiction *I*:[50]

(2)   The content of *P* in *I* = the set of scenes *S* such that *P* accurately depicts *S* in *I*.

The content of a picture, according to this view, is its *accuracy conditions* in a system. Note that although accuracy can come in degrees, my interest here is in *maximal* or *perfect* accuracy. The content of a picture is the set of scenes it perfectly accurately depicts.[51] Here and throughout, by "accurate depiction" I mean perfectly accurate depiction, or very near it.

The concept of accuracy at play here is that which is required to define pictorial content in the way specified by (2). But it is more than that: it is the same concept which figured in my initial defense of the idea of pictorial content in §2. As I argued there, not only do we easily make direct and robust judgements of accuracy, it is the standard of fidelity to which artists aim, and viewers expect them to achieve, particularly in high-stakes contexts of communication with pictures. In short, accuracy is the basic measure of success for pictorial representation.

I have defined the content of a picture in a system as the set of scenes or "conditions" it accurately depicts in that system. I thereby join a formidable rank of philosophers and linguists to characterize the content of a variety of representational kinds in terms of some type of *correctness* conditions.[52] Most prominently, natural language semanticists in the twentieth century have overwhelmingly specified the content of assertoric sentences in terms of *truth*-conditions. The same

---

[49]This way of defining content is implicit in Blumson (2009*b*).

[50]An accuracy-conditional definition of content is briefly considered in Haugeland (1991, p. 78) and Blumson (2009*b*, §3). Appropriately enough, parallel accuracy-conditional accounts of content has been proposed for perceptual experience. (Siegel 2009, §2; Burge 2010, pp. 34-42)

[51]It seems to me that a definition of perfect accuracy is *all* one needs to specify the content of a picture. Part of the content of the spiky-hair image of Obama is that Obama has precisely *that* shape of hair— not close-cropped hair, not even hair which is slightly less spiky, nor hair which is slightly more spiky. For any hair-shape not perfectly accurately represented by the picture, that shape is not part of the content of the picture. Methodologically, this is good news, for to specify the representational content of a picture, a semantic theory need only supply necessary and sufficient conditions for perfect accuracy without involving itself in the dalliances of gradability.

[52]See Stalnaker (1998), D. Kaplan (1999), Gunther (2003, pp. 4-6), and Burge (2010, pp. 38-9) for discussion of the connection between content and correctness conditions. Once again, I mean to include in this group those theorists who refuse to *identify* some signal as its correctness conditions— for example defenders of Russellian and Fregean theories of content— but who nevertheless accept that their preferred notion of content must determine correctness conditions.

kind of approach has been fruitfully extended to questions, commands, and even perceptual experiences.[53] In general, observing the connection between content and conditions of correctness has been the key to the development of formal semantics as a precise and progressive research program; here I follow this precedent.

The Pictorial Semantics Hypothesis can now be formally articulated as the claim that, for every system of depiction, that system determines a systematic mapping from pictures to accuracy conditions. A theory of pictorial semantics for a given system is therefore a theory of accuracy in that system. Given a system of depiction I, such a theory must complete a schema of the following form, for any picture P in the system I, and scene S:[54]

(3) P accurately depicts S in I if and only if _____.

Ultimately, the empirical plausibility of any semantic theory depends on the specific predictions it makes in particular cases. The primary import of the accuracy-conditional schema presented here is to provide a standard format by which the predictions of such theories can be expressed, and their success rigorously measured. It provides the basic framework for the development of a pictorial semantics in the next two chapters.

6   The scope of inquiry

In this chapter I have defined and defended the hypothesis that there are pictorial semantics, systems of depiction which determine systematic and conventional mappings from pictures to content. In the next two chapters I take up the question of what sort of semantics pictures actually have. While the final aim of any pictorial semantics is to complete the semantic schema articulated above, I will limit the scope of the following inquiry in two important ways.

First, philosophers of language distinguish between FOUNDATIONAL and DESCRIPTIVE semantic theories, each necessary components of a complete theory of content.[55] Foundational semantic theories attempt to explain, from the metaphysical ground up, how inanimate concrete objects come to have representational properties.[56] Such theories are characteristically concerned with

---

[53]The truth-conditional approach has dominated twentieth-century semantics. Contemporary theorists in the dynamic tradition have added to this a variety of "update" conditions; see, for example, Groenendijk and Stokhof (1991) and Veltman (1996). In the realm of non-assertoric sentences, Hamblin (1958) proposes that the content of questions be defined in terms of "answer-hood" conditions; see Groenendijk and Stokhof (1997, p. 24) for further discussion. And see Starr (2010, Ch. 4) for a recent non-truth-conditional analysis of imperatives. Perceptual content, for its part, is often defined in terms of accuracy conditions. (Siegel 2009, §2)

[54]This schema is based on Tarski's (1944, p. 344) T-schema for truth.

[55]See Stalnaker (1997, pp. 535-6) and Stanley and Gendler-Szabo (2000, pp. 220-6).

[56]This is the sort of theory taken up by Putnam (1975; 1981), Field (1978), Dretske (1983), Stalnaker (1984), and Fodor (1990*a*; 1990*b*). See Loewer (1997) for a useful overview.

how an intelligent agent's interactions with an object, in a social context, manage to imbue that object with representational content. Descriptive semantic theories, by contrast, simply assume a class of objects with representational properties, and aim instead to explain why these representations have the content that they do. Such theories are the norm in linguistics, which aims only to articulate the formal rules which map linguistic signs to their content. This is typical in the philosophy of depiction as well, where the dominant theories, such as resemblance and perceptual accounts, are primarily distinguished by their approaches to the descriptive analysis of pictorial content.[57] My ambitions in this paper are the same. I'll have nothing to say about what separates accidental marks on a piece paper from genuine pictures. Instead, I'll *begin* with the class of pictures, and proceed directly to a descriptive account of the rules by which they are associated with their content.

Second, as I indicated above in §3, the discussion that follows is only concerned with the context-invariant, or explicit, component of pictorial content. While a final theory of pictorial representation must also contend with the context-sensitive, extended content of pictures, any such theory necessarily depends on a prior account of explicit content. So it is there that I begin. Henceforth, when I speak of "accurate depiction" I mean accurate depiction with respect to explicit content alone. In this sense, the same picture may accurately depict Obama, a wax sculpture, or lines in the sand.

Eventually, a theory of pictorial semantics must come to grips with both the foundational and the context-sensitive aspects of depiction. But we are still very far from that achievement, and there is a danger of seeking completeness at the cost of rigor. In this dissertation, I propose instead to purse a rigorous account of the central, context-invariant semantic architecture of certain common systems of depiction. I hope this will provide the firm bedrock which will ultimately be necessary for a systematic account of all aspects of pictorial representation.

---

[57] Abell and Bantinaki (2010) provide a clear discussion of this point, in the context of philosophy depiction, including an useful survey of foundational theories of depiction from the literature.

Chapter 2

Beyond Resemblance

In the last chapter I presented evidence for the Pictorial Semantics Hypothesis, the claim that pictures are associated with their content by systematic and conventional interpretive rules. But I offered no concrete proposal about what kind of semantics pictures actually have. Yet until we can articulate a plausible and detailed theory of depiction, the Pictorial Semantics Hypothesis remains purely conjectural. So, in this chapter, I will examine one particular family of theories, according to which all forms of pictorial representation should be analyzed in terms of *resemblance*. Such resemblance theories are among the oldest and most historically popular accounts of depiction, and I will show how they can be construed as a precise and plausible theories of pictorial semantics. But I will ultimately argue that, once suitably precisified, any resemblance-based analysis fails. The investigation of this failure motivates the positive proposal which I develop in the following chapter.

Inquiry into the nature of pictorial representation traditionally begins with the metaphysical question: what is it for a picture 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. The scene and the picture would seem to be linked by many dimensions of similarity.

Mamie Garcia, AP

Understood as an account of pictorial content, resemblance theory also makes computational sense. 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 which

depicts that scene, for each has the same properties. Recovering information about a sign's subject by directly perceiving the sign 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.[1] As I will argue in the first part of the chapter, resemblance theories are best understood as theories of pictorial semantics: 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 pre-conditions 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 *ismorphism*, a kind of similarity with respect to abstract relational structure.

In proportion to their historical currency, resemblance theories have been the target of numerous objections. Far from dampening the spirit of resemblance, however, the proposed counterexamples 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.[2] I will show how resemblance theorists have successfully responded to all of the most important charges pressed against them thus far.

But in the final part of the chapter, I present a new and general argument against what I consider to be the best version of the resemblance theory. 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

---

[1] In the 20[th] and 21[st] centuries, resemblance theorists have included Peirce (1906), Wittgenstein (1921/1997), Morris (1946), Gibson (1960; 1971), Manns (1971), Novitz (1975), Fodor (1975), Lewis (1986, §3.3), Neander (1987), Peacocke (1987), Malinas (1991), Files (1996), Hopkins (1998), Newall (2003), Kulvicki (2006*a*), Blumson (2007; 2009*a*), 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, §4).

[2] See, for example, Blumson's (2007) dissertation and a recent article by Abell (2009).

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. Unlike some previous 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.[3]

The chapter is organized as follows: §1 argues that resemblance theory is properly understood as a theory of pictorial semantics— a theory of accurate depiction. §2 takes the reader through four successive versions of the resemblance theory, beginning with the naïve and implausible, and building up to a sophisticated and compelling treatment of linear perspective. §3 lays out the argument against all such resemblance theories of depiction, by examining in detail the cases of curvilinear perspective. I argue that although linear perspective can be accounted for in terms of resemblance, curvilinear perspective cannot. §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 Resemblance theory as pictorial semantics

Advocates of resemblance theory have not always been perfectly explicit about their subject matter. Resemblance theories are typically presented simply as theories of "depiction". But as we have seen, there are two salient notions of depiction which play independent roles in pictorial communication. The first is that which determines pictorial *reference*— the thing which a picture is *of* or *about*. The second is that which determines pictorial *content*— the set of properties the picture ascribes to its referent. As we observed in the last chapter, the two notions come apart in cases of misrepresentation. In such cases, the content of a picture represents its referent as being a certain way, even while its referent does not actually answer to this representation. We also saw how pictorial content can be characterized in terms of conditions of accurate depiction. In that idiom, misrepresentation arises when a picture refers to a given entity, but fails to accurately depict the same. The question at hand is whether resemblance theories aim to provide an analysis of pictorial reference or accurate depiction? In what follows, I'll argue that the only plausible interpretation of such theories is as accounts of accurate depiction; resemblance theories, in other words, are theories

---

[3]Goodman (1968, p. 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.

of pictorial semantics.

Given that reference and accuracy can come apart, what determines each? The examples of misrepresentation discussed in the previous chapter 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 one hand, and those of Obama, on the other, inflexibly determine that the drawing is inaccurate.[4]

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.[5] Anyone familiar with 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 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* refer to Obama, we need only a little knowledge of Obama's actual appearance to infer the image's obvious inaccuracy.

Such examples strongly suggest 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

---

[4]On the intention-relativity of pictorial reference, see Cummins (1996, Ch. 2), Cohen and Callender (2006, p. 14), and Van Fraassen (2008, p. 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. (1968, p. 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 footnote. See Costello and Phillips (2009, p. 16) for related discussion.

[5]This type of example was suggested by Jeff King.

of pictorial reference, but an interesting and plausible analysis of accurate depiction.[6] For consider again the first pair of drawings of Obama: the short-hair drawing at once accurately depicts Obama, depictively refers to him, and significantly resembles him. (*left*) 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. (*right*)

Thus, on one hand, accurate 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.[7] But this view is belied by cases of extreme misrepresentation such as the blindfold-drawing above. For that is a misrepresentation of Obama, hence refers to Obama, despite resembling him in almost no relevant respects.[8] 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.[9]

For these reasons I will understand resemblance theory as a theory of accurate depiction, hence as an account of pictorial semantics. In this guise resemblance theory is not a pictorial semantics for just one *one* or *some* systems of depiction. It is uncontroversial that there are some systems which

---

[6]The irrelevance of resemblance to pictorial reference, but plausible relevance to accurate depiction is noted by Cohen and Callender (2006, Section 4) and Van Fraassen (2008). Suarez (2003) argues against resemblance theories of pictorial reference, 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, pp. 122-4) for detailed exegetical discussion.

[7]In the words of Hopkins (1998, p. 30), "pictorial misrepresentation... has its limits." See also Abell (2009, p. 212).

[8]The primary evidence cited in defense of a 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.

[9]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.

are governed by resemblance. The system by which paint-chips are used to depict the color of paint from a bucket is a plausible example. Rather, resemblance theory must be understood as a general claim about the semantic structure of *all* systems of depiction. Like Tarski's theory of truth for language, resemblance theory embodies a view about the universal architecture of pictorial semantics. In the next section, I will show how this claim can be made precise, and plausible. But in the subsequent section I will argue that this view of pictorial semantics is ultimately untenable.

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 pictorial semantics. 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 which are viewer-dependent.[10] 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 chapter, a resemblance theory of pictorial semantics is a theory that grounds accurate depiction in *any* kind of similarity, be it viewer-dependent, viewer-independent, or some admixture.[11]

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 resemblance theory which is at once flexible and compelling. Since my overall goal is to critique resemblance

---

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

[11]Resemblance theories also divide up according to whether the similarity in question is real or merely experienced. For example, even if we set 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, p. 188) for discussion.) But since in this chapter 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 her perception and judgement is guaranteed to be accurate. Ultimately, I develop an enriched notion of imagery which 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.)

theories *in general*, at each step I will incorporate only those specifics which 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 which 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.[12] 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.[13] 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 depiction. For example, this might be similarity with respect to shape, or similarity with respect to the visual impression elicited among viewers.[14] To express this analysis, the resemblance theorist posits a fixed set of properties $\mathcal{F}$ which are those properties relevant to determinations of depiction in general. The initial resemblance theory of depiction is then formulated as follows.[15]

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 $\mathcal{F}$.

---

[12] 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, §3.5) and Van Fraassen (2008, pp. 17-20).

Here and throughout I shall assume a liberal conception 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.

[13] See Neander (1987, p. 214) for the same observation.

[14] The latter is the view of Peacocke (1987) and Hopkins (1998).

[15] Such a theory is described by Lopes (1996, Section 1.2).

We may illustrate this theory diagrammatically. Where a picture P accurately depicts a scene S, we shall draw: P ![arrow] S. Where P and S are similar with respect to $\mathcal{F}$, we draw: P ←--→ $_{\mathcal{F}}$ 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.

iff

$\mathcal{F}$

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 $\mathcal{F}$. (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. They 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.[16]

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 is ac-

---

[16]The objection below follows Neander's (1987, pp. 214-5) reasons 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, pp. 20-32), albeit with respect to a Peacocke's (1987) particular version of the resemblance theory.

curate. *D* clearly does not require similarity with respect to color, for this image is black and white while its subject is multi-colored. If *D* did require similarity with respect to color, then the picture would only be accurate 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 (e.g. *C*) *must* require similarity with respect to color for accuracy; but (ii) other systems of depiction (e.g. *D*) *cannot* require similarity with respect to color for accuracy. The conclusion is that R1 must be wrong, for there is no single 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.[17] 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:[18]

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

---

[17]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.

[18]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, Section 1.5) describes, but does not endorse the variable approach. Variants on the view are defended in detail by Malinas (1991, pp. 282-91) and Abell (2009).

a scene S relative to a system I, we shall draw: $P \xrightarrow{\quad I \quad} S$. 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 relative to the system of line drawing D iff the drawing and the scene are similar with respect to $\mathcal{F}_D$. (*left*). Whereas accurate depiction in the system of color drawing C is grounded in a different kind of similarity– similarity with respect to $\mathcal{F}_C$. (*right*)


While the schema described by R2 leaves much to be filled in, it provides a flexible guide for 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, p. 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.[19] To begin, consider a line drawing which accurately depicts a scene involving Obama in system D.[20] 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. (*left*) Yet similarity is reflexive: for any A, A is similar to A.[21] 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. (*right*) Instead, it depicts Obama. In this case, depiction and similarity come apart, contrary to the claims of R2.

---

[19]It is not entirely clear that Goodman would approve of this application of his argument, because he may have been attacking resemblance theories of depictive reference rather than of accurate depiction. Walton (1990, Ch. 3 fn. 11) surveys the textual evidence both for the view that Goodman's notion of pictorial represenation is most like Walton's notion of "representation" (≈ depictive 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, §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.

[20]Several of Goodman's original cases are susceptible to the complaint that resemblance theory only aims 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.

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

picture $\mathcal{F}_D$ scene   picture $\mathcal{F}_D$ picture

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.[22] 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 $\mathcal{F}_D$. (*left*) 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 in with respect to $\mathcal{F}_D$. Yet the color drawing does not depict the line drawing. (*right*) Instead, it depicts Obama. Once again depiction and similarity come apart, contrary to the claims of R2.

line drawing $\mathcal{F}_D$ color drawing   color drawing $\mathcal{F}_D$ line drawing

Here again we find a case of similarity without depiction; thus similarity alone is not sufficient for depiction.[23] The resemblance theory must be revised.

---

[22]Newall (2003, pp. 386-7) suggests that a picture titled "Picture 12" with the caption "A picture of Picture 12" could accurately depict itself.

[23]While this is the same conclusion I advocate in §3, Goodman's reasons and mine are basically different. Furthermore, my argument target all resemblance theories, even those which are amended so as to circumvent Goodman's objection.

2.3 The reference condition

While Goodman's objection is valid and illuminating, it is easily deflected, as Goodman himself passingly observed. (Goodman 1968, p. 6) The proffered 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.[24] 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 which 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, p. 211; Blumson 2009*a*) 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 requirements for both accurate and inaccurate pictorial representation. It is the task of the resemblance condition to determine the accuracy of the representation.

---

[24]This reply is considered by Lopes (1996, p. 18) and defended, more or less, by Files (1996, pp. 404-5). Blumson (2009*a*) 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, p. 6) – but that (the artist intend that) the representation be distinctively pictorial.

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 depictive 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 pictorial semantics 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 which any pictorial semantics 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 geo-metrical 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 ex-tended 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 dif-ficulty and over several centuries, by artists and scholars at-tempting to recreate human perceptual experience on paper.

This endeavor proved successful in several notable respects. (Gombrich 1960, pp. 4-30) For exam-ple, perceptual experience and linear perspective projection have in common that objects which 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 non-linear 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 chapter 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 which 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 cannot be analyzed straightforwardly in terms of intrinsic similarity. Instead, a successful resemblance theory must somehow incorporate the key notion of a projection relative to a viewpoint. This in turn will require a further revision of R3. The remaining discussion is devoted to motivating and explaining 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 pictorial semantics 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 chapter will we see that even this sophistication cannot save the theory from demise.

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 a 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 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 non-starter, 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' *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' original argument is disarmingly compressed (Descartes 1637/2001, p. 90):

T his 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 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' 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 that of (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 a bird's eye view of the scene it *actually* accurately depicts.
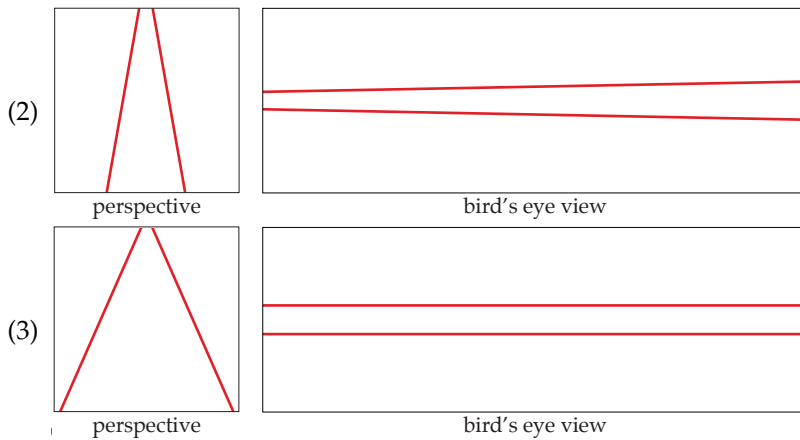
These examples now form the basis for a general argument that there can be no adequate analysis of perspective depiction in terms of similarity of intrinsic shape, no matter how nuanced or restricted the relevant notion of shape. The argument proceeds from three key premises, where *S* is the scene consisting of the two parallel lines described above.

Premise 1. Image (3) accurately represents *S* in linear perspective.

Premise 2. Image (2) *in*accurately represents *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 than Image (3) is similar to *S* in the relevant respects of shape. By Premise 3, it follows that (2) is similar to *S* in the *same respects*. By (1) again, it follows that (2) accurately represents *S*. But this contradicts Premise 2, that (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 (3) held in common with S— for example both are planes, both planes contain only two straight lines, both sets of lines are non-intersecting, both sets of lines are tilted at symmetrical angles, and so on. But note that all of these shared features are also ones that (2), the inaccurate image, holds in common with S. Of course, there *is* a difference between (2) and (3), the lines in the former define a more acute angle than those in the latter. But this does not make (3) more similar to S than (2) in any way. This is because, if anything, the lines of (2) are more nearly parallel than those of (3), while those in S are perfectly parallel; so (2) shares the property of minimal degree of parallelism with S, but (3) does not. Recall that the claim defended here is not that (3) and (2) are have all the same intrinsic shape properties in common with S, but only that whatever such properties (3) has in common with S, (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 which 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 (3) were inaccurate and (2) accurate, it would have been possible to object that (2) had something in common with the scene that (3) lacked, namely degree of parallelism. The objector could have claimed that this difference in degree of parallelism explained the accuracy of (2) and inaccuracy of (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 which 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 which incorporates these insights.

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 a 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.[25] To begin, suppose you are looking out flat a 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 window pane, 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 window pane 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 window pane 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.[26] To illustrate exactly how such a projection is defined, suppose the scene 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.[27] The viewpoint determines the "perspective" of the picture. We then introduce PROJECTION LINES which trace straight paths between select points on the cube and the viewpoint.[28]

---

[25]See Alberti (1991, Book I) and Da Vinci et al. (1970, II.83-5).

[26]The mathematical details of the following discussion are based on presentations in Sedgwick (1980), Hagen (1986, ch. 2-5), and Willats (1997, ch. 2, 5).

[27]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.

[28]The diagrams of this chapter are all drawn according to a highly imprecise system of depiction. The information they convey is impressionistic, not exact.

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 which 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 which are not occluded by another surface. For visual clarity in the diagrams above and below, I have only indicated those projection lines which connect the viewpoint to the accessible corners of the cube.
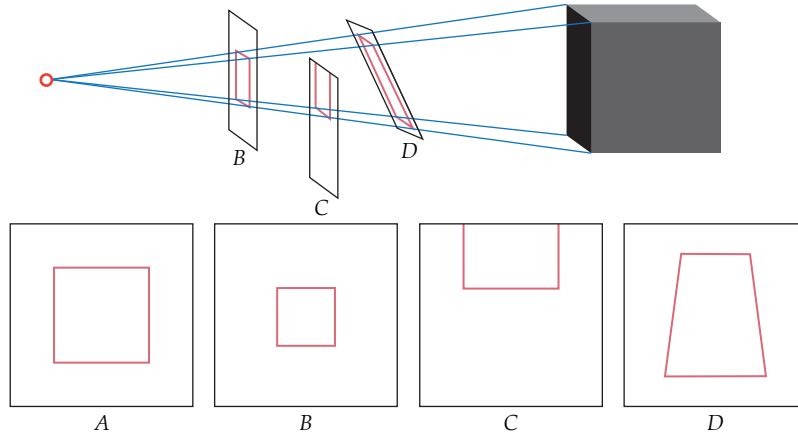
Into the spray of projection lines emanating from the viewpoint, we now introduce a PICTURE PLANE – analogous to Alberti's window. The picture plane can be positioned at any angle, but customarily, as suggested in the illustration below, it lies perpendicular to the line that connects the center of the 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, as at right, it reveals a side view of the cube.[29]


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

---

[29]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 crosshatch drawings, shaded surfaces in the scene correspond to regions in the drawing with a certain line density; and so on. Some of these cases are discussed in the next chapter.

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 which are in fact parallel are now represented by converging lines.

Each of the alterations just described changes the *way* the scene is represented, but none reveal 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 principle 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 owing 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 diagram above suggests, 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.[30] In any given case, we may label the scene S and projective index j. Finally, naming the projection function $\mathrm{lin}(\cdot)$, we shall say that lin applied to S and j returns an inscribed picture plane P, that is, $\mathrm{lin}(S, j) = P$.

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 which 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.

This suggests that we can state the accuracy conditions for L quite simply in terms of the pro-

---

[30] An explicit definition is supplied at the end of the next chapter.

jective formula $\text{lin}(S, j) = P$:[31]

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

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 $P = \text{lin}(S, j)$.

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 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 $P = \text{lin}(S, j)$.

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

---

[31]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.

vantage point of image A, and produce the following pair of pictures:

(11)

(12)

Intuitively, both (11) and (12) are both 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 $\text{lin}(S, j)$ 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.[32] 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 $\text{lin}(S, j)$.

One final emendation is required. According (13), any arrangement of shapes could depict a given cube S, relative to an index, so long as it 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.[33] 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 which is sufficient to guarantee depictive reference without generally determining whether the representation in question is accurate or not. This was just the mandate of the representation condition $\text{REF}(P, S)$ 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 $\text{lin}(S, j)$ & $\text{REF}(P, S)$.

We can illustrate this analysis with a moderate expansion of the chapter's diagramatic 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

---

[32] A more general and precise treatment of the same issue is presented in the next chapter.
[33] This problem was raised by H. Putnam 1981, p. 1. The solution adopted here is compatible with Putnam's own.

$\text{lin}(S, j)$. For convenience, I hereafter omit illustration of the reference condition.

the shape of P = the shape of $\text{lin}(S,j)$
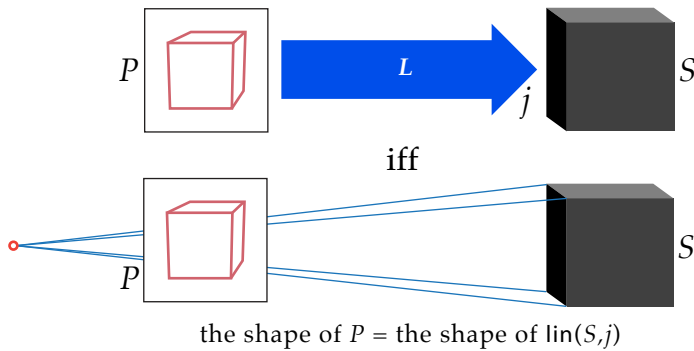
Proposition LP is a powerful tool for the theorist of depiction. By describing the exact relationship between a scene and a picture required by linear perspective, its sets a precise gold-standard for any pictorial semantic theory. The challenge for resemblance theorists is to reformulate LP with strict equivalence in terms that conform to the resemblance schema R3.

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 by 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 $\text{lin}(S, j)$ 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 (*right*), and then requiring that the result be perfectly alike P with respect to shape.[34]

---

[34]A proposal of this sort is defended by Hyman (Hyman 2000; Hyman 2006, Ch. 5).

Proposition (14) has the virtue of being equivalent to LP, 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 $\text{lin}(S, j)$

      & REF(P, S).

According to (14), the condition on accurate depiction is that P and $\text{lin}(S, j)$ are similar with respect to shape; that is, they share all the same shape properties. But this equivalent to the claim that the shape of P = the shape of $\text{lin}(S, j)$, which is exactly the condition imposed by LP. Unfortunately, despite the fact that (14) succeeds as an equivalent reformulation of LP, it is no 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 relata– 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 is patently *not* similarity-based in any straightforward sense, can be equivalently reformulated as a similarity relation under the transformation of one relata. Let the function m map individuals to their biological mothers, and let X be the set of properties of *being M* for every 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.[35] Thus similarity under the transformation of one relata has little to 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 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

---

[35] At the present time, I do not know what the limitations of this technique are– whether there are relations which cannot be reformulated in this way.

seems to covary with how a picture and the scene it depicts each appear in visual experience. This suggestion is illustrated below.
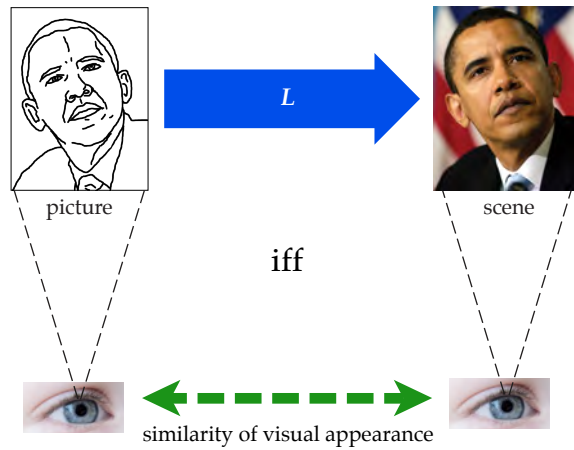
This observation suggests a framework for realizing a genuine resemblance analysis of linear persepctive, since we are asked to compare common properties of both relata– the visual appearance of the picture and the visual appearance of the scene. But how can we make this intuitive proposal sufficiently 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.[36]

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 which 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

---

[36]Proposals along these lines are defended by Peacocke (1987) and Hopkins (1998).

able to recapture LP. I shall term this analytic approach the method of SIMULTANEOUS PROJECTION. Unlike the asymmetrical analysis (14) which defines depiction in terms of similarity under selective projection of just one relata, 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 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 a projective index. In the method of simultaneous projection, it is natural to use this index to determine projection from the scene (at right in the diagram), 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 $\text{lin}(P, i)$ is similar to $\text{lin}(S, j)$ with respect to

shape & $\text{REF}(P, S)$.

According to (16), we determine whether P depicts S in three steps. First we take the perspective projection of S, $\text{lin}(S, j)$. Second we take the perspective projection of P, $\text{lin}(P, i)$. Finally we compare the shapes of 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 $\text{lin}(P, i)$ and $\text{lin}(S, j)$ are completely similar with respect to shape, then the shape of $\text{lin}(P, i)$ = the shape of $\text{lin}(S, j)$. Thus (16) is equivalent to:

(17) P accurately depicts$_L$ S relative to i, j iff the shape of $\text{lin}(P, i)$ = the shape of $\text{lin}(S, j)$

& $\text{REF}(P, S)$.

If we can further show that the shape of $\text{lin}(P, i)$ = 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).[37] For example, suppose that the projective index specifies that the plane onto which P is projected is titled forwards towards 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 similarity under simultaneous projection may come apart given a poor choice of the projective indices.

---

[37]Lopes (1996, p. 24) makes this point regarding Peacocke's (1987) theory of depiction, a view in many ways similar to (16).

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 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 which is parallel to it. In the second (B), S is projected onto a picture plane which 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)   (19)

A   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.[38] With the addition of this constraint, (16) successfully becomes equivalent to LP.

We now turn to the question of whether (16) is a genuinely resemblance-based analysis. One 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 the resemblance theorist has a more elegant solution at her 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.[39] 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

---

[38]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 projected shape, renders the operative notion of similarity in shape considerably less intuitive and direct.

[39]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 (2009*b*).

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.[40] It is plausible that these conditions are built into our conception 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.[41] 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 (16), as follows. (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.)

---

[40]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.

[41]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 stated using lambda abstraction: $\lambda X_y.H(\text{lin}(X, y))$.

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 especially common system of depiction, resemblance theory is able to provide a precise and descriptively adequate pictorial semantics. 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 pictorial semantics: 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 which can be grounded in resemblance, but rather, that there are some systems of depiction which cannot.[42] My ultimate diagnosis is that resemblance theory has mischaracterized the basic architecture of pictorial semantics. Rather than resemblance, accurate depiction in general is grounded in the more inclusive phenomena of geometrical transformation.

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

---

[42]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 there 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, pp. 68-9), and Newall (2006).

3.1 Curvilinear persective

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.[43] 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 phenomena remains a matter of dispute.[44] This fact should permanently quell any instinct to de-

---

[43]See Loverde and Weisstein (2009).

[44]Among the defenders of the view that natural perspective is curvilinear are none other than the 19[th] century perceptual scientist and physicist Hermann von Helmoltz (1867). 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).

nounce curvilinear perspective as "invalid" and linear perspective "valid". They are simply two systems of perspective, each with its own interest and utility.

Curvilinear perspective projection

Curvilinear perspective projection proceeds in two distinct steps. The first step is exactly like linear projection, save that the surface which 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.

view point

curved projection
surface

scene

The second steps 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.

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: curv$(S, k) = P$.

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 which 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 artists, deviates slightly from the intended rule in such 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 $\mathrm{curv}(S,k) = P$. 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 judgements 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 $\mathrm{curv}(S,k)$ & REF(P,S).

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

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 pro-

jection to successfully define accurate depiction 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 above makes this point obvious. This fact may be illustrated diagramatically 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.[45] However, it was also shown that when the parallel planes constraint was imposed on 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) which 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

---

[45]Lopes (1996, p. 24) notes this problem for Peacock's (1987) theory, a cousin of (20).

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, $\mathsf{curv}(P,i)$ and $\mathsf{curv}(S,j)$ do *not* have the same shape. For while the lines in $\mathsf{curv}(S,j)$ are curved projections from the straight edges of the cube, the lines in $\mathsf{curv}(P,i)$ 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.

Importantly, the fact that $\mathsf{curv}(S,i)$ and $\mathsf{curv}(P,j)$ 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 mis-drawn 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.

Lopes (p.c.) 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 0— that is, perfectly flat. This solution, while technically viable, comes at too high a philosophical cost. It effectively requires 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 0-curvature constraint. It is impossible to apply a 0-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* assymetrical, 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 which 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 which 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* depictive reference, but not accurate depiction. Hence they show that the conditions described by the best resemblance theory are insufficient for depiction.

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

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.

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, $R = \text{curv}(S, j)$.

Since R is a 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 uncurved. Thus it cannot be a curvilinear projection of S for *any* index. It follows that P_i does not accurately depict$_U$ S_j.

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 it failed. P is a picture *of* S, that is, a pictorial reference to S, thus the reference conditions is satisfied with respect to P and S.

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.[46] 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 apiece with the spirit of resemblance theory, and I know of no actual resemblance theory that violates it.[47]

The intuitive idea that the resemblance condition is only sensitive to qualitative properties of objects may be expressed more formally using a substitution principle. Objects that are qualitatively indistinguishable may be freely substituted into the resemblance condition with no change in truth value. 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:

---

[46]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 are properties such as *being that desk*. See Adams (1979) for discussion.

[47]Many resemblance theories make explicit commitments which seem to entail this assumption. For example, Files (1996, p. 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.

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 only sensitive 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 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 truth value. 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 which situates the viewpoint perpendicular to the center of the flat surface S, and perpendicular to the center of the picture P.[48] Thus, so far as *Qualita-*

---

[48]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 requirement of relational similarity between indices.

tive 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 a theories of resemblance conform with *Qualitative Indifference*, then we may derive our final premise with no further assumptions about the particular account of similarity at work.

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.

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.[49]

What does this argument reveal about R4's shortcomings? In the case above, the essential

---

[49]Two comments. First: The astute reader will notice a disconnect between the failure of the method of simultaneous projection for curvilinear perspective, illustrated above, and the argument just given. On one hand, the method of simultaneous projection failed because of cases of depiction without similarity. Thus the kind of similarity that it employed to ground depiction turned out not to be necessary for depiction. On the other hand, the argument just given shows that no degree of similarity is sufficient for depiction. Clearly these two results are not the same; what is there relationship? In fact, another general argument closely related to the present one can be given which corresponds to the failure of the simultaneous projection strategy above. According to this argument, if the resemblance condition in R4 is sufficiently substantive to explain cases of curvilinear *in*accuracy, then it cannot be necessary for accuracy. Whereas the argument above relies on the reflexivity of similarity, this one relies on the symmetry of similarity. Unfortunately, I have already taxed the endurance of the reader, so this argument must be presented at another time.

Second: The present argument relies on the reflexivity of similarity, as did Goodman's famous objection. What is the difference between the two? Goodman was either oblivious to any notion of depiction beyond depictive reference, or conflated reference with accurate depiction. (See footnote 19.) As a consequence, his arguments target cases which test the conditions for depictive reference alone. By contrast, I have taken care to focus cases where depictive reference is already guaranteed, and which instead test the conditions for accuracy alone. So whereas Goodman's argument revealed that a necessary condition for depiction is the irreflexive relation of reference, mine shows that accurate depiction requires some *further* irreflexive relation of fit.

problem for R4 was that *P* and *S* were *too* similar. As the discussion of Premise 1 made clear, *P* could only accurately depict scenes that differed from it in certain ways; and *S* could only be accurately depicted 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 non-similarity. 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.

3.3   Against resemblance

Resemblance theory is an account of the universal architecture of pictorial semantics, for it holds that *any* systems 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 general structure of pictorial semantics.

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.[50] 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 straight lines, viewed at the periphery of vision, appear curved.) If this is so,

---

[50]This is the position of Kulvicki (2006*a*). 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.

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 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 which resists analysis in terms of resemblance. Using the rubric of the argument presented here, the reader may go on to identify other counterexamples for herself: cases in which an accurate picture of a scene necessarily differs from its subject.[51] 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: overall shifts in lightness, darkness, and contrasts; false color and negative color; and the amplification of hue— as in technicolor— or suppression of it— as in the brown-tinted color of 19th century European painting. Of the last, Gombrich (1960, pp. 46-47) aptly wrote "it is a transposition, not a copy." In general, a relation which requires difference in some dimension cannot be analyzed in terms of one which 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. The problem with this claim is that similarity is cheap. Necessarily, any two objects are similar to each other in infinitely many respects. (Goodman 1972) For example, the number 3 and I are similar with respect to not being the number 2, the number 4, and so on... For her claim to be substantive, the

---

[51]Dilworth (2005, pp. 68-9), for one, makes precisely this objection against resemblance treatments of pictorial color, albeit rather briefly.

resemblance theorist 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. A proof of this claim is included in the appendix. The argument presented in this chapter undermines *any* resemblance-based account of depiction, wether it is rooted in first-order similarity or sophisticated isomorphism. It is therefore high time we looked beyond resemblance.

4   Beyond resemblance

I have argued that resemblance theory fails as a general account of pictorial semantics because there are actual systems of depiction which 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. In linear perspective, similarity with respect to linearity is a necessary condition. These 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.

If accurate depiction is grounded not in resemblance but systematic transformation, those resemblances which have seemed so powerful to so many theorists are not ruled out. Instead they are

understood as transformational invariants, properties of a scene which survive projection onto the picture plane. In this way, 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 judgements of individual viewers. 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 semantics, interpretive rules based on geometrical projection make more natural use of innate computational machinery involved in visual, perspectival perception.

The next chapter develops this projective theory of pictorial semantics in detail.

Appendix: Isomorphism as similarity

Here I show that isomorphism is a species of similarity. I will focus on isomorphism understood as a relationship between relational structures; a RELATIONAL STRUCTURE is a pair $\langle S, R \rangle$ where S is a set and R is a binary relation on S. Relational structures are useful in the present context, because both pictures and scenes can themselves be modeled as relational structures. Either kind of object can be understood as a pair $\langle S, R \rangle$ where S is a set of points with color or mass features, and R is a metric distance relation on those points. (A point with a feature can be modeled as the pair of a point and a feature.)

I now define isomorphism. For any relational structures $A = \langle S, R \rangle$ and $B = \langle S', R' \rangle$: A IS ISOMORPHIC WITH B iff there is a bijective function $f : S \to S'$ such that, for any $m, n \in S$: $R(m, n)$ iff $R'(f(m), f(n))$.

If A and B are isomorphic, they are similar with respect to certain abstract structural features. These abstract features are difficult to characterize directly. Instead, we will do so indirectly. Let Q be a set of all possible relational structures defined over some continuous class of points (for example, the real numbers). Let P be the set of properties *being isomorphic with q* for every $q \in Q$. More formally: $P = \{F \mid \exists q \in Q : F = \lambda x.x$ is isomorphic with $q\}$. We can think of the properties in P as abstract structural properties. Now we can show that similarity with respect to P is equivalent to isomorphism.

First, from similarity to isomorphism. Suppose A is similar to B with respect to P. I assume that any relational structure must have at least one property in P. Then there is at least one property F in P such that A is F and B is F. Then there is at least one element q in Q such that A is isomorphic with q and B is isomorphic with q. But isomorphism is symmetrical, so q is isomorphic to B. Isomorphism is also transitive, so A is isomorphic with B.

Second, from isomorphism to similarity. We proceed by reductio. Suppose A is isomorphic with B, but A and B do not share all the same properties in P. Then there is some property G in P such that A is G but B is not G. Then there is some structure q in Q such that A is isomorphic with q, but B is not isomorphic with q. But isomorphism is symmetric and transitive; since, by assumption, A isomorphic with B, and A is isomorphic with q, it follows that B is isomorphic with q. Contradiction.

Thus isomorphism is equivalent to similarity with respect to P. Isomorphism is therefore a species of similarity.

Chapter 3

Depiction as Projection

In the last chapter, I argued that the natural idea that pictorial representation is grounded in resemblance is, in conception, flawed. Any adequate account of depiction must be based instead in some more general notion of transformation. In this chapter, I'll apply that lesson. Inspired by the analysis of images in projective geometry and computational perceptual science, the theory I'll defend here holds that, in any system, depiction is defined in terms of *geometrical projection*: a picture accurately depicts a scene just in case the picture can be derived from the scene by a method of projecting three-dimensional scenes onto two-dimensional surfaces. The resulting *Geometrical Projection Theory* is a largely novel account of depiction, boasting an exactness and accuracy that make it a promising foundation for any future pictorial semantics.[1]

The proposal is developed in four parts. §1 presents the theory in its most general form: §1.1 briefly introduces the key concept of geometrical projection; §1.2 then articulates the core commitments of the Geometrical Projection Theory; §1.3 distinguishes the theory from the main alternative accounts of depiction. In §2 the theory is elaborated in greater detail for the particular system of perspective line drawing. §3 concludes the informal discussion. Finally, in §4 I develop the analysis of perspective line drawing as a precise formal semantics.

1   The Geometrical Projection Theory

According to Geometrical Projection Theory, accurate depiction is directly defined in terms of projection. I begin, then, with a brief review of the mechanisms and variety of geometrical

---

[1]There are, to my knowledge, only a very few explicit proposals for semantic accounts of pictorial representation. They include Malinas (1991) and Mackworth and Reiter (1989). Setting aside its affiliation with semantics, the novelty of the Geometrical Projection Theory depends on the comparison class. In projective geometry and perceptual science something like this account is simply *assumed* (see, e.g. Hagen 1986, Willats 1997, Durand 2002). But such inquiry tends to be inexplicit about the notion of depiction under analysis. On the other hand, in philosophical discussion, where the subject is more clearly set out, the theory has little precedent. The most similar account is that of Hyman (2006, ch. 6). The present theory generalizes, precisifies, and diverges from Hyman's proposal in a number of ways; see footnote 14 for discussion. A related account, which nevertheless differs fundamentally from the present theory is that of Kulvicki (2006*a*, ch. 3-4); see footnotes 8 and 35.

projections.

1.1 Geometrical Projection

GEOMETRICAL PROJECTION is a general method for transposing three-dimensional scenes onto two-dimensional picture planes, much the way a flashlight may be used at night to project the shadow of an extended object onto a flat wall. The method works by first defining an array of PRO-JECTION LINES which connect a scene to a picture plane according to a rule of projection; these lines are then used to map spatially distributed features of the scene to surface features of the picture. A simple example is illustrated below, with the resulting picture plane revealed at right.[2] (Note that the projection lines indicated here are only a representative sample of the full array).

In this particular method of projection, known as PERSPECTIVE PROJECTION, all the projection lines are defined by a single point of convergence (the ∘ above). This point can be shifted, with the effect that new features of the scene are revealed in the projected image. The result is a different *projection* from the one just illustrated, but the *method* of projection— perspective— remains the same.

Perspective projection was first fully codified by artists and scholars of the Italian Renaissance, but it is only one of a great many other methods of projection. The diagram below illustrates a version of PARALLEL PROJECTION, in which the projection lines, rather than converging on a single *point*, are all perpendicular to a single *plane*, hence parallel to one another. The resulting image is subtly but visibly distinct from the perspective projection.[3]

---

[2]The diagrams of this chapter are all drawn according to a highly imprecise system of depiction. The information they convey is impressionistic, not exact.

[3]The projection illustrated below is an oblique projection, a species of parallel projection.

PARALLEL PROJECTION

These examples illustrate how methods of projection may vary with respect to the configuration of their projection lines, hence distinct *structural* relationships between the scene and the projected image. But systems can also be distinguished according to *which* particular features of the scene are mapped to the picture plane, independent of the structure of the mapping. For example, in all of the drawings considered above, *visible edges* in the scene were mapped to *lines* on the picture plane. But there are many other possibilities: in "wireframe" projections, *all* edges in the scene are mapped to lines in the picture. In methods of color projection, colors in the scene are mapped to suitably related colors in the picture. These and other variations— by no means exhaustive— are illustrated below, in combination with the distinctive structure of perspective projection.[4]

| visible edges ⇒ lines | all edges ⇒ lines | surfaces ⇒ red | visible edges ⇒ lines <br> visible faces ⇒ pink | colors ⇒ colors |

In §2 I will explain the method of projection for perspective line drawings in considerably more detail, and in the Appendix the definition is explicitly formalized. For the present, I hope the preceding illustrations are sufficient to provide the reader with a working understanding of the general notion of geometrical projection. With this in hand, we are now in a position to examine this chapter's specific proposal for a semantics of pictures.

1.2   The Geometrical Projection Theory

What is the relation that characteristically holds between a picture and a scene it accurately depicts? As we'll see, philosophers of art have considered a variety of answers to this question. The orthodox view is that the relation is one of similarity with respect to a carefully constrained

---

[4]This explanation of geometrical projection, along with the notation for feature mappings, is drawn from Willats (1997).

set of spatial and chromatic properties. A radical reaction is that the relation is utterly arbitrary, determined by a library of picture-scene associations. A very different approach holds that the relation is defined in terms of the perceptual effects of the picture on a suitable viewer. The theory on offer here diverges markedly from all of these, holding instead that the relationship in question is simply that of geometrical projection.

Officially, the Geometrical Projection Theory has two components: the first describes the relationship between systems of depiction and methods of geometrical projection, while the second defines accurate depiction relative to a system in terms of projection.

(1) *Geometrical Projection Theory*

   (a) Every system of depiction determines a unique method of projection.

   (b) A picture accurately depicts a scene in a given system if and only if the picture can be derived from the scene by the projective method characteristic of that system.

I have in mind a formally precise rendering of this hypothesis, at least for the particular system of perspective line drawing, which I will develop in the next section. But for now, my aim is to elucidate the philosophical underpinnings of the theory while contrasting it with the principal alternatives.

The first claim of the Geometrical Projection Theory is that every system of depiction determines a unique method of projection. Recall that *systems* of depiction are interpretive conventions which establish standards of accuracy for pictures. Methods of *projection*, on the other hand, are geometrical techniques for mapping three-dimensional scenes onto two-dimensional surfaces. According to the Geometrical Projection Theory the two are directly correlated. For example, the system of perspective line drawing uniquely specifies the method of projection for perspective line drawings. Each system encodes a distinct and stable projective rule of interpretation, its use in communication established by convention.[5] To understand a system of depiction, on this view, is to be able to apply its characteristic method of projection. Such an ability need not be based on explicit knowledge. Just as language users rely on tacit knowledge of abstract grammar rules to guide interpretation, successful pictorial communication requires only a tacit knowledge of methods of geometrical projection.[6]

---

[5]Jakobson (1921/1987, p. 21) again: "The methods of projecting a three-dimensional space onto a flat surface are established by convention; the use of color, the abstracting, the simplification, of the object depicted, and the choice of reproducible features are all established by convention."

[6]See Hagen (1986, pp. 84-5). Admittedly, in the typical case, competent *expression* in a system of depiction requires formal training, while *comprehension* does not. This is rather different from the case of language, where mastery of expression and comprehension are regularly achieved without formal training. But all this shows is that implicit knowledge of the rules of pictorial interpretation does not automatically entail an ability to draw. And it should also be emphasized that even formal artistic training never takes the form of fully precise and explicit drawing rules, but rather of loose heuristics the likes of

The second component of the Geometrical Projection Theory specifies the role of geometrical projection in any system's definition of accuracy: a picture bears the relation of accurate depiction *to* a scene if and only if it is possible to derive the picture *from* the scene by geometrical projection. For an artist to intend to draw accurately is for the artist to intend (perhaps unconsciously) to perfectly execute a geometrical projection of some intended scene. Note that this is only an *in principle* projective mapping from pictures to scenes, not a *physical* process by which pictures are generated, nor a *computational* process by which projections are calculated. Like all descriptive semantic theories, the Geometrical Projection Theory articulates the conventional standards which govern the creation and interpretation of pictures, but imposes no constraints on how creation and interpretation are actually carried out.

The best positive evidence for the Geometrical Projection Theory is simple induction over cases. Suppose two artists set out to draw a cube S from a given vantage point, relative to the system of perspective line drawing. The first picture, (2), is created in such a way that it conforms with perspective projection. And it seems that (2) is an accurate pictorial representation of S in the system of perspective line drawing. But the second picture (3), even though it deviates only slightly from (2), is such that there is no way that (3) could be derived from a cube by perspective projection. (The validity of this claim will become apparent in the next section.) Furthermore, (3) is an *in*accurate pictorial representation of S in the system of perspective line drawing. Thus accuracy in a perspective system and perspective projection appear to covary precisely, exactly as the Geometrical Projection Theory requires. This thought experiment may be carried out indefinitely on any given scene built from simple geometrical solids, and multiplied indefinitely with appropriate modifications for any other system of depiction.[7]

(2)   (3)

A pictorial semantics ultimately determines a mapping from pictures to content. At the outset I defined the content of a picture as the set of scenes it accurately depicts. According to the Geometrical Projection Theory, a picture accurately depicts a scene just in case it can be derived from

which require a prior understanding of the relevant standards of accuracy.

[7]In particular, the evaluation can be carried out for the system of perspective line drawing with any given scene built from polyhedra— geometrical solids with only flat faces. The restriction is necessary because, in the crude method of edge-to-line mapping defined above, curved surfaces fail to register on the picture plane. Later we will consider more sophisticated projective methods which correct this problem. In the next section, we will also see that the claim of precise covariation requires some qualification.

the scene by geometrical projection. Putting these definitions together yields the following analysis of pictorial content:[8]

(4) ***Geometrical Projection Theory for Content***

The content of a picture in a given system is the set of scenes from which it may be derived by the projective method characteristic of that system.

Thus, to grasp what a picture depicts is to understand the space of possible scenes from which it could be projected. Those features held in common among the scenes in a picture's content correspond to how it represents the world as being; those features which vary among the scenes in a picture's content correspond to features of the worlds about which the picture is non-commital.

So, for example, the content of the drawing at right below will include all and only scenes containing shapes with visible faces which I will loosely describe as "square-like", because these are the scenes the picture could have been projected from. This "square-like" property is the feature which the picture represents the world as instantiating. But the picture does *not* represent the world as containing a *cube* or even a shape with a *square* face necessarily. For there are infinitely many irregular shapes, like those illustrated below, whose irregularities are not preserved in the given perspective projection.

Furthermore, because geometrical projection in general is insensitive to the identity and typological properties of the scene projected, such features are also left out of a picture's content. For example, the image of Obama presented earlier could equally well be a projection of Obama, a wax-sculpture of Obama, or suitably arranged lines in the sand. Thus its content would not differ-

---

[8]The analysis of content provided here is superficially like that proposed by Kulvicki (2006*a*, p.57-9). But Kulvicki and I reach these conclusions by very different routes. Kulvicki holds that content of a picture is (or is defined by) its projective invariants (but offers no explicit account of depiction)— a view I argue against specifically in Greenberg (2010*a*). I hold instead that content of a picture is its accuracy conditions, and define accurate depiction in terms of projection. These conceptual differences result in subtly but crucially divergent precisifications of the formulation here. See footnote 35 for discussion. In addition, Kulvicki holds that perspective projection is, in some sense, the prototypical method of projection, while I believe nothing of the sort.

entiate between these alternatives.[9] As a consequence, a purely projective semantics like the one offered here is only suitable as a theory of spatial and chromatic *explicit* content. It is not intended to capture the rich, typological *extended* content sometimes ascribed to pictures (e.g. that this is a picture *of Obama*, or *of a man*).[10]

It might seem incredible that a semantics whose output is content as abstract and unspecific as this could play any role in an explanation of how novel pictures are quickly and reliably converted into useful information. But it is important to remember that pictures, like any kind of sign, are always interpreted against a backdrop of world knowledge. The same content which is quite unspecific taken as a whole, also contains quite specific *conditional* information that is elicited under the right presuppositions. A semantics based on geometrical projection provides the systematic and conventional rules which are necessary to extract bare-bones representations of the world from marked, flat surfaces. In concert with appropriate background knowledge about the world— as opposed to further knowledge of interpretive conventions— these admittedly skeletal representations can yield impressively precise and useful information.

For example, we saw how the content of the square drawing above is indeterminate between an infinite space of radically different objects. But suppose that we are independently informed that this picture accurately depicts a scene containing exactly one Platonic solid. (Roughly, a three-dimensional solid is Platonic if all of its faces are regular polygons of the same kind.) On its own, this knowledge restricts the space of possible scenes to *five*, corresponding to the five Platonic solids, but not fewer. Yet the square picture could only be a projection of *one* Platonic solid, the cube. Thus the content of the picture and background knowledge together entail an exact description of the scene in question— information which is more specific than either source on its own.[11] In the same way, the Obama drawing considered in isolation could equally well depict lines in the sand, wax-sculptures, and infinitely many other oddities. But once we are informed that it accurately depicts a *human*, and given our background knowledge about humans, we are thereby provided with a mass of quite precise information about the physical characteristics of the person in question. When such supplementary correspondences between picture and world are specifically intended by the artist,

---

[9]Similarly, I understand "ambiguous" images like the Necker cube or duck-rabbit illustration to be merely semantically indefinite between, in the latter case, duck-scenes, rabbit-scenes, and infinitely many others. For psychological reasons anterior to the semantics, the duck and rabbit interpretations are especially vivid to human viewers.

[10]See Chapter 1 for further discussion of this distinction, and footnote 30 for its sources in the literature.

[11]Of course, the background knowledge in question need not put direct constraints on the content of an image in order to allow a viewer to extract rich content from it. Consider an agent who has antecedent knowledge that a picture accurately depicts some actual scene— because of her knowledge of the reliability of the artist— but who is otherwise wholly ignorant about what it depicts. Given background knowledge about what *the world* contains, she may come, after viewing the picture, to conclude that it accurately depicts a cube, because there is nothing else in the world it could accurately depict. Richer interactions, which make way for more measured confidence in the reliability of the artist and the furniture of the world, are the characteristic domain of Bayesian models of interpretation.

it may be appropriate to construct a model of *rich* pictorial content, where these mappings are built into the content itself. But this development lies beyond the scope of the present inquiry.[12]

1.3 Alternative theories

The Geometrical Projection Theory takes its inspiration from the study of projective geometry and recent work in computational perceptual science.[13] But the view differs sharply from the theories of depiction which have dominated contemporary Philosophy of Art, each of which has evolved from a correspondingly different point of origin.[14] In what follows I distinguish the Geometrical Projection Theory from the three main proposals discussed in this literature. Here two provisos are necessary. First, I do not assume that these theories are mutually exclusive; they represent distinct areas of emphasis in a spectrum of actual and possible proposals. Second, I will treat all the theories under discussion as theories of *accurate* depiction, though in the literature they are typically stated simply as theories of *depiction*. I hope the translation does not substantively distort the original intent. Given the constraints of space, I am not able to pursue detailed objections and replies; my aim is merely to distinguish The Geometrical Projection Theory from the alternatives.

(i) The *Symbol Theory* of depiction articulated by Nelson Goodman (1968, p. 5) begins with the recognition that communication with pictures, as with language, is governed by conventional systems of interpretation. Impressed by this analogy, Goodman went on to hold that depiction is based on a library of arbitrary correspondences between pictures and their content, much like the lexicons of language.[15] Yet this view is implausible on its face: the relationship between, say, a drawing of Obama and the scene it accurately depicts is much more intimate and direct than the admittedly arbitrary connection between, say, the word "tree" and the property of being a tree (or

---

[12] See Greenberg (2010*b*) for further discussion.

[13] An active research program in computational perceptual science over the last forty years has attempted to articulate ever more complete algorithms for the interpretation of line drawings. (Early contributions include Huffman (1971), Clowes (1971), and Waltz (1975). Two excellent overviews of this field are Willats (1997, ch. 5) and Durand (2002).) The principles evolved in this literature are not general standards of correct interpretation. Rather, they are implementable algorithms, based on reasonable heuristic assumptions, for determining reasonable interpretations. (e.g. Huffman 1971, pp. 297-9) Pictorial semantics occupies a distinct but compatible level of theoretical description: it articulates an abstract mapping from pictures to information which computational theories of interpretation aim to make practically computable. (See Feldman (1999) for a discussion of the analogous distinction for vision.) The present work owes a special debt to the discussion of depiction in Willats (1997).

[14] Among proposals by philosophers, the account of Hyman (2006, ch. 6) most nearly prefigures the Geometrical Projection Theory. Hyman's account is based on the recognition that projection is a necessary condition for depiction. But the similarities with the present view largely end there. For example, Hyman does not construe his account as a theory of accuracy; his proposal is highly informal, with little detail about the role of the projective index, or the connection between the geometry of projection and the definition of pictorial line and color (see §3 for my own position on these matters); and he seems to assume that perspective projection in particular is central to the definition of depiction, which I reject.

[15] Goodman would have objected to the label "symbol theory", but not the characterization in terms of arbitrary meanings. Lopes (1996, p. 65) notes that Goodman also resisted the view that systems of depiction are conventional in any obvious sense. Bach (1970) is among the few adherents of Goodman's view.

a given tree). Furthermore, as Wollheim (1987, p. 77) has pointed out, the theory fails to explain how we are able to interpret novel images with such facility, for there is no straightforward way of extrapolating from a list of arbitrary associations to new cases.

According to the Geometrical Projection Theory, by contrast, accurate depiction is based on *general* rules of projective transformation, not a long list of arbitrary meaning associations. On this view, the interpretation of novel images is straightforwardly explained as the application of a familiar general rule to unfamiliar particular cases. To determine whether a novel image accurately depicts a given scene, one must simply consider whether the former can be derived from the latter by geometrical projection. Meanwhile, the apparently intimate relationship between an accurate depiction and its subject is explained by the relative simplicity of such projection rules and the speed with which they are applied— aided no doubt by the computationally powerful visual system. Of course, the projection analysis is compatible with the fact that systems of depiction, like languages, define conventional rules of interpretation. But the theory rejects the conflation of these rules with the arbitrary lexicons of language. Instead, each system is a convention for interpreting pictures according to a one of the many possible methods of projection.

(ii) *Resemblance Theory* is the oldest and most orthodox view of depiction.[16] Its starting place is the powerful intuition that a picture, say of Obama, and the scene it accurately depicts are substantively similar with respect to certain color and shape properties. This observation has inspired the following view, roughly stated: a picture accurately depicts a scene just in case the picture and the scene are similar in the appropriate respects (and, in many accounts, the artist intends the picture to represent that scene).[17] Resemblance theorists have widely acknowledged the importance of geometrical projection in specifying the right relationship between pictures and scenes, but have attempted, sometimes with considerable effort, to fold this fact into a genuinely similarity-based analysis.[18]

The Resemblance Theory has been subjected to a host of objections, each met with ever more sophisticated revisions.[19] Yet as I have argued in Chapter 2, the very structure of the theory renders it empirically inadequate. The basic point can be made with respect to the system of CURVILINEAR

---

[16]In the 20th and 21st centuries, resemblance theorists have included Peirce (1906), Wittgenstein (1921/1997), Morris (1946), Gibson (1960; 1971), Manns (1971), Novitz (1975), Fodor (1975), Neander (1987), Peacocke (1987), Malinas (1991), Files (1996), Hopkins (1998), Newall (2003), Kulvicki (2006*a*), Blumson (2007; 2009*a*), 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, §4).

[17]Resemblance theorists vary considerably in how they revise or extend this rather crude formulation.

[18]For example: Peacocke 1987; Hopkins 1998, pp. 71-81; and Kulvicki 2006*a*, ch. 3-4.

[19]The most fundamental objections have been those posed by Descartes (1637/2001, p. 90), Goodman (1968, pp. 3-10, 34-9), and Lopes (1996, ch. 1). Recent, sophisticated forms of Resemblance Theory have been proposed by Kulvicki (2006*a*), Blumson (2007) and Abell (2009).

PERSPECTIVE, as in the "fish eye" photograph below. Due to its characteristic curvature, the image clearly does *not* resemble its subject in the intuitive and direct way exhibited by standard perspective depictions. Resemblance theorists have tried to preempt this concern by claiming that such curvature is semantically inert, inconsequential to the content of the image, which they insist on analyzing in terms of similarity. Yet I showed that just the opposite is the case. The accuracy of a picture in curvilinear perspective depends on the precise degree of deviation between the straightness of edges in the scene and curvature of lines in the picture. In this and many other systems, accuracy depends as much on a picture's systematic *differences* with the scene depicted as on its similarities. It follows that these systems cannot be analyzed in terms of similarity alone. Instead, a more general notion of transformation is required; I propose geometrical projection.

Both the Geometrical Projection Theory and the Resemblance Theory reject the view that accurate depiction is based on arbitrary associations between pictures and scenes; and both hold that it is grounded instead in a spatial and chromatic relation of "fit" between a picture and its subject. But they differ fundamentally in their construal of this relation. Formally, the Resemblance Theory holds that the relation underlying accurate depiction is a *symmetrical* relation of similarity, but for the Geometrical Projection Theory it is an *asymmetrical* relation of transformation. These relations are conceptualized very differently in each case: whereas the Resemblance Theory claims that accurate depiction arises from the *commonalities* between the picture and the scene, the Geometrical Projection Theory holds that depiction arises from the fact that the picture may be *derived* from the scene by projection. This shift in conception buys the latter a theoretical simplicity and empirical validity that resemblance accounts fall short of.

At the same time, those apparent similarities between pictures and scenes which have seemed so powerful to so many theorists are not ruled out. For any kind of transformation, there are INVARIANTS— properties which are always held in common between the inputs and outputs of the transformation.[20] The transformational invariants for common methods of projection are well known. For example, under perspective projection, straight lines in a scene are always mapped onto straight lines in the picture; so straightness of line is an invariant. (The inverse does not hold generally.) A consequence is that "betweenness" relations are also invariant. By contrast,

---

[20]See Hagen (1986, ch. 2).

101

parallelism is *not* an invariant; in perspective, parallel lines are usually mapped onto converging lines. By understanding the alleged resemblances between pictures and scenes as transformational invariants, the intuition of similarity may be substantiated even while the proposal to ground depiction in resemblance is rejected.

(iii) Finally, *Perceptual Theories* follow from the striking analogies between the process of picture interpretation and that of normal visual perception; depiction is then analyzed in terms of the kinds of perceptual effects a picture has on suitable viewers. The crudest perceptual account holds that a picture accurately depicts a scene if and only if seeing it causes a viewer to have the *same* perceptual experience she would have upon seeing the scene.[21] Depending on the theory, more sophisticated variants hold that a picture accurately depicts a scene if and only if it permits a viewer to see the scene *in* the picture, it facilitates the viewer's *recognition* of the scene, or it is a suitably vivid prop with which the viewer may *pretend to see* the scene, and so on.[22]

Despite their diversity, all Perceptual Theories face the same basic problem of giving precise articulation to their preferred notion of a "suitable viewer". The problem arises because not just any viewer will fit the designated theoretical role. There is no hope, for example, of defining accurate depiction in terms of the perceptual reactions of a chronic hallucinator or an acute myopic. Instead, Perceptual Theories require that the perceptual system of a suitable viewer must be, in some sense, ideal. In order to give substance to her proposal, it is incumbent on the Perceptual Theorist to specify the operative notion of ideal perception. And because perceptual abilities vary so widely, is not clear that such an idealization can be derived straightforwardly from the generalizations of contemporary perceptual science, or from statistical regularities. That is not all: in order to capture the difference in accuracy conditions between various systems of depiction, a Perceptual Theory must posit differences in perception for each. But then we are owed an account not merely of *ideal* perception, but of *ideal black and white* perception, *ideal color* perception, *ideal line drawing* perception and so on. What are these different modes of perception? Perceptual Theorists have had little to say in the way of specific answers to this question. But without these details explicitly worked out, Perceptual Theories cannot claim to offer a substantive analysis of accurate depiction.

The Geometrical Projection Theory, by contrast, offers a significant advance over Perceptual Theories by explicitly and precisely defining the conditions of accurate depiction in terms of geometrical transformation. (In the next section and in the formal presentation of §4, I show that

---

[21]Such a view is defended by Gibson (1954), and often ascribed to Gombrich (1960), though it is unclear if the ascription is entirely apt.

[22]The "seeing-in" theory is due to Wollheim (1987, pp. 46-59; 1998). Schier (1986, ch. 3) and Lopes (1996, ch. 7) defend recognition-based views. Walton (1990, ch. 8) is the originator the pretense theory. Loewer and Godbey (1978, pp. 339-40) and Wollheim (1998, pp. 217-8) succinctly capture the core commitments of this group of theorists.

this definition can be quite exact.) The many systems of depiction do not engender many modes of ideal perception, but merely many varieties of projection. Plausibly, if there were an ideal perceiver of pictures, her visual system would perforce follow the geometrical rules supplied by the Geometrical Projection Theory. But the theory need not posit any such creature; it characterizes the rules directly, without mediation. This is not to say that the forms of geometrical projection at work in common systems of depiction are divorced from visual perception. Instead, they are abstractions and strategic extensions of the *optical* projection of light into the eye. Thus the computations required to interpret a picture are broadly of a kind with those required of the visual system to interpret patterns of illumination reflected to the eye in normal perception.

Of course, Geometrical Projection Theory on its own does not explain why only a narrow range of projective methods are actually implemented in normal human pictorial communication. There are many possible systems of depiction, and some, like those that scramble their projection lines in eccentric ways, are never actually used. A plausible explanation is that although the human visual system can be trained to interpret a variety of projective methods that diverge from those that figure in normal perception, this capacity of limited. Ultimately, only a narrow range of projective methods can be computed by that part of the human visual system which is harnessed for the interpretation of pictures. Yet this appeal to visual biology is philosophically innocent. Whereas Perceptual Theories rely on the workings of *ideal* viewers, the Geometrical Projection Theory appeals only the capacities of *actual* viewers to explain why only a narrow range of projective methods are *actually* used. Systems of depiction are geometrical abstractions, just as generative grammars are logical abstractions; the extent and variety of such abstractions which can be successfully computed by the organs of human cognition is matter of ongoing empirical inquiry.

This concludes my general presentation of the Geometrical Projection Theory. I've shown how the theory yields a pictorial semantics— a systematic and conventional mapping from pictures to content— and I've indicated why it realizes the aims of pictorial semantics better than the alternatives. If the theory can be stated explicitly and precisely for at least one system of depiction, then I will have begun to meet the principle challenge facing the Pictorial Semantics Hypothesis. In the next section, I show that such precision can be achieved.

2   Semantics for Perspective Line Drawing

Having laid the foundations of the Geometrical Projection Theory, I now turn to developing it in detail for the particular case of perspective line drawing.[23]

Though this is just one of infinitely many possible systems of depiction, it has witnessed an extraordinary rise in global popularity since its codification in the Italian Renaissance. Today, drawings made in perspective are the norm, and nearly all mass market camera lenses are designed to mimic the geometry of perspective depiction.[24] Its most notable characteristic is that more distant objects are represented by smaller elements in the image, like the railroad ties at right; consequently, receding parallel lines in the scene are represented by converging lines in the image.[25]

This system is too complex to analyze systematically in the space of just one paper. Instead, my method here will be to develop an exact semantic theory for a recognizable *fragment* of the system of perspective line drawing.[26] According to the Geometrical Projection Theory, all systems of depiction are defined in terms of their characteristic method of projection. So I will begin in §2.1 by describing in some detail the method of projection associated with this simplified version of perspective line drawing. §2.2 then integrates this account into a precise semantics for the system of depiction in question. An explicit formalization of this analysis is included in §4. Finally, §2.3 considers how the semantics for this fragment might be extended to match the systems of line drawing found in ordinary pictorial communication.

---

[23]What I am calling "perspective" here is sometimes referred to as "linear perspective", in contrast with "curvilinear perspective" which involves projection onto a curved projection surface.

[24]More exactly, mass market camera lenses are designed to mimic *linear*, as opposed to *curvilinear* perspective; they yield merely perspectival representations more or less by default. See Pirenne (1970) for discussion.

[25]Here it is crucial to observe a distinction between the geometry of perspective *depiction* and the geometry of human visual *perception*. Although perspective was developed over many centuries by artists and scholars attempting to recreate perceptual experience on paper, the one is ultimately a product of convention, the other biology. (See: Da Vinci et al. 1970, II.107-109; Gombrich 1960, pp. 4-30; Gibson 1978, p. 232; Hyman 1995.) In fact, recent empirical research by Rogers and Rogers (2009) strongly suggests that some roughly perspectival but *non*-linear system best models the structure of visual perception. Many have been led to the same conclusion by less quantitative methods: Helmholtz 1867; Pirenne 1970; Hansen 1973; Arnheim 1974; and Hansen and Ward 1977.

[26]This is the approach pioneered by Montague (1970*a*, 1970*b*) in the study of natural language semantics. See Partee (1980, pp. 1-5) for discussion.

2.1 The method of perspective projection

A method of projection is a kind of algorithm for mapping three-dimensional scenes onto a two-dimensional picture planes.[27] Following Willats (1997), this method may be divided into two stages. In the first stage, an array of connecting lines are established, linking points in the scene to point in the picture plane. In the second, properties of the scene are mapped onto colors in the picture plane, according to a spatial distribution determined by the connecting lines of the first stage. Methods of projection can be defined by how they carry out these two stages: the first stage is specified by a given method's PROJECTION CONDITION, the second by its DENOTATION CONDITION.[28]

In any method of projection, we begin by positioning a finite PICTURE PLANE in relation to a given scene, and then defining a set of PROJECTION LINES which establishes a system of correspondences between points in the picture and points in the scene— illustrated below.[29] For every point in the picture plane, exactly one projection line passes through it. The projection condition structures this set of correspondences by specifying a certain configuration of projection lines. The projection condition characteristic to methods of *perspective* projection requires that all projection lines converge on a single point, commonly called the VIEWPOINT. This ultimately determines the "point of view" of the picture.[30]

Here it is often convenient to speak of the PROJECTIVE COUNTERPARTS of a point on the picture plane: these are the points in the scene to which it is connected by a projection line.[31] Typically, not *all* projective counterparts of a point in the picture are relevant to the final image. Only those that are ACCESSIBLE from the viewpoint in the following sense are of consequence: they can be reached by a projection line without passing *through* a surface in the scene. Projection lines from picture points to their accessible counterparts can be thought of as "lines of sight", since, like sight, they link the viewpoint to those surfaces in the scene which are not occluded by any other surface in the scene. For visual clarity, I have indicated only those projection lines below which connect the viewpoint to the accessible corners of the object in the scene.

---

[27]The mathematical details of the following discussion are based on presentations in Sedgwick (1980), Hagen (1986, ch. 2-5), and Willats (1997, ch. 2, 5).

[28]Willats (1997) correspondingly organizes methods of projection into "projection systems" and "denotation systems". See Durand (2002) for an extension and critical discussion of this approach.

[29]Admittedly, "picture plane" is a misnomer; I mean it to refer to a contiguous finite region of a plane.

[30]The viewpoint is sometimes called the "station point"; 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.

[31]Some points in the picture, for example at the periphery, may not have projective counterparts in the scene.

The next stage in the projection is specified by the denotation condition— the rule by which features of the scene are mapped onto features of the picture plane along the lines of projection. In this case I'll use a simplified version of the rule characteristic of *line drawing*: a point in the picture is colored red (for example) just in case it has a projective counterpart that lies on an accessible edge in the scene. I'll define an EDGE, for now, as the intersection of two flat planes. This rule is applied to every point on the picture plane; those points whose color is thereby left undefined assume the background color of the "page." If the picture plane is now seen face on, as at right, it reveals a side view of the cube, drawn in red.

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 which are in fact parallel are now represented by converging lines.

| A | B | C | D |

Each of the alterations just described changes the *way* the scene is represented, but none reveal 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 principle 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*).

| E | E |

Images (*A*)-(*E*) are produced by the same general method of projection, their differences owing 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. But *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.

Let me summarize. A method of projection is an algorithm for mapping a three-dimensional scene onto a two-dimensional picture plane. All methods of projection are defined in terms of an array of projection lines where exactly one passes through each point on the picture plane. The method of projection for perspective line drawing can then be identified first by how it structures these projection lines, and second by how it uses them to determine a mapping of features from scene to picture, its projection and denotation conditions respectively:

Projection: Every projection line intersects the viewpoint.

Denotation: A point in the picture is red if and only if it has a projective counterpart

in the scene which lies on an accessible edge.

Together, these conditions define a method of projection such that, given a scene and a projective index— that is, the position of the viewpoint and picture plane— the method determines a unique projected image. Thus where S is a scene and j a projective index, we may speak of *the* perspective projection of S relative to j. An explicit formalization of the method of perspective projection is provided in §4.

Other methods of projection can be derived by independently modifying the projection and denotation conditions. For example, methods of *parallel* projection differ with respect to the projection condition: the viewpoint of perspective projection is replaced by a view *plane*, such that all projection lines intersect it at right angles. On the other hand, methods of *color* projection differ with respect to the denotation condition: instead of correlating red points in the picture with edges in the scene, a rudimentary color system might hold that points in the picture picture simply have the same color as their accessible projective counterparts in the scene. And so on. Formal definitions of these alternatives are also included in §4. For now, I turn to integrating the method of perspective projection into a theory of accuracy.

2.2   Semantics for the system of perspective line drawing

The method of perspective *projection* just illustrated is an algorithm for deriving pictures from scenes. The *system* of perspective line drawing, on the other hand, determines a standard of pictorial accuracy. I shall say that every possible composition of red lines on a bounded white plane corresponds to a picture belonging to the system of perspective line drawing.[32] For any such picture and any scene, a definition of accuracy in this system determines whether the picture accurately depicts the scene. According to the Geometrical Projection Theory, the system of perspective line drawing defines accuracy in terms of perspective projection in roughly the following way. For any picture P in the system and scene S:

(5) P accurately depicts S in the system of perspective line drawing if and only if

   P is a perspective line drawing projection of S.

To state this formula more precisely, let L be the system of perspective line drawing suggested by the projective method above. Where a picture P accurately depicts a scene S in system L, we may write: "P accurately depicts$_L$ S". Now, since the specified method of perspective projection was entirely determinate, it can be represented by a mathematical function, what I will call a PROJECTION

---

[32]This does not imply that every marking of red ink on paper specifies a picture in this system, just as merely sounding like a word of English is not sufficient for being a word of English. The sign must be created with the proper intent under the appropriate conditions.

FUNCTION. A projection function takes as inputs (i) a scene and (ii) a projective index, and outputs an image— the unique perspective projection of the scene relative to the index. Where $\text{proj}(\cdot)$ is the perspective projection function, S a scene, and j a projective index I shall say that $\text{proj}(\cdot)$ applied to S and j returns an inscribed picture plane P, that is, $\text{proj}(S, j) = P$. Then a first attempt to render (5) in more exact terms may be articulated. For any picture P in L and scene S:

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

That is: P accurately depicts S in L if and only if the perspective projection of S relative to j is P. The problem with this definition is that it makes the projective index j part of the system itself. This would mean that all accurate perspective images would have to be drawn from the same viewpoint. But this is clearly false: both (8) and (9) below are accurate perspective depictions of S, but from different viewpoints. The solution is to allow some variability in j. One way to achieve this is 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 $\text{proj}(S, j) = P$.

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 with this analysis, but my interest here is in a more specific notion. According to this alternative, accurate depiction only obtains in relation to a projective index, and artists always intend that their pictures be interpreted relative to a particular 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) respectively, but both intend to represent the scene from the same vantage point, then only one succeeds in accurately depicting the scene. To accommodate this idea, I now analyze depiction as a relation between a picture and a scene relative not only to a system of depiction, but *also* to a projective index.[33] Where P accurately depicts$_L$ relative to index j, I will write "P accurately depicts$_L$ at j". Then for any picture P in L, scene S, and index j:

---

[33]See Hyman (2006, p. 82).

109

(10) P accurately depicts$_L$ S at j iff $\text{proj}(S, j) = P$.

Unfortunately, this implementation of the Geometrical Projection Theory is too demanding in a subtle but important way. It is not the case that an accurate picture can always be derived by projection from its subject; sometimes a picture is accurate when it is merely *similar* to some other picture which *can* be so projected. This principle is plainly illustrated by considerations of scale. Suppose that two artists intend to draw a cube from a given vantage point, and produce (11) and (12) below. Both images are perfectly accurate representations of S, despite their difference in size. Yet this flexibility is not allowed by (10), for $\text{proj}(S, j)$ yields a picture plane of fixed size, and any deviation from this picture leads to a failure of the biconditional. To accommodate this, (10) should be revised; rather than requiring that the picture itself be derived from the scene by projection, I instead require that some *copy* at the same or different scale be so derived.

(11)   (12)

Other systems of depiction require other dimensions of freedom besides scale, including flexibility with respect to choice of color, line smoothness, and in extreme cases, overall metrical structure.[34] We need a general way of handling these various cases. To do this, let us introduce a FLEXIBILITY FUNCTION, $\text{flex}(\cdot)$, which takes a picture and returns the set of suitably related pictures for that system of depiction. Now, instead of requiring for accuracy that the projection of S be *identical* with P, we merely require that the projection of S be *within* $\text{flex}(P)$, that is, the class of pictures appropriately similar to P. On this approach, the flexibility and projection functions play against each other, for while the projection function constrains the conditions on accurate depiction, the flexibility function weakens these constraints. But note that the idea here is not to introduce an open-ended parameter; rather, we must define specific flexibility functions for for specific systems. Thus, for L, we might define $\text{flex}(\cdot)$ as follows, for any picture P in L:

(13) $\text{flex}(P)$ = the set of pictures which are scale variants of P.

---

[34]Some color systems tolerate constrained variations in color among accurate depictions of the same scene. (Thanks to Dom Lopes for bringing this to my attention.) Some line drawing systems tolerate limited deviations from perfectly straight lines in the depiction of straight edges. Subway maps exhibit constrained deviations from the metrical structure of a strict projection, while preserving quasi-topographic features like station-betweenness. (Willats (1997, pp. 70-7); Giere 2004, §4). An alternative approach is rather natural at least the case of wobbly lines. Pictures, considered as abstract objects, may differ in interesting ways from the marks which express them— just as sentence structure can depart from the "surface structure" of an utterance. Then we might think that it is only marks on the page which are wobbly; they express pictures composed of perfectly smooth lines. I have no particular complaint with this approach, but presumably it should also be extended to scale and color. But then pictures are so abstract— lacking determinate size, for example— that they become difficult to discuss without excessive technicality.

With this amendment, we arrive at the final analysis:

(14)   **Accuracy in Perspective Line Drawing**

For any picture P in L, scene S, and index j:

P accurately depicts$_L$ S at j iff $\text{proj}(S, j)$ is in $\text{flex}(P)$.

This is a definition of accuracy for the system of perspective line drawing. It follows from this analysis plus my initial definition of depictive content that the content of a picture in perspective line drawing is, setting aside the flexibility function, the set of scenes from which it can be derived by perspective projection. To make this statement precise, one must account for the observation that accurate depiction holds not only relative to a system of depiction, but also to a choice of projective index. A natural view here is that pictorial content is defined relative to such an index as well.[35] Then just as linguistic content is determined in part by isolated features of conversational context, pictorial content depends on communicative context to specify an intended projective index. Formally, the content of a picture P in system L relative projective index j is, roughly, the set of scenes P could be derived from by perspective projection relative to j. More exactly:

(15)   **Content in Perspective Line Drawing**

For any picture P in L, scene S, and index j:

the content of P in L relative to j = the set of scenes S such that: $\text{proj}(S, j)$ is in $\text{flex}(P)$.

I have thus demonstrated, for one simple system of depiction, how the Geometrical Projection Theory determines a pictorial semantics— a systematic and conventional mapping from pictures to content. The analysis may be extended to other systems of depiction, and so to depiction in general, by defining alternative projection and flexibility functions. Once again, the details of such definitions are included in §4.

2.3   Natural systems of perspective line drawing

By demonstrating that it is possible to state a precise semantics for one simple system of line drawing, I have shown that unqualified skepticism about the Pictorial Semantics Hypothesis is un-

---

[35]This is not the only way. We might also define the content of P in L, independent of a projective index as the set of scene-index pairs $\langle S, j \rangle$ such that $\text{proj}(S, j)$ is in $\text{flex}(P)$. The difference between this view and the one expressed in the text is analogous to the difference between eternalist and temporalist theories of linguistic content. (Richard 1981) In each case there is a question about whether content is fixed by the value of given contextual parameter, or whether it generalizes over the possible values of this parameter. Kulvicki (2006*a*, pp. 57-9) holds yet another view, which derives from his assumption that content is determined by projective invariants; he defines the content of a picture, independent of projective index, as the set of scenes it could be projected from according to *some* index. (The resulting set is the set of scenes which share projective invariants with the picture itself.) But at least for the case of linear perspective, this proposal seem to me too crude: it misses the fact that by gaining information about a picture's viewpoint, we are able to extract more specific information from the content of the picture. Both the position in the text and that sketched above can accommodate this observation. For the case of curvilinear perspective, the view faces even more fundamental problems, as I argue in Greenberg (2010*a*).

warranted. But even this may not satisfy the entrenched doubter: perhaps artificially constructed systems of depiction like *L* have semantics, she admits, but the prospects seem dim for the unregimented depictive styles actually found in ordinary communication. The aim of this section is to offer a partial response to such a challenge. First I'll highlight the sense in which the skeptic is right: by design, *L* is only a fragment of any natural system of line drawing. But then I'll show how recent developments in computational perceptual science can be used to expand this fragment in promising ways. These findings give us reason to hope that, for all their complexity, even natural forms of depiction are undergirded by systematic rules, capable of mathematically precise articulation.

Following a precedent of linguistic theory, we can roughly distinguish NATURAL from ARTIFICIAL systems of depiction. Natural systems have evolved organically in ordinary communication, while artificial systems are introduced and disseminated by solely by explicit definition. For the scientist of human behavior, natural systems are the primary quarry, but because of their stipulative origins, the semantics of artificial systems are easier to define. For this reason, well-understood but artificial systems like *L* serve as useful stepping stones on the path toward the analysis of the not-yet-understood natural systems. Thus, while *L* is recognizable as a system of perspective line drawing, it is only a fragment of any natural drawing system, for it is suitable for interpreting only depictions of geometrical solids with flat sides.

The chief shortcoming of *L* lies in the denotation condition for the method of perspective projection upon which *L* is based, reproduced here:

(16)   A point in the picture is red if and only if it has a projective counterpart in the scene which lies on an accessible edge.

Recall that an EDGE was defined as the intersection of two flat surfaces. As a consequence, *L* cannot be used to depict curved objects like spheres, since spheres contain no intersecting flat surfaces. It doesn't help to expand the definition of edge to include what are sometimes called "creases"— abrupt intersections of any kind of surfaces— because spheres contain no such discontinuities. Indeed, given that spheres are completely uniform, it may seem puzzling how features of a sphere could be mapped to lines in a picture at all.

The answer is that we must introduce a very different concept of edge— one which applies not to intrinsic features of a scene, but instead to relational properties defined relative to the viewpoint.[36] Such is the notion of CONTOUR. Intuitively, a contour is any visible perimeter in a scene

---

[36]The idea of a "view-dependent" edge is described by Durand (2002, §6.2). Philosophical writing on depiction has

which occludes features of a scene behind it. Formally, we can define a contour as that region of a surface which is tangent to a projection line extended from the viewpoint— that is, the region of a surface which a projection line "barely touches", as illustrated below. With this, we can define the denotation condition for a revision of L in such a way that yields successful renderings of a sphere. (Here a CONTOUR LINE is just a line in the picture which corresponds to a contour in the scene.)

(17)   A point in the picture is red if and only if it has a projective counterpart in the scene which lies on an accessible contour.

Yet contours cannot replace edges in the analysis of line, for while (17) would successfully capture the *silhouette* of a cube, it would fail to depict any of the cube's *internal* edges— those edges facing the viewpoint which occlude nothing from view. Instead, we must combine the two definitions, allowing that a point is inscribed on the picture plane when it has a projective counterpart in the scene which lies *either* on an accessible edge *or* an accessible contour. This definition yields successful depictions both of spheres, like that above, and of cubes, as in (18).

But even this innovation will not suffice. Common line drawings in perspective contain certain kinds of line that correspond neither to edges in the scene, nor to contours. Instead, they merely *suggest* the shape of the object depicted. DeCarlo et al. (2003) dub these SUGGESTIVE CONTOURS. Consider, for example, the suggestive contour line tagged in (19). It does not indicate the sharp intersection of two surfaces (an edge) nor is it naturally interpreted as describing a jutting cranial ridge hiding a valley behind it (a contour). Instead it merely suggests a certain structural arc at a certain location on Obama's head.

employed the related notion of "occlusion shape", or its close kin (Peacocke 1987; Hopkins 1998, pp. 53-7; Hyman 2006, pp. 75-9).

(18)   (19)

contour line *and* edge line

edge line

suggestive contour line

edge line

contour line

One might think that we have here finally reached the limits of systematic inquiry, that line drawing is ultimately organic and improvised, and that a semantic analysis will never be able to account for such expressive fluctuations as suggestive contour. But this skepticism is premature. In a recent paper from the field of computer graphics, DeCarlo et al. (2003) propose an ingenious analysis of suggestive contour lines. They offer an algorithm according to which suggestive contour lines are inscribed not when they correspond to *actual* contours relative to the viewpoint, but when they *would* correspond to contours under constrained *permutations* of the viewpoint. On this account, the presence of the suggestive contour line in (19) indicates that the shape of Obama's head at that point is such that, if the viewpoint were shifted slightly to the left, that region would emerge as a genuine contour. Such "nearby" contours obviously reflect structurally significant features of a scene; the suggestion is that, in the system of perspective line drawing, the rules for rending contours are naturally extended to capture this additional structure.

These considerations are only a small indication of the challenges facing the analysis of natural line drawing. The moral is that there remain substantive matters to work out in the elaboration of pictorial semantics, from details to foundations— and not just in the theory of line reviewed here. Such developments will inevitably require a steep and steady descent from the relative simplicity of definitions like that of L into nuance and complexity.[37] On the other hand, mathematically precise accounts of ordinary drawing systems, like that of DeCarlo et al. (2003), show that pictorial semantics is not merely an analytic contrivance practiced on artificial examples. Despite the wide-spread skeptical attitude that depiction is basically unsystematic, such results show the potential for a progressive and rigorous analysis of natural systems of depiction. It remains to be seen whether this method can be extended to the myriad systems of depiction in use across history and culture. But the hypothesis of this chapter is that it can.[38] I hope to have shown that this hypothesis is at least plausible.

---

[37]For example, Pirenne (1970, pp. 116-35) and Hagen (1986, pp. 170-1) document systems in which the operative method of projection varies according to the *type* of object being depicted.

[38]Thanks to Tendayi Achiume for helping me see both the limitations and import of the present work in this way.

3 Conclusion

Carnap (1963, p. 938) once described language as "an instrument... like any other... useful for a hundred different purposes." Arguably the most basic such purpose is the mobilization of public signs to disseminate private beliefs. According to the the Pictorial Semantics Hypothesis defended in Chapter 1, systems of depiction embody a parallel technology of communication. Both languages and systems of depiction are based on *semantics*: systematic and conventional mappings from signs to representational content. By contrast, the thesis of this chapter, the Geometrical Projection Theory, suggests that these semantics are profoundly divergent. Whereas the semantics of language are based on arbitrary associations and rules of *logical* transformation (intersection, function application, and the like) the semantics of depiction, on this account, is based on non-arbitrary rules of *geometrical* transformation.

In informal discussion, pictorial representation is often described as "natural", in analogy with the phenomena of "natural meaning" exhibited shadows, tree rings, or foots prints.[39] On the view of depiction as projective transformation, accurate pictorial representation falls between linguistic meaning and natural meaning in a precise and intuitive way. Like linguistic meaning, depiction is conventional; there are many possible systems of depiction which must be selected amongst for successful communication. But like natural meaning, and unlike language, depiction is not mediated by arbitrary associations of sign and content. Yet no such arbitrary correspondences are exploited by natural meaning or depiction. Instead, depiction and natural meaning alike convey information in a more direct and unmediated way, in virtue of general, rule-based transformations. In the case of natural meaning these rules are the laws of physics; foot prints and shadows are related to their sources by causal transformation. The rules underlying systems of depiction are the abstract laws of geometry; the transformations are geometrical projections.

These conclusions provide compelling evidence that semantic analysis can thrive outside of the linguistic arena. They suggest that general semantics, the inclusive study of signs, is a viable endeavor. Such a study would apply the method of semantic analysis across the representational spectrum, covering languages, diagrams, pictures, and all the varying complexity in which these modes can be combined. It would bring into the light phenomena rarely studied, and shed new light on those which have come to seem familiar.

This concludes the informal contribution of this chapter. In the next section I develop a formal

---

[39]The term "natural meaning" is due to Grice (1957, p. 378).

semantics for the system of linear perspective line drawing *L*.

4   Formal semantics for *L*

This section partially formalizes the semantics for perspective line drawing presented in §3. I begin by modeling the key notions of picture, scene, projective index, and perspective projection. These components are then integrated into definitions of accuracy and content for the system of perspective line drawing, which are subsequently shown to be fully compositional. My aim in what follows is to formalize the relevant notions only as much as is required to show that a fully precise and algebraic analysis is possible, but not so much as to provide every detail of such an analysis.

As we embark, it is necessary to remember that formal models, in linguistics and here, are by definition abstract, mathematical representations of real phenomena. Consider two prominent examples: (1) Linguists typically model *sentences* as abstract logical structures which they distinguish from physically instantiated *utterances*. Here too I will model *pictures* as abstract geometrical objects, distinct from physically instantiated *inscriptions*.[40] (2) Mathematical linguists define languages in part by infinitely large sets of sentences— not merely those that have actually been uttered, but all the possible sentences of that language. In the same way, I will define systems of depiction in part by infinitely large sets of pictures. From the perspective of naturalistic metaphysics, such formally-motivated definitions may seem problematic. But I flag these issues only to set them aside: the definitions here are intended only to give my account precise articulation while abstracting away from irrelevant details.

The analysis employs a basic geometrical framework. I will be discussing Euclidean spaces of 2 and 3 dimensions, of finite and infinite extent, which I will refer to as 2-SPACES and 3-SPACES respectively. For every dimensionality, there are infinitely many distinct spaces of that dimensionality, individuated by the chromatic properties of and spatial relations among the points that make up each space. A REGION is a part of a space, such as a point, line, line segment, plane, convex solid, and so on. A COLORED REGION is a region where every point in that region is associated with exactly one color, but I also allow for colorless regions.[41] An OBJECT is defined as a connected colored region. Finally, an EMBEDDING is a function which maps a region in one space to an isomorphic

---

[40]Willats (1997, pp. 98-100) argues that we should distinguish pictures from inscriptions on the ground that not all marks on a page are semantically significant. For example, there may be stray marks or smudges on a piece of paper that are irrelevant to interpretation. One can imagine an entire science devoted to identifying the correct mappings from marks on paper to lines in the picture. But such details only distract from the project of semantics, which is to provide general rules of interpretation.

[41]Colors can be modeled as real numbers associated with positions on a color wheel.

region in another space. The only kinds of embeddings relevant here are those that preserve metrical relations between points and their chromatic properties. An embedding may occur between spaces of the same or different dimension. For example, one embedding may locate lines of a particular 2-space within a particular 3-space; another may locate lines of a particular 2-space within a distinct 2-space.

Definition of picture and scene

Let us define a PICTURE as a colored, rectangular segment of plane in a 2-space.[42] The decision to define pictures in this way is obviously an idealization, for it rules out pictures which are round, pictures which are not perfectly flat, and so on. Perhaps most artificially, I assume that the color of regions in a picture can be characterized independent of viewers or viewing conditions.[43] We shall model SCENES in a parallel fashion. At the outset we defined scenes as time- and location-centered worlds. We will now assume that every such scene determines a 3-space, which may be populated by objects, or not, in any way. We will refer interchangeably to scenes, considered as time- and location-centered worlds, and the spaces they determine.

Definition of projective index

Recall that in order to construct a perspective projection of some scene one must specify both the position of a viewpoint from which the projection lines are drawn, and the position and orientation of a picture plane onto which the image of the scene is projected. The specification of these parameters is the job of the PROJECTIVE INDEX. Different kinds of projective indices are required for different methods of projection. In the case of perspective projection, the projective index consists only of two components $\langle v, r \rangle$, where v is a VIEWPOINT FUNCTION, and r is a PICTURE POSITION FUNCTION.

The viewpoint function embeds a unique viewpoint— understood as a point— within the 3-space of a given scene. Formally, v is a function from scenes to points in the 3-spaces of those scenes. Where S is a scene, we shall refer to v applied to S, the viewpoint in S, as "v_S". The picture position function correspondingly embeds a picture— understood as a colored region within a 2-space— into the 3-space of a given scene. Formally, r is a function from pictures and scenes to colored flat

---

[42]A richer representation of pictorial structure would parse pictures as arrangements of *parts*. Such an account would allow the theorist to classify "ambiguous" images like the Necker Cube as *syntactically* ambiguous, rather than merely semantically unspecific. (Necker 1832; Wittgenstein 1921/1997, §5.5423) Associating pictures with part structure may also play a role in intelligent interpretation of abstract images. (Mi, DeCarlo, and Stone 2009; DeCarlo and Stone 2010)

[43]This last assumption would be unsustainable in any careful study of color painting or photography. But here my goal is to take on as few complications as possible. Since I will only seriously examine systems of line drawing, such simplistic modeling choices will help, not hinder understanding. I hope that readers interested in color depiction can see how the definitions here could be extended to treat their preferred subject matter.

regions in the 3-spaces of those scenes. Where S is a scene and P is a picture, we will refer to the picture embedded in S by r as "P_{r_S}". The following diagram shows a possible arrangement of the viewpoint and a picture P, as specified by v and r relative to the scene S.

Definition of perspective projection

I now define perspective projection in more exact terms. All forms of projection are defined in terms of an array of projection lines which pass through every point on the picture plane. Relative to this array of projection lines, methods of projection are defined by their PROJECTION and DENO-TATION conditions. In §3.1, we described these conditions for the method of perspective projection as follows:

Projection: Every projection line intersects the viewpoint.

Denotation: A point in the picture is red if and only if it has a projective counterpart in the scene which lies on an accessible edge.

To formalize this proposal, we begin by defining the array of projection lines passing through every point on the picture plane, which these conditions presuppose. Where P_{r_S} is a picture situated relative to the space of the scene:

(20) for every point p in P_{r_S}, there is a projection line l intersecting p

We will treat these projection lines as rays. Then we can state the projection condition (i) as the requirement that all projection lines have their endpoints at the viewpoint v_S. That is, for any projection line l:

(21) l is a ray with its endpoint at v_S

The diagram shows three projection lines $l1$, $l2$, and $l3$ passing through the picture plane at points

$p1$, $p2$, and $p3$ respectively. Of course, these are only a representative selection of the total array; according to (20), some projection line passes through *every* point in the picture. As per (21), each projection line has its endpoint at the viewpoint. The interactions between these lines and the scene will become relevant shortly.

Next, the denotation condition (ii) specifies a rule for mapping edges in the scene onto lines in the picture plane. For line drawing (with red ink), the rule is that a point in the picture is red just in case it has some projective counterpart in the scene which is (a) accessible and (b) lies on an edge. Recall that a picture-point's "projective counterparts" are just those points in the scene which lie along the projection line which passes through that point. Then we may state the denotations condition as follows, where S is the scene, p is any point on the picture plane, and l is the projection line intersection p:

(22) p is red iff there is a point s on l such that:

   (a) s is accessible from v_S along l;

   (b) s lies on an edge in S.

A point in the scene s is ACCESSIBLE from the viewpoint v_S along a line of projection l just in case v_S and s can be connected by l without passing through any other point in the scene. An EDGE is the line segment at the intersection of two flat surfaces. (Note that I assume that every point in a picture is either red or white, thus the schema above determines the color for every point in the picture.[44]) The diagram below illustrates the application of (22) to each point on the picture plane, along with the resulting picture at right.

Note how the three projection lines included here interact with the scene in different ways. First, $l1$ fails to intersect the scene at all; thus $p1$, the point at which $l1$ passes through the picture is left white. Next, $l2$ intersects the scene at two points; the first, $s1$ is accessible from the viewpoint, but does not lie on an edge; the second, $s2$, lies on an edge, but is not accessible from the viewpoint; thus $p2$, the point at which $l2$ passes through the picture is also left white. Finally, $l3$ intersects the

---

[44]The definition here fails if this assumption is relaxed, since on that condition, two pictures with different background colors would be compatible with the scene and index. Of course, the "whitespace" condition can be written into the definition of the projection function, but I find this distracting.

scene at $s3$, which lies on an accessible edge; thus $p3$ is colored red.

Finally, by combining each of these components, we may define perspective projection. For any picture P, scene S, and projective index $j = \langle v, r \rangle$:

(23)   The perspective projection of S relative to j, $\text{proj}(S, j) = P$ iff

for every point p in P_{r_S}, there is a projection line l intersecting p such that:

(i) l is a ray with its endpoint at v_S;

(ii) p is red iff there is a point s on l such that:

(a) s is accessible from v_S along l;

(b) s lies on an edge in S.

That is: a picture P is a perspective projection of a scene if and only if for each point p in P, as embedded in S by r, there is a projection line l such that (i) l originates at v_S and passes through p; and (2) p is red just in case there is some point s in the scene lying along l such that (a) s is accessible from v_S along l, and (b) s happens to lie on an edge in the scene.

By independently varying the projection and denotation conditions, one can derive a spectrum of projective methods besides that of perspective line drawing. Consider first the projection condition (i), which states that every point p in the picture defines a projection line l which intersects the viewpoint v_S and p. This simple stipulation characterizes all forms of perspective projection. Other methods of projection define the projection lines in other ways. For example, in the method of parallel projection, the projection lines are defined as rays extended perpendicularly from a *plane*.[45] Letting o_S be such a plane, we could then write the projection condition for such as method as follows, where l is any projection line passing through a point in the picture P_{r_S}:

(24) l = the ray extending perpendicularly from o_S

The denotation condition (ii) describes how features of the scene determine features of the

---

[45]Note that this would require changes to the definition of the projective index.

picture plane, given a certain projective mapping. The condition has two components. The first (a) states *which* of a picture-point's projective counterparts in the scene it is defined by; the second (b) states *how* it is defined by these counterparts. Alternatively, the first states *which* of a picture-point's projective counterparts it "says something" about; and the second states *what* it "says" about them.

In most cases, (a) isolates those projective counterparts which are accessible from the viewpoint along the line of projection. Occasionally, however, this condition is modified to reveal occluded edges and faces. This results in the style of so-called "wire-frame" pictures like (26) below. To achieve this effect, we may weaken condition (ii) by eliminating the restriction to accessible points in S. The denotation definition then states that a point in the picture is red just in case it has *some* projective counterpart in the scene which lies on an edge. The second part of the denotation condition, (b), determines the relationship between the spatial or chromatic properties of a picture-point's projective counterparts, and the chromatic properties of the picture point itself. In the definition above, the picture-point is determined to be red just in case its accessible counterparts lie along an edge. But we may introduce some simple variations of this condition. To redefine the projection function as one that yields silhouettes, we may replace the denotation condition with (27). Then a picture-point is determined to be red just in case its accessible counterparts lie on *any* surface in the scene. To reveal the faces of the cube, but distinguish them from the edges, we must introduce two conditions— one for the edges, and one for the faces— as in (28).

(25) p is red iff there is a point s on l such that:

    (a) s is accessible from v_S along l;

    (b) s lies on an edge in S.

(26) p is red iff there is a point s on l such that:

    (a) ——-

    (b) s lies on an edge in S.

(27) p is red iff there is a point s on l such that:

    (a) s is accessible from v_S along l

    (b) s lies on a surface in S.

(28)   *p* is red iff there is a point *s* on *l* such that:

    (a) *s* is accessible from v_S along *l*;

    (b) *s* lies on an edge in *S*.

    p is pink iff there is a point *s* on *l* such that:

    (a) *s* is accessible from v_S along *l*;

    (b) *s* lies on a face in *S*.

These kinds of associations between surface type and point color are quite arbitrary. An edge could be mapped to a point of some other color, or something other than a point, such as a patterned region. But the relation between point color and properties of the scene may also be much more direct. For example, in a system reminiscent of color painting, the color of the points in the picture are simply inherited from the color of points in the scene:[46]

(29)   for any color *C*:

    p is *C* iff there is a point *s* on *l* such that:

    (a) *s* is accessible from v_S along *l*;

    (b) *s* lies on a surface in *S* and *s* is *C*.

In general, modifications to the denotation condition result in changes to the kind of information a picture records about its subject. Silhouette drawings are probably the least informative of the options just considered— they record merely the outline shape of a scene. Color projections contain the same outline information, but add to this the colors of the accessible surfaces in the scene. In the original system of edge-to-line projection, the picture reflected no color information about the scene, but it did contain additional spatial information not explicitly recorded in the color image— to wit, the location of edges in the scene.

Finally, it must noted that the original definition of perspective projection from §3.1 results in images composed of true *lines*, colored regions of zero width. Ideally, however, the definition should allow us to adjust line thickness as a parameter.[47] There are various ways to achieve this; the approach adopted here will require a slight change to the overall definition. The idea is that a point is colored red— not, as before, if it has a projective counterpart on an appropriate edge— but instead, if some *nearby* point has a projective counterpart on an appropriate edge. The result is that, for every one of the points that was assigned red by the original definition, there now will be a cloud of nearby points also assigned red. The metric of nearness defines the size of the "pen nib".

---

[46]Willats (1997, ch. 6) argues that many forms of color depiction map color features of *light* from the scene to points on the picture plane, rather than intrinsic color features of the scene itself (if such there be).
[47]See Durand (2002, §6.3) for discussion.

To formalize this suggestion, let $n(\cdot)$ be a NIB FUNCTION which takes a picture P and point p in P and returns a circular region of P which includes p. Different nib functions may specify the size of this region differently. We then revise the definition of perspective projection as follows. For any picture P, scene S, and projective index j:

(30)   The perspective projection of S relative to j, $\mathrm{proj}(S, j) = P$ iff

for every point p in P_{r_S}, there is a projection line l intersecting some p' in $n(P_{r_S}, p)$ such that:

(i) l is a ray with its endpoint at v_S;

(ii) p is red iff there is a point s on l such that:

(a) s is accessible from v_S along l;

(b) s lies on an edge in S.

From perspective projection, we now turn to semantics.

Definition of system of depiction

A SYSTEM OF DEPICTION has three components. It specifies (i) the set of all possible pictures belonging to that system; (ii) the set of projective indices compatible with that system; and (iii) an accuracy function. The first component, the set of all possible pictures in the system, effectively fixes the system's syntax. In principle, the definition of this set might be quite complex, but at least for the systems included here, it is not. For our purposes, a picture belongs to the system of red on white line drawing just in case every region of the picture is a red region of a given minimum diameter or a white region.[48]

The second component of a system of depiction describes the range of projective indices compatible with that system. As we've seen, different systems of depiction require different kinds of projective indices. In some cases, systems put only technical constraints on indices, like substituting a plane of projection for the viewpoint characteristic of perspective projection. Other constraints are more "stylistic": some systems differ only in the orientation of the picture plane relative to

---

[48]Researchers in computer science have sometimes claimed that the constraints on well-formedness for line drawings are much more substantive. (e.g. Huffman 1971, p. 313) They have deemed apparently uninterpretable drawings, like the Penrose triangle to be ill-formed, and have produced intricate formal theories of well-formedness. But I think this diagnosis of the situation is premature. In such cases it is possible to explain this uninterpretability semantically, by the fact that there are no objects that such a picture could depict. My methodological assumption is that it is always preferable to explain uninterpretability as semantic impossibility, rather than by fiat of syntactic rule. If I am right, the results of research on impossible drawings from computer science are significant, but they are not what their discoverers claim them to be. Rather, this field is the analogue of proof theory for pictures. In logic, proof-theory attempts to isolate the syntactic transformations which will produce or preserve semantic properties such as contradiction, theoremhood, and contingency. Similarly, research into impossible drawings attempts to isolate graphic transformations which will produce or preserve the semantic properties of impossibility and possibility. (As Wittgenstein (1921/1997, §2.225) observed, drawings cannot be *a priori* true; thus they cannot be theorems.) Further discussion of this question must be postponed.

whatever is considered the "dominant" plane of the scene.[49]

Finally, the accuracy function simply formalizes the relation of accurate depiction. It maps a picture P, scene S, and index j to 1 if and only if P perfectly accurately depicts S at j. (Thus the accuracy function is the correlate of the valuation function familiar from the semantics of propositional logic.) In this chapter, we are concerned only with perfect accuracy and its failure, but in principle one might set out to give conditions for all degrees of accuracy between 0 and 1. The main task in analyzing a system of depiction is to define its accuracy function; here we shall do this implicitly, by specifying necessary and sufficient conditions on perfectly accurate depiction.

Semantics for *L*

In §3.2 we presented a schematic definition of accuracy for the system of simple line drawing in perspective, L:

(31) P accurately depicts$_L$ S at j iff $\mathsf{proj}(S, j)$ is in $\mathsf{flex}(P)$.

And the correlate definition of content:

(32)   the content of P in L relative to j = the set of scenes S such that: $\mathsf{proj}(S, j)$ is in $\mathsf{flex}(P)$.

To obtain a explicit semantic analysis of L, we need only substitute into each of these formulae the definition of the projection function $\mathsf{proj}(\cdot)$ provided above, along with the definition of the flexibility function $\mathsf{flex}(\cdot)$ given in §3.2:[50]

(33) $\mathsf{flex}(P)$ = the set of pictures which are scale duplicates of P.

Alternative systems of depiction can be derived by suitably redefining the projection function and flexibility functions. We have already considered several variations of the projection function corresponding to familiar alternative systems of depiction. Such variations can be multiplied by altering the flexibility function so as to allow some degree of freedom with respect, for example, to straightness of line or choice of color.

---

[49] As Hagen (1986, pp. 242-7) notes, in some systems there are clear conventions about the definition of the dominant face, dependent on object type. Modeling such a system would require a richer representation of scenes than I have elected for here.

[50] The substitution is achieved perhaps more perspicuously with following equivalent reformulation of (33): P accurately depicts$_L$ S at j iff there is a P' in $\mathsf{flex}(P)$ such that $P' = \mathsf{proj}(S, j)$.

Compositional semantics for *L*

The semantics on offer here describes a general rule for interpreting pictures. The ability to interpret novel pictures is straightforwardly explained as the application of the general rule to novel particular cases. But it has been argued in the linguistic case that for this explanation to work, a semantics must be COMPOSITIONAL, that is, the content of the whole sign must be determined by the content of its parts and the way they are put together.[51] In this section I show how this principle may be realized for the case of pictures.

My treatment is shaped by two observations, both adapted from Fodor (2007, pp. 108-9). First, I assume that every contentful part of a picture is also a picture. This is very different, note, from the linguistic case, where sentences have contentful non-sentential parts. A compositional semantics for pictures must show how the content of the whole picture can be derived from the content of pictures which are its parts. Second, I assume that, unlike complex linguistic expressions, pictures have no "canonical decomposition". That is, the content of a picture may be compositionally derived from *any* division of the picture into parts– so long as the parts are members of the same system as the whole. For any given picture, there are many possible decompositions; at the limit, at least for *L*, one may decompose a picture into points. Illustrated below are three arbitrarily selected possible decompositions of a given image.

To begin, let $D = \langle \mathbf{P}, \mathbf{m} \rangle$ be a DECOMPOSITION of a picture P relative to a system I. Intuitively, $\mathbf{P}$ is the set of picture-parts that results from cutting P up into pieces, and $\mathbf{m}$, which we'll call the ARRANGEMENT, shows you how to put them back together again. Formally, $\mathbf{P}$ is a sequence of finite, colored 2-spaces belonging to I; where the member of $\mathbf{P}$ at the nth position, is referred to as "p_n". Next, the arrangement $\mathbf{m}$ is a sequence of embeddings, each of which embeds a ranked member of $\mathbf{P}$ into the same 2-space; again the member of $\mathbf{m}$ at the nth position, is referred to as "m_n". $\mathbf{P}$ and $\mathbf{m}$ are coordinated in the following way: for every m_n in $\mathbf{m}$, $m_n(p_n)$ is a disjoint subset of P, and the union of every embedding in $\mathbf{m}$ applied to its corresponding member in $\mathbf{P}$ is identical to P— that is, $\bigcup_{i \leq n} m_i(p_i) = P$.[52] The order of parts in $\mathbf{P}$ doesn't matter; but since it may be

---

[51]See Fodor and Lepore (2002) and Szabo (2010).

[52]This way of defining the members of $\mathbf{m}$ requires that the spaces in question encode absolute position. To accommodate spaces that support only relative positions, $\mathbf{m}$ could be redefined so as to map *pairs* of members of $\mathbf{P}$ to relative locations in a single 2-space.

necessary to map intrinsically identical regions to different locations, **m** distinguishes the different mappings by their order. According to the principle of compositionality, the content of the whole is determined by the content of the parts and the way they are combined; we will capture the way they are combined by the arrangement **m**.

The diagram below illustrates an example decomposition. Here the picture is divided into two parts such that $\mathbf{P} = \langle p_1, p_2 \rangle$ and the arrangement **m** maps these parts to their original positions within P. An alternative arrangement $\mathbf{m}^*$ maps the same parts to different positions, at right.

Having defined the relevant notions of a picture's parts and mode of composition, we now define the *content* of a part. Here we encounter a decision point: the content of a picture is defined relative to a projective index, which effectively fixes the vantage point of that picture. Is the content of picture's part defined relative to the same vantage point, or is it effectively "cut free" from the evaluative context of the original? Both seem to me reasonable ways of thinking about the content of a picture's parts. In what follows, I shall mainly elaborate the first approach, assuming that the projective index of the whole constrains the projective indices relative to which the the content of the parts are determined. Afterwards I will show how to execute a compositional analysis using the latter definition.

This tack engenders some technical complications. Recall that we defined the content of a picture relative to a projective index $j = \langle v, r \rangle$ where v is a viewpoint function and r is the picture position function. Given a scene S, r_S maps pictures in 2-space to flat regions in the 3-space of S. But we can't apply such a function straightforwardly to members of **P** in a given decomposition, for two reasons. First, there is a size mismatch: r_S takes whole pictures as inputs; the members of **P** are parts of pictures. The solution here is simply to let r_S range over regions of P, where the outputs are corresponding regions of $r_S(P)$.

A second challenge is more fundamental. The members of $\mathbf{P}$ are independent spaces, not regions of P; they contain no intrinsic information about their original relative position in P. Where in the scene then, should r_S embed such regions? How do we guarantee that they won't get misplaced in the process? The solution here is to first apply the embeddings in $\mathbf{m}$ to the members of $\mathbf{P}$. Each such function cannily locates a member of $\mathbf{P}$ within a unified picture space. At that point r_S can be brought in to map subregions of this space to their corresponding regions in the space of S. The diagram below illustrates this technique for a simple example where $\mathbf{P} = \langle p_1, p_2 \rangle$.

Given a member m_n of $\mathbf{m}$ and a picture position function r_S, we can define a function $r_S m_n$ which maps p_n into the space of S. (Correspondingly, rm_n takes a scene S and yields $r_S m_n$.) In this way we can say that the content of p_n is defined relative a projective index $\langle v, rm_n \rangle$. In general, where $j = \langle v, r \rangle$ is the projective index for a picture P, let $j_n = \langle v, rm_n \rangle$ be the projective index for any part p_n of P relative to a decomposition D. Thus we have arrived a definition of the content of a part of a picture.

In summary: we began with the idea of the content of a picture P in a system I at index j. We then introduced the concept of a decomposition D of P, which consists of an arbitrary division of the picture into parts, and a record of how these parts are combined in the original. Then, for every part p_n of P relative to the decomposition, we showed how to define the content p_n in I at index j_n, where j_n is determined by the decomposition D and the original index j. Thus given a picture, a system, an index, and a decomposition, we gave a general definition of the content of each part. To show that a pictorial semantics is compositional, we must further demonstrate that the content of the whole can be derived from the content of the parts and the way they are combined, as reflected in the decomposition. As we proceed, it will be convenient to refer to the content of a picture P in system L at an index j as "$[\![P]\!]_{L,j}$", and the content of a picture part p_n relative to decomposition D as

"$[\![p_n]\!]_{L,j_n}$". What we aim to show is that, for any picture P in system L, index j, and decomposition D of P, we can derive $[\![P]\!]_{L,j}$ from $[\![p_n]\!]_{L,j_n}$ for every p_n in **P** and the mode of composition reflected in the arrangement **m**.

To motivate our final analysis, let us see how a simple first suggestion fails. According to this proposal, the content of a whole picture is the generalized *union* of the content of its parts. That is, given the system L, for any picture P in I and decomposition D:

(34)　$[\![P]\!]_{L,j} = \bigcup\{[\![p_n]\!]_{L,j_n} \mid p_n \text{ in } \mathbf{P}\}$

Consider the predictions of this analysis for the case illustrated below. Here a picture P is decomposed into two parts p_1 and p_2. Image p_1 accurately depicts scene A and scene B, but not scene C. Image p_2 accurately depicts scene B and scene C, but not scene A. If we take the union of the sets of scenes that both p_1 and p_2 accurately depict, we get a set that includes all of A, B, and C. But obviously only B is in the content of P; P does not accurately depict A and C. Evidently, the content of P cannot be so defined; so (34) is false.

The failure of this method clearly suggests a better one. The content of P in the example above includes A, the scene accurately depicted by both p_1 and p_2, but none of the scenes accurately depicted by p_1 alone or p_2 alone. Thus it seems we should define the content of the whole, not as the generalized union of the content of its parts, but as their generalized *intersection*:

(35)　$[\![P]\!]_{L,j} = \bigcap\{[\![p_n]\!]_{L,j_n} \mid p_n \text{ in } \mathbf{P}\}$

This analysis is evident enough, but here is an informal proof, with indices freely suppressed for brevity. Let D be a decomposition of a picture P into two parts, p_1 and p_2 (the number of parts

is chosen arbitrarily for the example). Then proposition (35), applied to P is equivalent to the following claim: for any scene S, P accurately depicts S if and only if p_1 accurately depicts S and p_2 accurately depicts S. By the Geometrical Projection Theory, this in turn is equivalent to the proposition that P is a perspective projection of S if and only if p_1 is a perspective projection of S and p_2 is a perspective projection of S (here I suppress mention of the flexibility function). But this proposition follows directly. From left to right: if a picture can be derived from a scene by projection, then each of its parts can be so derived as well. From right to left: if all of a pictures parts can be individually derived from a scene when located at their original position in the picture, then the picture itself can also be derived from the scene. Thus (35) correctly defines the content of the picture in terms of the content of its parts.

But is the proposed semantics *compositional*? Though it is clear how the content of the picture is derived from the content of its parts, one might wonder what happened to the "mode of composition" in this compositional semantics. In fact, it is encoded in the definition of the index j_n in terms of original index j and the arrangement **m**. An unfortunate consequence of this approach that the conclusions derived above appear trivial; unlike the sophisticated combinatorics of language, the compositional rule for pictures amounts to little more than simple intersection. So it may prove more intellectually satisfying to see how compositionality can also be derived when the mode of composition is not built into the definition of part-content, but instead imposes a substantive constraint on how the content of the whole is derived from the content of its parts.

To this end, we will replace the definition of content at work in the preceding discussion, according to which the content of a picture is defined relative to a projective index, with the alternative conception of picture content discussed in footnote 35 which abstracts away from projective index, as follows:

(36) $[\![P]\!]_L = \{\langle S, j\rangle \mid P \text{ accurately depicts}_L S \text{ at } j\}$

Here the content of a picture is modeled as the set of all scene-index pairs such that the picture accurately depicts the scene at that index. The definition applies to the content of parts of pictures, without modification. The result is that the content of a picture part is "cut free" from any projective index that might be used to evaluate the whole. As a consequence, picture parts do not encode their mode of composition in their content, and so there is no very simple technique for deriving the content of a whole picture from a set containing the content of each of its parts. As in the familiar linguistic case, merely intersecting the relevant contents is too unconstrained, leaving compositional semantics for pictures a substantive matter.

These observations are illustrated by the following case. In the diagram below, the scene D is incompatible with the content of the picture P, but compatible with each of its parts p_1 and p_2 considered individually— even when the viewpoint and orientation of the picture plane is held fixed across cases.

An unsophisticated method of intersection would incorrectly allow D (along with the common viewpoint) into the content of P, because it is in the content of both parts considered individually. What goes wrong with such a method? A clue is that the only conditions in which p_1 and p_2 can be projected from D at the same viewpoint are those in which they arranged relative to one another in a way that *differs* from their original arrangement in P. The diagram below makes this fact clear.

While the scenes in the content of the whole will indeed be contained in the intersection of the content of the parts, this case illustrates that such intersection is too permissive. Instead, we must specify a spatial arrangement of the contents so as to reflect the spatial arrangement of the parts in the original picture. This can be achieved by defining the content of the whole not directly in terms of the contents of its parts, but in terms of the content of its parts when their projective indices are suitably constrained to conform with their position in the picture's compositional structure. Thus for any system I, picture P and decomposition D of P:

(37) $[\![P]\!]_L = \{\langle S, \langle v, r \rangle \rangle \mid \forall p_n \text{ in } \mathbf{P} : \langle S, \langle v, rm_n \rangle \rangle \in [\![p_n]\!]_L \}$

That is: the content of a whole picture is the set of all scene-index pairs such that each pair is in the content of *every* part of the picture, when the index is suitably constrained to reflect the position

of the part within the whole picture's compositional structure. Now more clearly than before, the content of the whole is determined by the contents of the parts ($[\![p_n]\!]_L$ for each p_n) and the way they are put together (by $\mathbf{m}$). Thus the proposed semantics for the system of perspective line drawing is fully compositional.

Chapter 4

Iconic and Symbolic Semantics

Since the pioneering work of Charles Sanders Peirce, theorists of signs have distinguished two basic forms of representation. ICONIC representation is exemplified by pictures, maps, sculptures, videos, and audio recordings. SYMBOLIC representation is exemplified by words and sentences, but also maritime flags, smoke signals, and sequences of lanterns used as coded warnings at night. As a very rough heuristic: symbolic representations are those which bear a merely arbitrarily relation to their content, while iconic representations are those which bear a more direct relationship to their content, one often glossed as "resemblance." Thus the relationship between the English word 'tree' and the property of being a tree is utterly arbitrary; but the relationship between an accurate depiction of a tree and its subject is much more immediate and rich. This chapter is an investigation of this distinction.

The examples cited above are case of symbolic and iconic representation in public communication, but the distinction also arises in the more fundamental domain of mental representation. Some form of iconic representation seems to be operative in low-level perception. For example, early regions of the visual cortex have a retinotopic organization: the spatial distribution of neural activity literally corresponds to the pattern of light reflected to the eye. (Pylyshyn 2006) Still, the exact role of such apparently iconic representations is not well understood. Symbolic representation, on the other hand, has secured a foundational place in the contemporary understanding of human cognition and communication. The computational theory of mind— the account of the mind presupposed by contemporary cognitive science— characterizes higher-order mental representations as mental symbols, and defines processes of thought as transformations of such symbols.

Let us term the feature of a public or private sign which determines whether it is iconic, symbolic, or the like, its REPRESENTATIONAL MODE. As I noted in the introductory chapter of this dissertation, it seems likely that there is a full spectrum of representational modes that lie between those which are entirely symbolic and those which are entirely iconic. There may be other dimen-

sions of representational mode as well. Peirce, for example, believed that INDEXICAL representation was third, distinct such category. But my focus in this chapter is on the apparently distal cases of iconic and symbolic representation; explanation must begin somewhere. And of course, the application of terms such as "symbolic" and "iconic" is pre-theoretical. A theory of representational mode should help to illuminate these untutored distinctions, but it may ultimately revise them as precision and generality require.

By now, we have encountered concrete and examples of semantics for both iconic and symbolic representational systems. Semantic theories of natural language, familiar from contemporary linguistic research, describe one family of symbolic systems in detail. The semantic theory of pictorial representation, developed over the course of the last three chapters, describes one family of iconic system. Though I have stressed their common foundations, the semantics of language and the semantics of pictures are undeniably different. There seem to be too many differences to profitably enumerate. Still, we want to know, what are the fundamental and distinguishing differences which set these forms of representation apart? And what explains their commonalities with other systems of the same representational mode? Answering these questions is the task for a theory of representational mode.

Until such a theory is secured, we lack the necessary means to adjudicate debates about whether this or that representation is symbolic or iconic. Even more pressing, we are unable to give precise substance to the theory that the mind is symbol-manipulating computational device, or that language users communicate by means of coordination on public symbols. Ultimately, a theory of representational mode should be more than diagnostic; it should help explain the empirical distribution of representational modes in human communication. Why have both symbolic and iconic representation thrived in contemporary media? What special communicative needs do each serve? Why does so much of cognition and communication depend on symbolic, as opposed to iconic representation— if indeed this is true? Despite their evident importance, the attempts to answer these questions have been relatively few.

This chapter sketches the foundations of new account of representational mode. I will ultimately argue that symbolic and iconic representation are distinguished by the architectures of the systems in which they are physically instantiated. Symbolic systems, I suggests, exploit regularities which associate *particular* signs structures with content, whereas iconic systems exploit *general* regularities in the associations between sign and content. This conclusion substantiates the view that symbolic and iconic representation constitute different representational strategies, by which an agent may encode content in a physical sign.

Section §1 reviews theories in the formalist tradition, according to which representational mode is merely a matter of the structural features of signs and their contents. §2 addresses theories in the resemblance-theoretic tradition, which attempt to analyze representational mode in terms their relative dependence on relations of resemblance. §3 investigates Saussure's promising suggestion that symbolic, but not iconic representation, is based on arbitrary links between sign and content. §4 sketches an alternative, computational approach to the theory of representational mode.

1   The formalist tradition

Theories in the FORMALIST tradition treat representational mode merely as a matter of formal structure. After considering and then rejecting some of the most prominent examples from this tradition, I'll define this family of theories more precisely, and argue that none can succeed.

To the naked eye, the most obvious difference between typical symbols, such as words, and typical icons, such as pictures, is how they *look*. This suggests that icons and symbols differ essentially in their physical instantiation. But this suggestion is easily refuted, as Fodor (1975, pp. 178-9) as noted: we could easily construct a novel lexicon by borrowing a collection of pictures from some system of depiction, and arbitrarily assigning these pictures to our choice of denotations. The result would be a system of representation which, for all intents and purposes, was symbolic, though its signs had the same physical instantiation as those of some system of depiction. Thus the intrinsic physical characteristics of a sign is at best a poor indicator of its representational mode. (The same argument can be recapitulated for mental representations: if there are iconic mental representations, then duplicates of the neural realization of these states could also realize symbolic representations in some *other* cognitive system.)

A more refined approach to the same intuitive idea, distinguishes representational mode by the kinds of *formal structures* which are provided to interpretation. Formal structures here are abstractions from physical structure; they reflect all and only those features which are salient for interpretation. For example, in a typical semantic theory for a natural language, the inputs to interpretation are PARSE TREES, complex objects whose structure corresponds to the underlying syntactic structure of each disambiguation of a linguistic string. By contrast, in the pictorial semantics articulated in Chapter 3 took colored two-dimensional planes as inputs. Even if this particular proposal is ultimately revised, something quite like this, and quite different from linguistic parse trees, surely serve as the input to pictorial interpretation.

Based on these kinds of observations it becomes tempting to distinguish representational modes

on the basis of formal structure. Symbolic signs, perhaps, have *logical form* or *syntactic structure*, while iconic signs do not. But what is the relevant notion of logical form or syntactic structure? By way of elaboration, one approach holds that symbolic signs are uniquely those with finitely many atomic parts, while iconic signs have uncountably many parts. But this suggestion immediately runs afoul of digital imagery: pictorial representations with finitely many atomic parts. Fodor (2007, p. 108) has a more sophisticated suggestion: symbolic signs are uniquely those with a "canonical decomposition". A canonical decomposition is the *unique* structure associated with any sign for the purposes of interpretation. As Fodor observes, linguistic expressions have parse trees as canonical decompositions. But there is not uniquely correct way to parse a picture; the compositional semantics offered in Chapter 3 gave formal ratification to this claim. Yet, despite these success, Fodor's proposal faces basic problems. Again, there is the issue of digital imagery. It is true that images in such systems do not necessarily have a unique canonical decompositions, but it is worrying that it would be easy to stipulate such a decomposition, in which the basic inputs to interpretation are the pixels, together with their spatial arrangement. A second problem is that the claim that linguistic structures have canonical decompositions is itself an contentious empirical hypothesis. There are viable theories of linguistic interpretation which reject the principle. (Steedman 1996; Steedman and Baldridge 2009)

More generally, *all* attempt to representational mode in terms of the formal structure of the representation face the problem that distinctions in mode arise even among compositionally *simple* signs. We can imagine codes which map modulations in the brightness of a single lamp to some gradable quantity, such as the amount of water in a reservoir. Some such codes are distinctly iconic— brightness is mapped to water level according to transformational code reminiscent of pictorial or diagrammatic systems of representation. Other codes are distinctly symbolic— brightness levels are arbitrary mapped to different quantities of water. These codes are described in much more detail later in this chapter; suffice for now to observe that they seem to differ with respect to representational mode, but their signs are exactly alike with respect to formal structure. They are all unitary lamp states at different levels of brightness. Thus attempt to define representational mode in terms of the formal structure of signs has little promise.

This conclusion illustrates an important lesson: the distinction between symbolic and non-symbolic representation is not a distinction among *signs* considered in isolation. It is at least a distinction among signs evaluated *relative* to systems of representation. To simplify matters, we may stipulate that an entire system is symbolic just in case every sign in that system is symbolic with respect to it. Then the mandate of the present inquiry is to identify criteria by classify *systems*

as symbolic or not-symbolic. (On this way of proceeding, if there are systems which combine symbolic and non-symbolic elements, those systems will be counted as non-symbolic, but this is acceptable.)

These kinds of consideration give rise to a more popular and sophisticated class of theories which analyze symbolic representation in terms of the structure of the *domain* of signs in a system, rather than the structure of individual signs. Holistic accounts like these typically seek to reduce the distinction between symbolic and iconic representation to the distinction between digital and analog signals, or close counterparts. The simplest such approach defines symbolic systems as those with only a finite or countable domain of signs, while non-symbolic systems have continuous domains. Yet all such accounts are flummoxed by the existence of digital imagery: genuine systems of depiction with finite domains of images, each composed of finitely many atomic parts. Since there is a functional mapping from images to contents, such systems also have only finitely many contents. So distinguishing symbolic systems by the cardinality of their range won't work either.[1] Thus the advent of digital imagery shows that the suggestive connections between the symbolic and the digital are misleading. Analog and digital are categories of signal; symbolic and iconic are categories of representational system.

All of the accounts of symbolic representation discussed thus far belong to an theoretical tradition which I will call FORMALIST.[2] Theories in the formalist tradition treat systems of representation abstractly, as interpretation functions defined set theoretically— sets of ⟨*sign, content*⟩ pairs. The actual mechanism of interpretation remains a black box. Formalist theories contend that a functional description is sufficient to determine the mode of representation of a given system. Theories which focus especially on the sign domain or content domain, like those discussed above, are species of formalism; and we have seen how these fail. But perhaps there is some more sophisticated way of describing the functional *connections* between the sign domain and the content domain which will successfully capture the concept of symbolic representation? Indeed, this was the strategy elected by none other than Nelson Goodman (1968). Yet I now argue that *no* purely formalist account can successfully define symbolic representation.

Consider a grid of a fixed and finite size. Let S be the set of every possible way of filling this grid with black and white squares. Since the grid is finite, and there are only two ways of filling in each square on the grid, S itself is finite. Now let D be a system of digital black and white drawing,

---

[1]Haugeland (1991) is one of the few authors to analyze representational mode in terms of types of content alone.

[2]The nice distinction between theories of representational mode that focus on the sign domain, the content domain, or the connection between them is due to Kulvicki (2006*b*). Kulvicki's discussion there has strongly influenced the direction of the present discussion.

whose domain is S. For every image in this domain, D assigns it a content (modeled as a set of of ⟨ *scene, index*⟩ pairs) according to some method of projection. Let M be the range of D; since D has finitely many inputs, and is deterministic, M is also finite.

Now suppose that everyday we develop a new code that randomly assigns a member of S to a member of M. Such codes are symbolic: they depend on a library of arbitrary associations that match signs with their content, in the same manner of linguistic lexicons. If we go on long enough (at most, $|S|!$ days), eventually, on day n we will hit upon a code whose inputs and outputs are paired exactly as they are by D. Call this *Code n*. Even though Code n matches D, a pictorial system, in terms of inputs and outputs, it associates its signs with contents by an utterly arbitrary mechanism quite unlike that of D. Interpretation in Code n, like that of the codes created just before and after it, requires that we consult a library of sign-content associations, not that we apply some method of geometrical projection. Indeed, there is no principled distinction between Code n and its other, obviously symbolic sequence-mates. We should therefore conclude that Code n is also a system of symbolic representation.

The case involves two systems of representation— D and Code n— which are identical with respect to input and outputs, but differ with respect to mode of representation; one is iconic, the other symbolic. It follows immediately that the defining feature of symbolic representation cannot be described in terms of sets of input and output pairs. No formalist theory of symbolic representation can succeed. (Some may object that, despite their duplicate physical instantiation, expressions of Code n are distinct from those of D. But the same argument can be rerun using the system of structurally simple iconic representations described above, as the reader may verify for herself.)

The failure of formalism reveals a principle which will guide the remainder of my investigation. The defining feature of symbolic representation is not to be found in the nature or structure of the signs that a system takes as inputs, or the contents it delivers as outputs. Rather, it must be something about the *relation between* signs and content determined by that system which defines symbolic representation. Symbolic and iconic representation, we may conclude, correspond to different kinds of sign-content relations; that is, they correspond to different kinds of representation relations. This is a conclusion we should have anticipated. The kind of semantics for pictures described in Chapter 3 differs fundamentally from the more familiar semantics of natural languages; that difference is not merely a matter of the internal structure of the inputs to interpretation. There is an obvious difference in the *mechanics* of the semantic analysis in each case. But it remains to be seen what this difference is.

2   The resemblance tradition

One prominent family of views attacks the problem of representational mode by focussing on iconic representation. Reflecting the considerations of the last section, this tradition holds that iconic representation is distinguished by the relation it imposes between iconic sign and content. Indeed, the iconic connection between a picture of a tree and the tree itself is undeniably intimate and direct when contrasted with the very thin symbolic connection between the *word* 'tree' and the property of being a tree. But how to define this intimate connection? A seemingly obvious answer holds that the connection is one of resemblance, understood as sharing of properties. This would explain the "intimacy" and "directness" of iconic representation. Symbolic representation is then defined, contrastively, as representation which is not based in resemblance.

A first worry with this account is that it is not clear how it overcomes the problems facing theories in the formalist tradition. The resemblance requirement on iconic representation cannot be a merely extensional; it must be somehow reflected in conditions that extend beyond the intrinsic features of the sign and content so related. This is a challenge, not a fatal problem. Perhaps, as some theorists of depiction have recently suggested, iconic representation depends on an agent's *intending* to bring about a relation of resemblance between sign and content. (Abell 2009; Blumson 2009*a*) This strategy cannot be carried over to the domain of mental representation, though it might be held that it is the cognitive *function* of some representational system to bring about resemblance between sign and content. Obviously, more must be said to avert the anti-formalist worries.

In any case, fundamental problems render the entire approach unfeasible. As I have argued at length in Chapter 2, resemblance cannot ground pictorial representation. There are forms of pictorial representation, thus forms of iconic representation, which cannot be analyzed in terms of resemblance alone. And there are no obvious species of resemblance which can even be claimed as uniquely necessary to pictorial representation. Consequently, symbolic representation cannot be defined in terms of the absence of resemblance either. These conclusions only complicate matters. The most straightforward criterion by which symbolic and iconic representation might have been differentiated cannot succeed. We must search elsewhere.

3   The Saussurian tradition

A progenitor of modern linguistics, Ferdinand de Saussure famously proposed that the distinctive feature of linguistic representation is the *arbitrary* relation between sign and content. He contrasted this with systems of representation, such as mime, which determine a "natural" relation

between sign and content. Following is a canonical passage from Saussure's *Course in General Linguistics*. (1922/1983, I.2) Note that Saussure's terminology here is idiosyncratic: he calls the inputs to interpretation "signals"— corresponding approximately to what I have called "signs"; the outputs of interpretation are "significations"— corresponding to my "contents"; finally, he reserves the term "sign" for a complex whole that includes the signal, the signification, and the representational *relation* between them.

> T he link between signal and signification is arbitrary. Since we are treating a sign as the combination in which a signal is associated with a signification, we can express this more simply as: *the linguistic sign is arbitrary*.
>   There is no internal connexion, for example, between the idea 'sister' and the French sequence of sounds *s-ö-r* which acts as its signal. The same idea might as well be represented by any other sequence of sounds. This is demonstrated by differences between languages, and even by the existence of different languages. The signification 'ox' has as its signal *b-ö-f* on one side of the frontier, but *o-k-s* (*Ochs*) on the other side.

Though Saussure uses his language as his primary example, it is clear that he intends the analysis in terms of arbitrariness to extend to all symbolic representations. Indeed, I will construe Saussure's proposal as a general theory of representational mode, albeit one which cleaves the space of modes into those which are symbolic and those which are not.

As a theory of representational mode, Saussure's proposal succeeds exactly where theories in the formalist tradition fail. He holds that that the distinctive feature of symbolic representation is located in the *relation between* sign and content. And it is natural to understand arbitrariness as a property of the causal or historical connections which link physically instantiated signs to contents. Representational mode therefore depends on factors which extend beyond the mere structural characteristics of signs and contents alone. Furthermore, Saussure's view is powerfully intuitive. The relationship between, for example, a drawing of a cat and the cat itself is in some way obviously motivated or natural; the relation between the *word* "cat" and the biological family of cats (or any member of that family) is completely arbitrary. Another term could have served to denote that category just as well.

I will argue that Saussure's theory is correct but incomplete. In so far as we have a pretheoretical grasp of the notion of arbitrary representation, those cases of representation which seem to be arbitrary are just those which seem to be symbolic. Yet the account is incomplete, for Saussure supplies no explicit account of *what it is* for such a relation to be arbitrary, nor of the nature of those relations which are not arbitrary. The danger is that Saussure's account replaces the pre-theoretical notion of symbolic representation with another, equally opaque concept. Nor is it trivial to fill the lacuna. To dramatize the difficulty of this task, we will consider two possible elaborations of the

arbitrariness analysis. Each can be detected in Saussure's own text, but both are ubiquitous in discussions of the topic. Each pursues the hypothesis that arbitrariness is a naturalistic feature of the sign-content relation, to be understood in terms of its profile of causal or historical connections to the environment.

The Conventionality Analysis. Perhaps the most popular interpretation of of Saussure's proposal takes arbitrariness and conventionality to be equivalent.[3] According to this analysis, the relationship between a sign and its content is symbolic if and only if the sign is mapped to the content by an interpretive social convention. Yet despite its widespread appeal, this condition is neither necessary nor sufficient for symbolic representation.

Fodor (1975, p. 178) was first to observe that conventionality is unnecessary for symbolic representation: "English would be a discursive (i.e., a symbolic; i.e., a noniconic) representational system even if it were innate (i.e. nonconventional)." Indeed, symbolic, non-conventional representation is regularly deployed by non-human animals. For example, fireflies of different species employ distinct flash patterns to indicate species identity, thereby facilitating coordination with potential mates. The use of a given flash pattern by a species is symbolic. It is an arbitrary representation of membership in a given species; another flash pattern would have worked just as well, so long as it wasn't in use by another species. But it is also genetically determined, therefore non-conventional. (Lewis and Cratsley 2008) Hence symbolic, arbitrary representation can exist without social convention.

The point can be extended by to the human mind (at least, according to the computational theory of mind). Mental symbols, like their public counterparts are thought to bear a merely arbitrary relation to their contents. To make the point vivid, consider an alien mind very much like a computer; in this mind the computational symbol "101" represents the property of being a cat. But "011" would have served the creatures purpose just as well. The mental symbols are arbitrary. Yet the representational relation between a mental symbol and its content is not in general determined by social convention. Instead it is a product of biological evolution and the causal history of the organism which is the bearer of that symbol.

The insufficiency of the conventionality analysis was observed by Eco (1979), who noted that a code may be selected for use by culture even when it is fully "motivated". Indeed, in Chapter 1, I argued at length that there are systems of depiction which are established by social convention. Thus there are systems of representation which are definitionally *non*-symbolic, yet conven-

---

[3]Saussure himself appears to have shared this conviction: "The main object of study in semiology will none the less be the class of systems based upon the arbitrary nature of the sign. For any means of expression accepted in a society rests in principle upon a collective habit, or on convention, which comes to the same thing." (Saussure 1922/1983, I.2)

tional. As I stressed in that discussion, we should not confuse the concept of arbitrariness associated with the open-ended selection of an interpretive rule by a population— a prerequisite for social convention— with the concept of arbitrariness used to describe the internal structure of a representaitonal relation— the defining characteristic symbolic representation. The two concepts are essentially independent.

The Counterfactual Analysis. Instead of attempting to specify specific social relations underlying symbolic representation, an alternative is to look to the causal networks in which physically instantiated symbols are embedded. These in turn are spelled out in terms of counterfactual conditions on representation: the relationship between a sign and its content is symbolic if and only if users of the sign *could* have employed some *other* sign to designate the same content. This suggestion captures our intuitive idea of arbitrariness rather well; Saussure himself relied on such counterfactuals to elucidate the notion in his original discussion. But taken as a serious analysis of symbolic representation, it is unclear how to make this proposal precise while maintaining its plausibility.

One one hand, if the counterfactual modal is read narrowly, then the condition appears not be necessary. For example, if we allow that a representation is symbolic just in case its bearer could have *chosen* to use a different sign for the same content, or could have taken some course of action that would have lead to the use of a different sign, then we run afoul of biologically determined symbolic representation. The bearer of a mental symbol, for example, typically has no choice about what symbol she employs to represent a given content, nor, typically, is there any discreet alternative action which could have lead to that result. Indeed, in order to coherently specify a situation in which the bearer of a mental symbol used some other symbol to represent the same content, we may have to re-imagine the agent's entire lifespan; for biologically basic symbolic representation, such revision may engulf entire branches of the evolutionary tree.

On the other hand, if the counterfactual modal is read broadly, then the condition appears not to be sufficient. For example, suppose we allow that representation is symbolic just in case, starting arbitrarily far back in history, a situation could have evolved in which the agent used a different symbol to represent the same content. Then the condition appears to be necessary for symbolic representation, but it also includes a variety of non-symbolic representational forms. A given population may rely on a system of depiction which maps a picture P to a certain content C; if history had evolved differently, the same population might have used a different system of depiction, relative to which a different picture P' depicted C. Such a counterfactual would satisfy the proposed condition, but the representation in question would not be symbolic. As stated, the counterfactual analysis fails to provide a productive solution to the problem at hand.

Neither precisificaiton works, and it is not clear how better to proceed. Without an explicit theory of arbitrariness, Saussure's project remains incompletely realized. In the next section I aim to provide a theory of symbolic representation which is broadly Saussurian, but gives substance to the motivating idea of arbitrary representation.

4   A computational approach

In this section, I sketch an alternative theory of representational mode. By dint of the nature of this chapter's subject matter, and the time I have left myself to write it, what follows is necessarily speculative. I ask for the reader's patience.

The moral of the discussion of formalism was that differences in representational mode must be traced to the variety of representation relations. Saussure offered such an account, cast in terms of representational arbitrariness; this was promising in part because it suggested an avenue by which representational mode might be grounded in the physical conditions that give rise to representation. But as we saw, this idea could not be straightforwardly elucidated in terms of social convention, nor causal etiology.

The computational alternative I'll develop here holds instead that systems of representation belonging to different modes are distinguished by the kinds of algorithms that compute them. The view is most clearly articulated with respect to public systems of representation. I propose that public systems of representation should be understood as interpretive algorithms that map signs to content.[4] There are many kinds of algorithm, corresponding to variety of representational systems. The various algorithmic kinds are distinguished principally by their architecture— the structure of rules they enlist to compute a content, given a sign. According to the computational approach to representational mode advocated here, iconic and symbolic representation are reflections of fundamentally different algorithmic architectures. The algorithms grounding symbolic representation exploit *individuated* rules which map signs to content in a piecemeal fashion, by applying a distinct procedure to each sign. The algorithms grounding iconic representation exploit *uniform* rules which map signs to content by a applying the same procedure to each sign. On this view, symbolic and iconic representation are distinguished by different kinds of computational architecture. They reflect different strategies for storing and retrieving content from signs.

This analysis presupposes that systems of representation can be identified with interpretive *algorithms*, which I understand, following Fodor (1975), as systematic mappings from public signs into signs of a mental code. The computations involved in such a mapping take public repre-

---

[4]Thanks to Matthew Stone for helping to clarify this basic orientation.

sentations as inputs and deliver mental representations as outputs. Symbolic representation is based on one family of such computations, while non-symbolic, including iconic representation is based on another. Yet ultimately, the theory must be extended to the mental representations themselves. Here my discussion will be especially imprecise. Quite roughly, I propose that the analysis of representational mode cast in terms of differences of algorithmic structure can be extended by generalizing the idea of an algorithm. Algorithms are one among a family of mechanisms that deterministically map inputs to outputs. Others include physical laws, chains of causal interaction, and teleological norms. Letting the blanket term *rule* cover all these phenomena, I propose that the same architectural distinctions arise among rules. And these in turn ground the differences in mode of mental representations.

To simplify the ensuring discussion, I will focus almost entirely on compositionally simple signs. In the final section, I will indicate how the proposed analysis might be extended to complex signs as well.

4.1   Simple iconic and symbolic representation

As advertised, I will initially confine my discussion to SIMPLE signs. A sign is simple in a system of depiction if its content is *not* computed as a function of the content of its parts. Since simple signs with full-blown representational content are unusual in any mode, I pause in this section to provide detailed examples of simple iconic and symbolic representational systems. They will serve as points of reference in the remainder of the chapter.

Suppose I must communicate to you information about the the amount of water in a reservoir, over a great enough distance that merely shouting will be ineffective. I have a lamp which can be set at ten evenly distributed levels of brightness. We agree to use different settings of the lamp to represent the percentage of the reservoir that has been filled with water. Of course, there are many different kinds of codes we might establish to facilitate this communication. Here is one. According to the *I-Code* ("I" for *iconic*), the brightest setting of the lamp indicates that the reservoir is between 100% and 90% full; the next lowest setting indicates that the reservoir is between 90% and 80% full, and so on, down to the dimmest setting, which corresponds to the reservoir being between 10% and 0% full. To express the semantics for such a system, let level() be a function which maps worlds to the percentage of the reservoir which is filled in that world, rounded up to the nearest 10-percentile. Let bright() be a function from states of the lamp to a numerical measure of the brightness of the lamp in that state, ranging from 1 to 10. Then the semantics for the I-Code is as follows. For any state of the lamp *S*:

(1) $I(S) = \{w \mid \text{level}(w) = \text{bright}(S) * 10\}$

That is, the content of a given setting of the lamp is the set of worlds in which the percentage of the reservoir which is full is at most 10 percentage points lower than the measure of the brightness of the lamp multiplied by 10.

A very different code would exploit randomly stipulated correspondences between levels of brightness and ranges of fullness. Call one such system the *S-Code* ("S" for *symbolic*). Here is its semantics, for any state of the lamp *S*:

(2) $I(S) = \{w \mid \text{level}(w) = 10\}$ if $\text{bright}(S) = 7$
$\{w \mid \text{level}(w) = 20\}$ if $\text{bright}(S) = 3$
$\{w \mid \text{level}(w) = 30\}$ if $\text{bright}(S) = 6$
$\{w \mid \text{level}(w) = 40\}$ if $\text{bright}(S) = 1$
$\{w \mid \text{level}(w) = 50\}$ if $\text{bright}(S) = 5$
$\{w \mid \text{level}(w) = 60\}$ if $\text{bright}(S) = 8$
$\{w \mid \text{level}(w) = 70\}$ if $\text{bright}(S) = 4$
$\{w \mid \text{level}(w) = 80\}$ if $\text{bright}(S) = 9$
$\{w \mid \text{level}(w) = 90\}$ if $\text{bright}(S) = 2$
$\{w \mid \text{level}(w) = 100\}$ if $\text{bright}(S) = 10$

The S-Code corresponds to a system of representation which appears to be straightforwardly symbolic. The mappings from signs to content in that system are arbitrary in exactly the way Saussure associated with symbols. By contrast, the I-Code is highly iconic. The brightness of the light represents the state of the reservoir in a manner which is at least diagrammatic, if not simply pictorial. (Nautical maps often associate variations in color lightness with water depth in essentially the same way.) All signs in either system are *simple* because in neither case do the signs have component parts which are assigned content by the system of representation in question. This is not deny that they have structure; signs must be differentiated on the basis of *something* after all. But the component parts of this structure are not assigned content in the calculation of the content of the whole.

When we go to actually use the I-Code or the S-Code, the receiver of the message must interpret it. It will be useful to establish a reasonably explicit representation of the procedure an agent might use when interpreting a message in each of these codes. This would be an algorithm which computes the interpretation function of each, by taking signs as inputs, and returning content as

output.[5] Again, I assume a framework in which the process of interpretation for public signs is one of systematically mapping public signs into signs of a mental code. The computations involved in such a mapping take public representations as inputs and deliver mental representations as outputs. These computations map signs to *content* only in the derivative sense that the representations which are outputs have an explicit and immediate connection to the environment, and therefore real representational content. (This is the same derivative sense in which an algorithm which computes the successor function delivers *numbers* as outputs.)

Here I will describe algorithms that take signs as inputs and output *sets of worlds*. Strictly speaking, these algorithms should be understood as outputting data structures which contain unordered *lists* (i.e. sets) of *names* of worlds, and assume that that, for the purposes of computation, there are only finitely many worlds. I take this to be an especially explicit and reasonably meta-language-neutral representation of content. Those who wish to imagine alternative computational architectures whose native tongue does not include names of worlds should have no trouble porting the examples provided here into their preferred framework.

```
I-Code
READ sign
WRITE { w | level(w) = bright(sign)*10 }
```

```
S-Code
READ sign
IF bright(sign) = 7
    WRITE { w | level(w) = 10 }
ELSE IF bright(sign) = 3
    WRITE { w | level(w) = 20 }
ELSE IF bright(sign) = 6
    WRITE { w | level(w) = 30 }
ELSE IF bright(sign) = 1
    WRITE { w | level(w) = 40 }
ELSE IF bright(sign) = 5
    WRITE { w | level(w) = 50 }
ELSE IF bright(sign) = 8
    WRITE { w | level(w) = 60 }
ELSE IF bright(sign) = 4
    WRITE { w | level(w) = 70 }
ELSE IF bright(sign) = 9
    WRITE { w | level(w) = 80 }
ELSE IF bright(sign) = 2
    WRITE { w | level(w) = 90 }
ELSE IF bright(sign) = 10
    WRITE { w | level(w) = 100 }
```

In both cases, the READ operation retrieves the value of a variable input; the WRITE operation outputs a new value.

So much for public representation. Could there be simple iconic and symbolic mental repre-

---

[5]I shall use the terms "algorithm" and "procedure" more or less interchangeably; I take these terms to denote what theorists of computation sometimes call an "effective procedure"— a step-wise process so simple and deterministic that it can be carried out by a mechanical device.

sentations as well? I assume that the idea of a simple symbolic mental representation is familiar enough. Typically, simple symbolic mental representations in humans are thought of, in analogy with the words of public language as having denotational content, rather than fully truth-conditional content. But it is not difficult to imagine (perhaps much simpler) creatures that would. The same considerations apply to simple iconic representation. I can see no in-principle reason why intelligent agents could not have iconic mental representations. A mental code much like the I-Code, might use intensity of chemical or electrical signals to carry content about some internal or external state. Unfortunately, there is no consensus about the structure of the causal network which relates mental representations to their content, so I cannot express this relation as simple algorithm, as I did for the interpretive rules of public language. At most, we can allow that the mapping of mental signs to content is governed by deterministic regularities, though the nature and structure of these regularities remains obscure.

4.2   Individuated and uniform representation

I now turn to the analysis of representational mode, again confining my discussion to simple representations. The heart of the analysis is a distinction between two kinds of rule. Initially I'll draw this distinction in a quick and informal way; later I'll discuss whether and to what extent the definition can be precisified. To begin, let us temporarily distance ourselves from the specific problem of representation. Consider the following two algorithms, roughly described.

```
+1 Algorithm              Table Algorithm
READ input                READ input
WRITE input + 1           IF input = 1
                              WRITE 2
                          ELSE IF input = 2
                              WRITE 3
                          ELSE IF input = 3
                              WRITE 4
                          ...
```

Modulo spatial constraints on physical implementation, each procedure maps the same inputs to the same outputs. Each computes the same function, *successor*, but in very different ways. The +1 Algorithm computes the successor function by applying the same rule to each possible input. Because it applies the same procedure to *all* inputs, it is an example of what I will call a UNIFORM RULE. By contrast, the Table Algorithm treats each input differently, applying a *different* procedure to each individual case. It is an example of a INDIVIDUATED RULE. By "rule" I mean any schematic procedure which maps inputs to outputs by a process sufficiently incremental and deterministic that it can be carried out by a mechanical device. When referring to computational algorithms in

particular, I will speak of uniform and individuated algorithms

In general, uniform rules apply the same procedure to each input in order to derive its output. Individuated rules apply a different procedure to each input in order to derive its output. This distinction is difficult to make entirely exact, in part because a rule is itself a procedure, so there is one sense in which even an individuated rule applies the "same" procedure to each input. To fend off this misunderstanding, we might say that, in the case of a individuated rule, the procedure applied to each input is specified by a different part of the global rule; while in a uniform rule, the procedure applied to each input is same part of the global-rule. Even this way of putting things faces tricky counterexamples, but it is enough for now to convey a reasonably clear idea of the distinction at hand.

My analysis of representational mode among simple representations can now be stated succinctly:

(3)   *I* is a system of symbolic representation iff *I* is a individuated rule.

The analysis is most clearly elaborated with respect to the interpretive algorithms underlying public systems of representation. After explaining this application I'll briefly discuss the ways in which the same analysis may be extended the domain of mental representation. Consider again the algorithms used to compute the S-Code and I-Code:

```
I-Code                                        S-Code
READ sign                                     READ sign
WRITE { w | level(w) = bright(sign)*10 }      IF bright(sign) = 7
                                                  WRITE { w | level(w) = 10 }
                                              ELSE IF bright(sign) = 3
                                                  WRITE { w | level(w) = 20 }
                                              ELSE IF bright(sign) = 6
                                                  WRITE { w | level(w) = 30 }
                                              ELSE IF bright(sign) = 1
                                                  WRITE { w | level(w) = 40 }
                                              ...
```

The S-Code is a clear example of a individuated rule. It applies a distinct procedure to each possible input to derive its content. And as I argued above, the S-Code is a symbolic system of representation. By contrast, the I-Code is a clear example of a uniform rule. It applies the same procedure to all possible inputs in order to derive their content. And it is an iconic system of representation. Thus, at least for these simple cases, our pre-theoretic judgement about the distribution of iconic and symbolic modes covaries closely with the individuation and uniformity of the interpretive algorithms which correspond to those systems. The same analysis applies smoothly to other examples of simple symbolic representation. Paul Revere's famous symbolic code, for exam-

ple, was obviously individuated, for it contained a distinct clause for each possible input sign: *one* if by land, *two* if by sea. In the same way, linguists typically specify the lexicons of natural language by tables which separately match each possible lexeme with its denotation. Lexicons are core cases of symbolic representation, and they are realized by individuated algorithms.

This account of symbolic representation allows for a precise rendering of Saussure's intuitive ascription of arbitrariness, which I describe briefly here, and return to in more detail later. In their most natural implementation, individuated algorithms include, for each possible sign input, a single clause which assigns content to that input. In the S-Code, for example, each sign is mapped to its content by a single clause of the form "IF f(sign) = n, WRITE *content*". Because of the modularity of this algorithmic architecture, it is always possible to replace one such clause with a new clause which maps a different sign to the same content— without disrupting the normal functioning of the rest of the procedure. And it is in this sense that symbolic representation is arbitrary: for any sub-routine of an algorithm which maps a given sign to its content, an alternative sub-routine could be used to map a different sign to the same content, while holding fixed every other sub-routine of the algorithm. Such a construal of arbitrariness is broadly counterfactual, since it conforms to the general idea that a sign is arbitrary just in case a different sign *could* have been used to represent the same content. But the range of possible arbitrary variations is fixed by the architecture of the interpretive algorithm, not the causal history of the representation.

As expected, the I-Code is not arbitrary in this way, for there is no obvious modular transformation which can discretely replace one clause for another, without disrupting the rest of the procedure. For example, suppose I set out to modify the I-Code so that a lamp state of brightness 7 is mapped to the reservoir being 10% full, but the rest of the algorithm remains fixed. The most natural way to do this is to *add* an exception clause to the procedure, as illustrated at right. There is no straightforward way of replacing the existing WRITE clause with another to achieve the same result.[6]

```
I-Code
READ sign
WRITE {w | level(w) = bright(sign)*10}
```

```
Modified I-Code
READ sign
IF bright(sign) = 7
    WRITE {w | level(w) = 10}
ELSE
    WRITE {w | level(w) = bright(sign)*10}
```

This is the result we should expect. The I-Code is most akin to a system of diagrammatic or pictorial

---

[6]But it *can* be done— by replacing bright() itself with a individuated algorithm. This particular maneuver could be avoided by making the definition of bright() explicit in the code itself. Yet even then there may be no fool-proof way of ruling out tricks like these.

representation. It does not determine an arbitrary relation between its signs and their content. And it fails the characterization of arbitrariness on offer here.

Naturally, the question arises whether the present analysis can be carried over to *iconic* systems of representation as well. Do simple iconic systems correspond to all and only the uniform algorithms? I am not sure. Consider the interpretive rules which govern indexical representations like directional arrows and pointings. These are certainly uniform, in the sense that that the same rule maps every arrow and its position to an indicated direction. But I suspect that there are good reasons for distinguishing such indexical representations from iconic representations. If so, then rule individuation is not a sufficient condition for iconic representation.

On the other hand, it does seem that rule uniformity is a *necessary* condition on iconic representation. We saw that the I-Code is both iconic and uniform. Although, officially, I am not yet considering higher-order systems of representation, it is easy to see how a semantics of pictorial representation based on geometrical projection, such as that presented in Chapter 3, lends itself to a uniform algorithm as well. Setting aside certain details, for any given picture p and projective index i, the algorithm outputs {s | lin(s,i) = p}. It too is uniform. If all iconic systems are uniform algorithms, this makes sense of the widespread characterization of iconic representation as "motivated" or "natural". In a unfiorm algorithm, every sign is mapped to its content by the same procedure. Since the same procedure can be applied to every sign, it must be general and systematic. Thus iconic representation is general and systematic, a surrogate for "natural". And since the same procedure is applied to every input, the variation in content must derive from the variation in the intrinsic features of the sign. Thus the content of the sign is "motivated" by the structure of the sign itself.

These considerations suggest that the distinction between symbolic and non-symbolic representation is the primary taxonomic distinction among representational modes, corresponding as it does to the basic distinction between individuated and uniform rules. Further distinction, for example, among pictures, diagrams, and indexical signs, correspond to differences among the variety of uniform interpretive rules.

I have defined an interpretive algorithm as a translational mapping from signs in a public representational system to those of a private, cognitive representational system. What if it turns out that all mental representation is symbolic? Does this imply that all iconic representations are therefore symbolic? It does not. The statuses of public and mental representational modes, while parallel, are autonomous. The distinction between iconic and symbolic public signs is pre-theoretical, clear without significant empirical understanding of private representation. Even *if* all mental represen-

tation is symbolic, there remain obvious and unaccounted for differences between, say, the systems of natural language, and those of perspective depiction. While the ultimate mapping from iconic public representations to content may go via symbolic mental representation, the first, public-to-private stage of this mapping, is what significantly distinguishes public systems of different modes. Differences in this stage of the mapping are what I have described above.

The more pressing challenge for the rule-based analysis of representational mode is whether it can be applied to the potentially various modes of *mental* representation. Thus far I have only discussed the individuation/uniformity distinction with respect to interpretive rules; but mental signs do not, in general, acquire their content through a process of interpretation. Rather, mental content is a product, broadly speaking, of a sign's causal and functional role in an environmentally embedded computational system. Still, I believe the present analysis can be extended to the representational mode of mental signs. But I regret that my comments here will be at best sketchy and circumspect.

I begin with a series of basic assumptions which I hope are held in common by all representational theories of mind which are more or less naturalistic. First, I assume that the mapping of signs to content, whether mental or public, is deterministic, so can be modeled by an interpretation function. Second, I assume that these mappings emerge from rules, broadly construed, again, as any type of schematic mechanism for mapping inputs to outputs. Rules may be realized by algorithms, patterns of causal determination, laws, or norms. Third, I assume that these rules themselves have internal structure which can be richer than the pairings of inputs and outputs. For example, I assume that the rules which govern the interpretation of mental lexemes have a different and simpler structure than the compositional rules which govern the interpretation of complex mental representations. Finally, I assume the rule structure may vary by representational systems. In the context of mental representation, such systems may correspond to the various modules or sub-modules posited by contemporary cognitive scientists. In sum: I assume that the mappings from signs to content, whether mental or public, are grounded in internally structure rules, which vary according to representational system.

The many available theories of intentionality posit very different kinds of rules grounding mental representation. *Whatever* theory ultimately prevails, I propose that the distinction between individuated and uniform rules, applied to the psychosemantic conditions of that theory, will track the distinction among representational modes for mental representations.

Consider, for example, the causal lineages linking singular signs to environment properties posited by causal theories of reference. For each mental "word", a distinct lineage links that sign to

150

its referent. The entire mental "lexicon" (for a particular cognitive system) is based on the collection of all such lineages for the signs of that system. But these causal paths need not result from any one mechanism; instead, the lexicon is grounded in a sequence of largely independent causal threads linking signs to content. It is governed by an individuated causal rule; thus it is a symbolic system of representation. By contrast, consider the causal mechanism, whatever it is, which underlies the correlation of perceptual content with perceptual representations. It is obvious that there is no dedicated causal mechanism linking, say, a perception of an edge at a particular position with the content that there is an edge at that position. Rather, there is a general causal mechanism, embodied in the visual system, which links any of a range of perceptual representations to any of a range of distal stimuli. The psychosemantics of perception, whatever form it ultimately takes, and however it is ultimately implemented, is governed at least in part by a uniform causal rule; thus it is, at least in part, an iconic system of representation.

I mention these examples by way of suggestive illustration; they are too noisy to ground any robust theoretical conclusions. (For example, the inputs to perceptual interpretation are not simple representations.) The following more idealized cases may make the point more vividly. Consider two machines. One maps temperature to binary strings according to a tally system, reminiscent of the I-Code: for every full degree fahrenheit above 0, the system outputs a 1. When the temperature is n degrees fahrenheit, the system outputs n 1's. Call this the I-Machine. Another machine maps temperatures to tally marks according to a tabular system, reminiscent of the S-Code: 1°F is encoded as "11111", 2°F as "111", 3°F as "111111", and so on. The outputs of each machine probably do not qualify as *representations* unless they are embedded in more complex computational devices which use the outputs to guide action. But suppose they are so embedded. The outputs of each are representations.

It is clear here that the representational content of the outputs of each machine are grounded by their causal connection to the environment, and also that this causal connection has a very different character in each case. The S-Machine establishes a causal relation between each possible temperature and a string of 1's, but in a way that is piecemeal. For each sign, a distinct causal mechanism connects it to its environmental counterpart. Thus the S-Machine is governed by a causally realized individuated rule. A representational system which depends on the S-Machine is therefore dependent on an individuated representational rule. By the analysis proposed here, the representational mode of the system is symbolic. By contrast, the I-Machine contains a mechanism which can causally relate any of a variety of states in the world to any of a variety of strings of 1's, in a manner which is entirely general. There is no dedicated mechanism which relates one particular

string to a particular temperature in the environment. Thus the I-Machine is governed by a causally realized uniform rule. A representational system which depends on the I-Machine is therefore dependent on a uniform representational rule. By the analysis proposed here, the representational mode of this system is non-symbolic, and possibly iconic.

Without being able to provide much more detail, I hypothesize that the symbolic representational mode, manifested in both public and mental representation, is grounded in individuated interpretive rules. Non-symbolic modes, including iconic representation, are grounded in uniform interpretive rules. Each employs a fundamentally different strategy for storing information in a physical medium, reflected by correspondingly distant causal and algorithmic architectures. If the analysis proposed here is correct, than the core semiotic distinction between symbolic and iconic representation reflects this basic difference in semantic organization.

4.3 The architecture of representational rules

I now turn to discussing, in greater detail, the architecture of individuated rules which underly symbolic representation. I have leaned heavily on the informal distinction between a uniform and individuated rule. Can we make this idea more precise? Yes and no. In this section, I'll elaborate on the key distinction as applied to algorithms, and indicate some of its important consequences and properties. Disappointingly, I know of no way to make this distinction perfectly rigorous, and I suspect that there is no way to do so.[7]

One very natural way to implement an individuated algorithm is by using a LOOK-UP TABLE. A look-up table is an algorithmic architecture in which an explicit representation of an input is paired with an explicit representation of the output. (Gallistel and King 2009, p. 51) Mapping a sign to its content using a look-up table requires testing the sign against every entry in the table until the sign is found. If the table is long, this could be time consuming. On the other hand, no more intelligent operation need be performed to compute the output than looking down the table. By contrast, a typical uniform algorithm requires the application of much more complex computations.

Individuated algorithms based on look-up tables differ from uniform algorithms in other ways that are relevant to interpretation. A look-up table requires an explicit representation of each and every possible input. The algorithm itself must store, or have access to, the entire suite of possible signs. As a consequence, an individual who acquires the ability to interpret in a given symbolic

---

[7]In practice, drawing the distinction between systems that follow a uniform or individuated rule seems may depend on facts about when a given event or action counts as following a given rule. This idea in turn has been the subject of considerable philosophical scrutiny. (Saul Kripke 1982) The moral of that discussion seems to be that the concept of following a rule cannot be fully formalized.

system must have a mental representation for every symbol in the system. By contrast, it is always possible to specify a uniform algorithm which has *no* representations of any inputs. Competence with an interpretive algorithm of this kind requires the ability to apply a general rule to a range of inputs and nothing more.

In any physical implementation of a look-up table, a representation of each input-output pair occupies a distinct position in physical memory. This fact too has important consequences. Since each entry in the table is physically independent, entries can be added and removed piecemeal, without impairing the overall function of the computational machinery. For symbolic systems, this means that it is easy to add new simple symbols to the lexicon, and easy to drop or replace old ones. By contrast, in iconic systems, there is no straightforward way of adding new signs or jettisoning old ones with making significant changes to the interpretive rules.

The fact that entries in a table can be easily replaced with new ones leads to the definition of arbitrariness offered in the previous section. It is interesting to make this definition explicit. Let I be a system of representation which assigns a sign S to the content C. The representation of C by S is arbitrary in I because it is possible to replace S with alternative sign S' such that, in the resulting system I', S' carries C, but in all other respects, I' is just like I. I take this to be a reasonable precisification of Saussure's guiding idea. Here is a first attempt to render it more exactly:

(4) **Symbolic Representation** — First Attempt

For any system of representation I and sign S in I: S is symbolic in I iff

there is system $I' \neq I$ and sign S' not in I and structurally distinct from S such that:

(a) I' is derived from I by replacing S with S' throughout;

(b) For every sign $E \neq S$ in I: $I(E) = I'(E)$, and I(S)=I'(S').

According to this definition, a sign is symbolic in a system of representation just in case that sign could be replaced with a structurally distinct alternative sign, and, modulo this change, the system could go on functioning as before. Clause (a) describes the replacement component of the definition; clause (b) describes the "go on functioning as before" component.

To see how the definition applies to symbolic systems of representation, consider again the S-Code. It is possible to make a modular change to this code, by replacing the representation of one input with another, without engendering any more holistic alteration. In the example below, the content assigned to brightness level 7 is re-mapped to brightness level 0 in the modified code. Hence, according to the definition, brightness level 7 is symbolic with respect to the S-Code.

153

```
S-Code                              S'-Code
READ sign                           READ sign
IF bright(sign) = 7                 IF bright(sign) = 0
    WRITE { w | level(w) = 10 }         WRITE { w | level(w) = 10 }
ELSE IF bright(sign) = 3            ELSE IF bright(sign) = 3
    WRITE { w | level(w) = 20 }         WRITE { w | level(w) = 20 }
ELSE IF bright(sign) = 6            ELSE IF bright(sign) = 6
    WRITE { w | level(w) = 30 }         WRITE { w | level(w) = 30 }
ELSE IF bright(sign) = 1            ELSE IF bright(sign) = 1
    WRITE { w | level(w) = 40 }         WRITE { w | level(w) = 40 }
...                                 ...
```

I illustrated in the previous section how similarly modular changes cannot be made to uniform codes such as the I-Code.

Yet this definition fails to correctly classify signs which are components of other, complex signs that are inputs to the same system. Though my aim, for now, is only to define symbolic representation for atomic signs, the definition should apply to atomic signs even when they figure in systems which also interpret complex signs. The definition fails because if a complex expression E contains S, and is not identical to S, then it will not be the case, once the modification of clause (a) is implemented, that $I(E) = I'(E)$.

Luckily the problem is easily solved by altering clause (b). Instead of requiring that the two systems be exactly alike for every sign, save for their treatment of S and S', we now require that the two systems be exactly alike for every sign *when S' replaces S in every case*.

(5) *Symbolic Representation* — Second Attempt

For any system of representation I and sign S in I: S is symbolic in I iff

there is system $I' \neq I$ and sign S' not in I and structurally distinct from S such that:

(a) I' is derived from I by replacing S with S' throughout;

(b) For every sign E in I: $I(E) = I'(E[S/S'])$.

Insofar as the key concepts of the definition— such as the replacement of a sign within an algorithm— can be made explicit, the definition here offers a significant clarification of Saussure's original idea, at least as applied to symbolic systems which rely on look-up tables to compute the content of their simple signs.

Explicit representation of inputs, piecemeal extension, and arbitrary replaceability— all of these are relevant properties of individuated algorithms implemented as look-up tables. They are particularly useful because each can be precisely defined. Unfortunately, none of these definitions can serve as a definition of individuated algorithms in general, because it is possible to implement individuated algorithms *without* using look-up tables, and many such implementations

satisfy none the conditions described above. (I must postpone a detailed elaboration of the evidence for this claim.) The variety of ways in which this can be done resists formal characterization.

But there is one possible definition which is not implementation-bound in this way. In a great many cases, uniform algorithms are exactly those which computer scientists term COMPACT. A compact algorithm is one whose code takes up less physical space than a pair-wise list of its possible inputs and outputs. (Gallistel and King 2009) Since many finitely-specifiable uniform rules have infinitely many inputs and outputs, they are trivially compact. Even when the domain of interpretation is explicitly restricted, the generalization typically holds. Furthermore, a pair-wise list of inputs and outputs is, within a reasonable margin, the same size as the code for individuated algorithm with the same inputs and outputs. So in a great many cases, individuated algorithms are exactly those which are not compact. Thus compactness is a useful heuristic for diagnosing rule type. But it is not an analysis. It is possible to devise apparently uniform rules applied over narrow finite domains which are not compact. Thus the concepts of uniformity and individuation remain formally elusive.

4.4 The hierarchy of representational modes

Thus far I have discussed representational mode only with respect to simple signs. We could stop there. Perhaps representational mode is a feature of simple signs alone: systems which countenance complex signs contain subsystems which countenance only simple signs, and these are the primary bearers of representational mode. On this view, though pictures and sentences are complex, they inherit their representational mode from that of their simplest parts.

But there are compelling reasons to reject this proposal. Most striking is the observation of complex signs which appear to employ different modes of representation at different levels of composition. For example, consider a seating chart consisting only of names spatially arranged on a page to reflected the intended seating position of the individuals named.[8] The simple constituents of the sign are names— paradigmatic symbolic representations. But the way in which the names are arranged is determined by some system of geometrical projection— characteristic of iconic representation. In this case it seems that the complex sign is iconic with respect to composition, but symbolic with respect to simple parts.

This order can be reversed. Imagine a code that used a sequence of lamp flashes to indicate the water level in each of a group of reservoirs. Each lamp flash is governed by something very much like the I-Code described above, save that the identity of the reservoir is a variable parameter;

---

[8]The example is due to Jeff King.

but the correspondence between sequential position and reservoir identity is simply stipulated (1st flash → reservoir X, 2nd flash → reservoir Y, and so on). Here, the simple parts are iconic representations, yet their mode of composition is distinctly symbolic. After all, the correspondence between sequential position and reservoir identity is entirely arbitrary. In this case, it seems the complex sign is symbolic with respect to composition, but iconic with respect to simple parts.

These examples demonstrate that representational mode is not a matter of simples alone. It manifests in different ways at different levels of sign composition. Therefore we must replace the blanket concepts of iconic and symbolic representation with those of ATOMICALLY ICONIC (SYMBOLIC) representation and COMPOSITIONALLY ICONIC (SYMBOLIC) representation. The seating chart is composed of atomically symbolic representations; but the chart itself is a compositionally iconic representation of its subject. Note that, because composition is hierarchical, each application of a compositional rule has the potential to introduce a new variation of representational mode. A sequence of seating charts, for example, could be used to represent each of several seating arrangements in different rooms, where sequence position is arbitrarily associated with room identity. Such a complex sign would be compositionally *symbolic*; it would have parts (the seating charts) that were compositionally *iconic*; these in turn would have parts (the names) which were atomically *symbolic*. Thus a hierarchy of representational modes emerge. (It would be profitable to spell out the semantics for such a system in detail, but I postpone this task for now.)

This analytical framework can illuminate the mechanisms of more familiar complex signs. In a color digital photograph, for example, each pixel can be thought of as an iconic simple: the color of the pixel corresponds to the color of a surface region in the scene in a general, characteristically iconic manner, reminiscent of the I-Code. Meanwhile, the arrangement of pixels on the picture plane is determined by a different general rule of geometrical projection. The signs of such a system are both atomically iconic and compositionally iconic.[9]

A plausible, if somewhat surprising, proposal is that line drawing is compositionally iconic, but atomically *symbolic*. In such systems, colored points in the picture are associated with contours at a point in the scene. Allowing the points to be the simples of this system, we can detect a kind of arbitrariness in this rule. Instead of points of a given color, we could equally well associate points of another color, or regions (of various shapes) with contours. The point is even more vivid for the case of systems of depiction which use solid lines to depict unoccluded contours, but dotted lines to depict occluded contours. The mapping here is obviously arbitrary; an informationally equivalent

---

[9]In the presentation of the pictorial semantics of Chapter 3, I did *not* quite follow this characterization. I did not explicitly assign content to points independent of their position. Still, a semantics of this form is not hard to envision. I hope to implement this change in the near future.

system could map dotted lines to unoccluded contours, and solid lines to occluded contours. Thus line drawing and color photography are alike in compositional iconicity, but arguably different with respect to the representational mode of simples.

In natural languages, the lexical simples are symbolic representations. But the composition rules are also plausibly symbolic. Their arbitrariness is evidence for this claim. For example, suppose a semantics includes the following composition rule:

(6) For all $\phi, \psi : [\![\ulcorner \phi\psi \urcorner]\!] = [\![\phi]\!]([\![\psi]\!])$.

That is, given a concatenation of two, possibly complex signs, the content of the whole is the result of applying the content of the first to the content of the second. Yet the same semantic function could be served by the following composition rule (so long as counterpart changes were made in the syntax). Here, the function-functor order is reversed: the content of the whole results from applying the content of the second to the content of the first.

(7) For all $\psi, \phi : [\![\ulcorner \psi\phi \urcorner]\!] = [\![\phi]\!]([\![\psi]\!])$.

And even if, in decomposition, sister nodes are unordered, alternative composition rules could have required the presence of a special concatenation symbol— but would again have served the same semantic function:

(8) For all $\phi, \psi : [\![\ulcorner \phi \circ \psi \urcorner]\!] = [\![\phi]\!]([\![\psi]\!])$.

Thus, at the clausal level, natural languages are typically compositionally symbolic, while their simple parts are atomically symbolic. Interestingly, considerable evidence shows that, in the case of narrative discourse, language is *iconic* at the discourse level. Specifically, sentence order is used to represent temporal sequence, much the way spatial order is used in a time line. (**cummings2010narrative**) Thus narratives are compositionally iconic, while their sentential parts are compositionally symbolic, and their simple, lexical parts are atomically symbolic.

The distinction between compositionally iconic and symbolic representation can be drawn fairly neatly by appeal to the distinction between individuated and uniform interpretive rules. The application of this framework to rules of composition might at first seem confusing. *All* composition rules of practical use are uniform, in a certain sense, since any such rules applies the same compositional computation to any of a range of possible constituents. (If the composition rules are recursive, the point is even more vivid.) But here it is important to recall that signs themselves are signal *types*. By further abstraction, we arrive at higher-order *sign types*— those which corre-

spond to the inputs to composition rules. It is these sign types which are mapped to content by compositional rules which are either uniform or individuated.

We must accordingly distinguish between two kinds of individuated rule: on one hand, there are rules which are classified as individuated *with respect to simple signs*— rules which map each simple sign to its content according to a different rule. On the other hand, there are rules which are classified as individuated *with respect to complex sign types*— rules which map each complex sign type to its content according to a different rule. Using this distinction, I revise the initially proposed analysis of symbolic representation as follows:

(9) (a) *I* is a system of atomically symbolic representation iff

  I is an individuated rule with respect to simple signs.

  (b) *I* is a system of compositionally symbolic representation iff

  I is an individuated rule with respect to complex sign types.

By way of illustration, consider a language with two composition rules, one for binary branching trees, and one for ternary branching trees, as follows. (For convenience, I'll represent binary trees by pairs, and ternary trees by triples)

(10) For all ϕ, ψ, χ:

  (i) $[\![\langle \phi, \psi \rangle]\!] = [\![\phi]\!]([\![\psi]\!])$.

  (ii) $[\![\langle \phi, \psi, \chi \rangle]\!] = [\![\psi]\!]([\![\phi]\!], [\![\chi]\!])$.

Here, a distinct rule maps each complex sign type to its content. Relative to the level of typological abstraction which these rules themselves exploit, the composition rules of this language are compositionally individuated, hence symbolic. The same considerations apply to languages which employ just one composition rule.

By contrast, even at the level of abstraction of complex sign types, the composition rules for pictures are uniform. Assume a pictorial semantics in which the simples are colored points, and the complex signs are spatial arrangements of colored points. The semantic rules for the simples map each point of a given color to a property; the rules of composition map each point and its position in the overall spatial arrangement to the set of scenes in which that point's viewpoint-relative projective counterparts instantiate the point's simple content.[10] Such composition rules are emphatically uniform: for every possible position on the picture plane, they map the simple at that position to a content, according to a general projective rule. It is not as if there are distinct rules for the position $\langle 0, 0 \rangle$, another for $\langle 0, 1 \rangle$, another for $\langle 1, 1 \rangle$ and so on. Thus the compositional iconicity

---

[10]See Chapter 3, and footnote 9 for details.

of complex signs in the system of line drawing is reflected by the uniformity of its compositional rules.

I do not pretend that this discussion has been sufficiently precise to qualify as a clear position statement on the representational mode of complex signs. I have merely gestured at the semantics of the systems under discussion here; ultimately, these semantics must be spelled out explicitly and cast in a standard idiom. Only then can the distinction between singular and uniform be made more explicitly for compositional rules. Instead, the discussion here is a road map for future work. I hope, at least, to have convinced the reader that the rule-based analysis of representational mode has the potential generality to profitably illuminate simple and complex signs accross the semiotic spectrum.

5   Conclusion

I argued initially that theories of representational mode which are concerned only with the formal features of signs and their contents are fatally impoverished. One way or another, representational mode is a matter of the representational relationships that obtain between signs and their content. I considered two prominent families of theory which promised to make good on this lesson: the resemblance-based tradition, and the Saussurian arbitrariness-based tradition. The former I found empirically invalid, the latter wanting in detailed explication. Finally, I sketched an alternative approach, according to which representational mode is a property of the causal and computational rules that underly the association of content with physical signs. Iconic and symbolic representation correspond to fundamentally different possible architectures which these rules may assume. I argued that these structural features faithfully track representational mode in both simple and complex signs. I hope thereby to have detected the rudiments of a useful representational taxonomy.

Chapter 5

The Semiotic Spectrum

At the beginning of this dissertation, I hypothesized a continuous spectrum of semantic kinds, extending from the purely symbolic to the purely iconic. Over the course of the last four chapters, I have taken initial steps towards establishing this claim. For I have argued that there are representational systems whose internal mechanisms are profoundly divergent, but which are all genuinely semantic. Thus I have sketched the outer limits of semantic variation.

Symbolic representation, the most basic and familiar semantic kind, exploits interpretive rules which modularly associate particular signs with their content. Symbolic representation is opposed with systems which rely on general transformations to map any one in a field of signs to their corresponding content. I have argued that this distinction marks a basic difference in the strategies available for encoding content in sign structure. Among the latter, transformational systems, an important class exploit methods of geometrical projection from a viewpoint. These are the pictorial systems: they are abstractions and extensions of the psychosemantics of perception. I have shown that pictorial representations can be analyzed in terms which are just as precise and systematic as those commonly applied to symbolic systems. All are subject to analysis under a general semantics.

Of course the phenomena studied here scratch only the surface of the variety and complexity of representational kinds that figure in our daily communications. Some of this complexity simply results from the combination of signs that belong, on their own, to purely pictorial or purely linguistic systems. Formal semantic analyses of such *multi-modal* representations have been carried out recently for the case of speech with iconic gesture. (Lascarides and Stone 2009*a*; Lascarides and Stone 2009*b*) By bringing together the insights of these studies with the account of pictorial representation proposed here, it may be possible to arrive at productive analyses of complex multi-modal signs like the two illustrated below. The first, a daily weather broadcast, contains a rich blend of text, spoken word, video imagery, and map depiction. The second, part of a detailed visual explanation

provided by my dentist prior to a root canal, integrates perspective line drawing with sequenced textual annotations.

Other species of sign systems fuse iconic and symbolic elements more thoroughly, and in a manner less tractable to theoretical framework of this dissertation. Consider the two Mayan figurines pictured below. The three-dimensional representation on the left is recognizably pictorial in a sense which is now familiar. For it is plausible that there exists a projective transformation which would map a real human head to the figure on display. The transformation would involve a reduction in size and an abstraction from structural detail. But this form of representation falls broadly within the range of phenomena I have been discussing. More puzzling is the sculpture pictured at right. This highly stylized representation of a human is the result of no generally applicable transformation. For consider the fact that there is no obvious way to extend the same system of representation to the portrayal of, say, an elephant or a television. The figure is the product of a system which is at once arbitrary and broadly pictorial. Stylization of this kind is anathema to the kind of analysis I have been developing here.


And yet, such stylization is ubiquitous. More contemporary examples are rendered in the popular "DOT style", in which human limbs and torsos are depicted by connected black oblongs, while the head is depicted by a disconnected black circle. Such images result from what Gombrich (1960) called "schemas"— predetermined recipes for drawing specific kinds of objects. In these cases, a

schema dictates that heads are represented by floating circles.

Are such representations icons or symbols? They seem to involve aspects of both modes. The gross spatial arrangement of parts is a product of projective geometry, as evidenced by the difference in spatial configuration of the two bodies depicted above. But the arbitrary choice to render the head as a circle— as a opposed to a square or an oval— is strongly reminiscent of the sort of singular procedure I have connected with symbolic representation. How should we construct a semantics for such images? And where do they stand in relation to pictures, on one hand, and words, on the other? I do not know.

However these questions are resolved, the answers must also illuminate the connections between these images and graphics which use spatial structure in ways that are both more abstract and more concrete. The directional arrows pictured on the left, for example, derive their meaning in part from their own spatial orientation. The graph on the right, sketched by my banking representative to illustrate the relationship between "linked" bank accounts, also relies on spatial relationships to make its point, but invokes neither orientation nor projection.

It is foolish to expect that all forms of representation will fit into neat and rational categories; new systems of representation are continually invented to fit the evolving demands of communication. Yet the variation they exhibit is not random. No student of language or of depiction can

afford to leave aside the basic question which such signs force upon us: what are the dimensions along which the myriad systems of representation vary? I have not managed to find an answer to this question, or even to settle on any reasonable conjectures. Instead, I hope that my contribution has been to ask the question more clearly.

References

Abell, C. (2007). "Pictorial realism". In: *Australasian Journal of Philosophy* 85.1. Pp. 1–17.
— (2009). "Canny Resemblance". In: *The Philosophical Review* 118.2. P. 183.
Abell, C. and K. Bantinaki (2010). *Philosophical perspectives on depiction*. Oxford University Press.
Adams, R.M. (1979). "Primitive thisness and primitive identity". In: *The Journal of Philosophy* 76.1. Pp. 5–26.
Alberti, L.B. (1991). *On Painting*. Ed. by C. Grayson. Penguin Classics.
Arnheim, R. (1974). "Hansen's Curvilinear Perspective". In: *The Journal of Aesthetics and Art Criticism* 32.3. Pp. 424–424.
Bach, K. (1970). "Part of what a picture is". In: *The British Journal of Aesthetics* 10.2. Pp. 119–137.
Balog, K. (2009). "Jerry Fodor on non-conceptual content". In: *Synthese* 170.2. Pp. 311–320.
Barthes, R. (1968). *Elements of semiology*. Hill and Wang New York.
Beardsley, M.C. (1958). *Aesthetics: problems in the philosophy of criticism*. Harcourt, Brace and Company.
Block, N. (1982). *Imagery*. MIT Press.
— (1983). "The photographic fallacy in the debate about mental imagery". In: *Nous* 17.4. Pp. 651–661.
Blumson, B. (2007). "Resemblance and Representation". PhD thesis. Australian National University, Canberra.
— (2009*a*). "Defining Depiction". In: *British Journal of Aesthetics* 49.2.
— (2009*b*). "Pictures, perspective, and possibility". In: *Philosophical Studies*.
Borg, E. (2006). "Intention-based semantics". In: *The Oxford Handbook of Philosophy of Language*. P. 250.
Burge, T. (2010). *Origins of Objectivity*. Oxford University Press.
Carnap, R. (1963). *The Philosophy of Rudolf Carnap*. Ed. by P.A. Schilpp. Open Court Publishing Company.
Carroll, N. (1999). *Philosophy of art: A contemporary introduction*. Psychology Press.
Chomsky, N. (1965). *Aspects of the Theory of Syntax*. The MIT press.
Clowes, M.B. (1971). "On seeing things". In: *Artificial Intelligence* 2.1. Pp. 79–116.
Cohen, J. and C. Callender (2006). "There is no special problem about scientific representation". In: *Theoria* 21.55. P. 67.
Costello, D. and D.M. Phillips (2009). "Automatism, Causality and Realism: Foundational Problems in the Philosophy of Photography". In: *Philosophy Compass* 4.1. Pp. 1–21.
Cummins, R. (1996). *Representations, targets, and attitudes*. The MIT Press.
Da Vinci, L. et al. (1970). *The Notebooks of Leonardo da Vinci*. Dover Publications.
DeCarlo, D. and M. Stone (2010). "Visual explanations". In: *Proceedings of the 8th International Symposium on Non-Photorealistic Animation and Rendering*. ACM. Pp. 173–178.
DeCarlo, D. et al. (2003). "Suggestive contours for conveying shape". In: *ACM Transactions on Graphics (TOG)* 22.3. Pp. 848–855.
Descartes, R. (1637/2001). *Discourse on method, optics, geometry, and meteorology*. Ed. by P.J. Olscamp (trans.) Hackett Publishing.
Dilworth, J. (2005). "Resemblance, Restriction, and Content-Bearing Features". In: *Journal of Aesthetics and Art Criticism*. Pp. 67–70.
Dretske, F. (1981). *Knowledge and the Flow of Information*. MIT Press, Cambridge, MA.
— (1983). "Knowledge and the Flow of Information". In:

Dummett, M. (1996). "What Is a Theory of Meaning? (I)". In: *The Seas of Language* 1.9. Pp. 1–34.

Durand, F. (2002). "An invitation to discuss computer depiction". In: *Proceedings of the 2nd international symposium on Non-photorealistic animation and rendering*. ACM. Pp. 111–124.

Eco, U. (1979). *A Theory of Semiotics*. Indiana University Press.

Feldman, J. (1999). "Does vision work? Towards a semantics of perception". In: *What is cognitive science?* Ed. by E. Lepore and Z. Pylyshyn. Basil Blackwell. Pp. 208–229.

Field, H.H. (1978). "Mental representation". In: *Erkenntnis* 13.1. Pp. 9–61.

Files, C. (1996). "Goodman's rejection of resemblance". In: *The British Journal of Aesthetics* 36.4. Pp. 398–412.

Fodor, J. (1975). *The language of thought*. Harvard University Press.

— (1990*a*). "A theory of content, I: The problem". In: *A Theory of Content and Other Essays*. Pp. 51–87.

— (1990*b*). "A theory of content, II: The theory". In: *A Theory of Content and other Essays*. Pp. 89–136.

— (1992). *A theory of content and other essays*. MIT Press, Cambridge, MA.

— (2007). "The Revenge of the Given". In: *Contemporary Debates in Philosophy of Mind*. Ed. by B.P. McLaughlin and J. Cohn. Blackwell.

Fodor, J. and E. Lepore (2002). *The compositionality papers*. Oxford University Press, USA.

Frege, G. (1918). "The Thought: A Logical Inquiry". In:

French, S. (2003). "A Model-Theoretic Account of Representation (or, I Don't Know Much about Art – but I Know it Involves Isomorphism)". In: *Philosophy of Science* 70. Pp. 1472–1483.

Gallistel, C.R. and A.P. King (2009). *Memory and the computational brain: Why cognitive science will transform neuroscience*. Vol. 3. Blackwell Pub.

Gibson, J.J. (1954). "A theory of pictorial perception". In: *Educational Technology Research and Development* 2.1. Pp. 3–23.

— (1960). "Pictures, perspective, and perception". In: *Daedalus* 89.1. Pp. 216–227.

— (1971). "The information available in pictures". In: *Leonardo* 4.1. Pp. 27–35.

— (1978). "The ecological approach to the visual perception of pictures". In: *Leonardo* 11.3. Pp. 227–235.

Giere, R. (2004). "How Models Are Used to Represent Reality". In: *Philosophy of Science* 71. Pp. 742–752.

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

Goodman, N. (1968). *Languages of Art: An Approach to a Theory of Symbols*. The Bobbs-Merrill Compnay, Inc.

— (1972). "Seven strictures on similarity". In: *Problems and projects*. Pp. 437–447.

Greenberg, G. (2010*a*). "Beyond Resemblance". Unpublished manuscript. Available at: `http://eden.rutgers.edu/~gabrielg/`.

— (2010*b*). "The Semiotic Spectrum". PhD thesis. Rutgers University.

Grice, P. (1957). "Meaning". In: *Philosophical Review* 66.

— (1975). "Logic and Conversation". In: *Syntax and Semantics, Volume 3: Speech Acts*. Ed. by Peter Cole and Jerry Morgan. New York: Academic Press. Pp. 41–58.

Groenendijk, J. and M. Stokhof (1991). "Dynamic predicate logic". In: *Linguistics and philosophy* 14.1. Pp. 39–100.

— (1997). "Questions". In: *Handbook of Logic and Language*. Ed. by J. van Bentham and A. ter Meulen. http://staff.science.uva.nl/ stokhof/q.pdf: Elsevier/MIT Press.

Gunderson, Keith, ed. (1975). *Language, Mind and Knowledge*. Vol. 7. Minnesota Studies in the Philosophy of Science. Minneapolis: University of Minnesota Press.

Gunther, Y.H. (2003). *Essays on nonconceptual content*. The MIT Press.

Hagen, M.A. (1986). *Varieties of Realism: Geometries of Representational Art*. CUP Archive.

Hamblin, C.L. (1958). "Questions". In: *Australasian Journal of Philosophy* 36.3. Pp. 159–168.

Hansen, R. (1973). "This Curving World: Hyperbolic Linear Perspective". In: *The Journal of Aesthetics and Art Criticism* 32.2. Pp. 147–161.

Hansen, R. and J.L. Ward (1977). "Perception of Pictorial Space in Perspective Pictures". In: *Leonardo* 10.3. Pp. 261–262.

Haugeland, J. (1991). "Representational Genera". In: *Philosophy and Connectionist Theory*. Ed. by W.M. Ramsey, S.P. Stich, and D.E. Rumelhart. Erlbaum.

Helmholtz, H. (1867). *Treatise on Physiological Optics*. Ed. by J.P.C. Southall. Vol. III. Dover.(Translated from the 3 rd German edition, English edition 1962).

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

Howell, R. (1974). "The logical structure of pictorial representation". In: *Theoria* 40.2. Pp. 76–109.

Huffman, D.A. (1971). "Impossible Objects as Nonsense Sentences". In: *Machine intelligence*. P. 295.

Hyman, J. (1995). "Perspective". In: *A Companion to Aesthetics*. Blackwell Publishers.

— (2000). "Pictorial art and visual experience". In: *The British Journal of Aesthetics* 40.1. Pp. 21–45.

— (2006). *The objective eye: color, form, and reality in the theory of art*. University of Chicago Press.

Jakobson, R. (1921/1987). "On realism in art". In: *Language in Literature*. Harvard University Press.

Kaplan, D. (1968). "Quantifying In". In: *Synthese* 19. Pp. 178–214.

— (1999). "The Meaning of Ouch and Oops".

Kaplan, David (1989). "Demonstratives". In: *Themes from Kaplan*. Ed. by Joseph Almog, John Perry, and Howard Wettstein. New York: Oxford University Press.

Kennedy, J.M. (1974). *A psychology of picture perception*. Jossey-Bass Publishers, London.

King, J.C. (2007). *The nature and structure of content*. Oxford University Press, USA.

King, J.C. and J. Stanley (2005). "Semantics, pragmatics, and the role of semantic content". In: *Semantics vs. pragmatics*. Pp. 111–64.

Kivy, P. (1984). *Sound and semblance: Reflections on musical representation*. Princeton University Press.

Knight, H. (1930). "Aesthetic experience in pictorial art". In: *The Monist* 40. Pp. 74–83.

Kripke, S. (1977). "Speaker's Reference and Semantic Reference". In: *Midwest studies in philosophy* 2.1. Pp. 255–276.

Kripke, Saul (1982). *Wittgenstein on Rules and Private Language*. Cambridge, MA: Harvard University Press.

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

— (2006*b*). "Pictorial Representation". In: *Philosophy Compass* 1.6. Pp. 535–546.

Lascarides, A. and M. Stone (2009*a*). "A Formal Semantic Analysis of Gesture". In: *Journal of Semantics* 26.4. P. 393.

— (2009*b*). "Discourse coherence and gesture interpretation". In: *Gesture* 9.2. Pp. 147–180.

Lewis, D. (1969). *Convention: A Philosophical Study*. Oxford: Blackwell.

— (1970). "General Semantics". In: *Synthese* 22. Pp. 18–67.

— (1975). "Languages and Language". In: *Language, Mind and Knowledge*. Ed. by Keith Gunderson. Vol. 7. Minnesota Studies in the Philosophy of Science. Minneapolis: University of Minnesota Press. Pp. 3–35.

— (1979). "Attitudes de dicto and de se". In: *The Philosophical Review* 88.4. Pp. 513–543.

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

Lewis, S.M. and C.K. Cratsley (2008). "Flash signal evolution, mate choice, and predation in fireflies". In: *Annu. Rev. Entomol.* 53. Pp. 293–321.

Loewer, B. (1997). "A guide to naturalizing semantics". In: *A Companion to the Philosophy of Language*. Pp. 108–126.

Loewer, B. and J.W.J.R. Godbey (1978). "Representational Symbol Systems". In: *Semiotica* 23.3-4. Pp. 333–342.

Lopes, D. (1992). "Pictures, Styles and Purposes". In: *The British Journal of Aesthetics* 32.4. P. 330.

— (1995). "Pictorial Realism". In: *Journal of Aesthetics and Art Criticism* 53.3. Pp. 277–85.

— (1996). *Understanding Pictures*. Oxford University Press.

— (1999). "Pictorial color: aesthetics and cognitive science". In: *Philosophical Psychology* 12.4. Pp. 415–428.

— (2009). "Drawing in a Social Science: Lithic Illustration". In: *Perspectives on Science* 17.1. Pp. 5–25.

Loverde, T. and E.W. Weisstein (2009). "Fisheye Perspective". In: *MathWorld–A Wolfram Web Resource*. URL: `http://mathworld.wolfram.com/FisheyePerspective.html`.

Lubbock, T. (2006). "Philippe de Champaigne: Vanité (1660s)". In: *The Independent*. URL: `http://www.independent.co.uk/arts-entertainment/art/great-works/de-champaigne-philippe-vanit-1660s-798027.html`.

Mackworth, A.K. and R. Reiter (1989). "A logical framework for depiction and image interpretation". In: *Artificial Intelligence* 41. Pp. 125–155.

Malinas, G. (1991). "A Semantics for Pictures". In: *Canadian journal of philosophy* 21.3. Pp. 275–298.

Manns, J.W. (1971). "Representation, Relativism and Resemblance". In: *The British Journal of Aesthetics* 11.3. Pp. 281–287.

McCloud, S. (1993). *Understanding Comics: The Invisible Art*. Kitchen Sink Press.

Mi, X., D. DeCarlo, and M. Stone (2009). "Abstraction of 2D shapes in terms of parts". In: *Proceedings of the 7th International Symposium on Non-Photorealistic Animation and Rendering*. ACM. Pp. 15–24.

Millikan, R.G. (1984). *Language, thought, and other biological categories: New foundations for realism*. MIT Press, Cambridge, MA.

Montague, R. (1970*a*). "English as a formal language". In: *Linguaggi nella societae nella tecnica*. Pp. 189–224.

— (1970*b*). "Universal grammar". In: *Theoria* 36.3. Pp. 373–398.

Morris, C.W. (1946). *Signs, Language and Behavior*. Prentice-Hall Englewood, NJ.

Neander, K. (1987). "Pictorial representation: A matter of resemblance". In: *The British Journal of Aesthetics* 27.3. Pp. 213–226.

Necker, L.A. (1832). "Observations on some remarkable optical phænomena seen in Switzerland; and on an optical phænomenon which occurs on viewing a figure of a crystal or geometrical solid". In: *Philosophical Magazine Series 3* 1.5. Pp. 329–337.

Newall, M. (2003). "A Restriction for Pictures and Some Consequences for a Theory of Depiction". In: *Journal of Aesthetics and Art Criticism*. Pp. 381–394.

— (2006). "Pictures, Colour and Resemblance". In: *The Philosophical Quarterly* 56.225. Pp. 587–595.

Novitz, D. (1975). "Picturing". In: *Journal of Aesthetics and Art Criticism*. Pp. 145–155.

Partee, B. (1980). "Semantics– mathematics or psychology?" In: *Semantics from Different Points of View*. Ed. by Bäuerle, Egli, and von Stechow. Springer-Verlag, Berlin.

Peacocke, C. (1987). "Depiction". In: *The Philosophical Review* 96.3. Pp. 383–410.

Peirce, C.S. (1906). "Prolegomena to an apology for pragmaticism". In: *The Monist* 16.4. Pp. 492–546.

Pinker, S. (1995). *The language instinct*. HarperPerennial New York, NY.

Pirenne, M.H. (1970). *Optics, painting & photography*. Cambridge University Press.

Putnam, H. (1981). *Reason, truth, and history*. Cambridge Univ Pr.

Putnam, Hilary (1975). "The Meaning of Meaning". In: *Language, Mind and Knowledge*. Ed. by Keith Gunderson. Vol. 7. Minnesota Studies in the Philosophy of Science. Minneapolis: University of Minnesota Press. Pp. 131–93.

Pylyshyn, Z. (2006). *Seeing and visualizing: It's not what you think*. The MIT Press.

Richard, M. (1981). "Temporalism and eternalism". In: *Philosophical Studies* 39.1. Pp. 1–13.

Rogers, B. and C. Rogers (2009). "Visual globes, celestial spheres, and the perception of straight and parallel lines". In: *Perception* 38.9. P. 1295.

Ross, J. (1997). *The Semantics of Media*. Kluwer Academic Pub.

Saussure, F. (1922/1983). *Course in general linguistics*. Ed. by R. Harris. Open Court Publishing Company.

Saussure, F. et al. (1922). *Cours de Linguistique Générale*. Payot, Paris.

Schier, F. (1986). *Deeper Into Pictures: An Essay on Pictorial Representation*. Cambridge University Press.

Sebeok, T.A. (1977). *How animals communicate*. Indiana University Press.

Sedgwick, H.A. (1980). "The Geometry of Spatial Layout in Pictorial Represenation". In: *The Perception of Pictures*. Vol. 1. Academic Press. Pp. 33–90.

Shin, S.J. (1994). *The logical status of diagrams*. Cambridge Univ Pr.

Siegel, S. (2006). "Which Properties are Represented in Perception?" In: *Perceptual Experience*. Ed. by Tamar Szabó Gendler and John Hawthorne. Oxford: Oxford University Press. Pp. 481–503.

— (2009). "The contents of perception". In: *Stanford Encyclopedia of Philosophy*.

Soames, S. (1989). "Semantics and semantic competence". In: *Philosophical Perspectives* 3. Pp. 575–596.

Sorensen, R. (2002). "The art of the impossible". In: *Conceivability and Possibility*. Ed. by T.S. Gendler and J. Hawthorne. Oxford University Press. Pp. 337–68.

Speaks, J. (2010). "Theories of Meaning". In: *Stanford Encyclopedia of Philosophy*.

Stalnaker, R. (1984). *Inquiry*. Cambridge, MA: MIT Press.

— (1997). "Reference and necessity". In: *A Companion to the Philosophy of Language*. Pp. 534–554.

— (1998). "What might nonconceptual content Be?" In: *Philosophical Issues* 9. Pp. 339–352.

Stanley, J. (2000). "Context and logical form". In: *Linguistics and Philosophy* 23.4. Pp. 391–434.

Stanley, J. and Z. Gendler-Szabo (2000). "On Quantifier Domain Restriction". In: *Mind and Language* 15. Pp. 219–61.

Starr, William B. (2010). "Conditionals, Meaning and Mood". PhD thesis. New Brunswick, NJ: Rutgers University.

Steedman, M. (1996). *Surface structure and interpretation*. Vol. 30. MIT press.

Steedman, M. and J. Baldridge (2009). "Combinatory categorial grammar". In: *Non-Transformational Syntax*. Pp. 181–224.

Suarez, M. (2003). "Scientific Representation: Against Similarity and Isomorphism". In: *International Studies in the Philosophy of Science* 17. Pp. 225–244.

Szabo, Z.G. (2010). "Compositionality". In: *Stanford Encyclopedia of Philosophy*.

Tarski, A. (1944). "The semantic conception of truth: and the foundations of semantics". In: *Philosophy and phenomenological research* 4.3. Pp. 341–376.

Toker, F. (2003). *Falling Water Rising: Frank Lloyd Wright, EJ Kaufmann, and America's Most Extraordinary House*. Knopf.

Tufte, E.R. and G.M. Schmieg (1985). "The Visual Display of Quantitative Information". In: *American Journal of Physics* 53. P. 1117.

Tversky, B. (2000). "Some ways that maps and diagrams communicate". In: *Spatial Cognition II*. Pp. 72–79.

Van Fraassen, B. (2008). *Scientific Representation: Paradoxes of Perspective*. Clarendon Press.

Veltman, F. (1996). "Defaults in update semantics". In: *Journal of philosophical logic* 25.3. Pp. 221–261.

Von Frisch, K. (1967). *The Dance Language and Orientation of Bees*. Harvard University Press.

Walton, K.L. (1990). *Mimesis as Make-Believe: On the Foundations of the Representational Arts*. Harvard University Press.

Waltz, D. (1975). "Understanding line drawings of scenes with shadows". In: *The psychology of computer vision*. Citeseer.

Willats, J. (1997). *Art and representation: New principles in the analysis of pictures*. Princeton University Press.

Wittgenstein, L. (1921/1997). *Tractatus Logico-philosophicus*. Ed. by D. Pears (trans.) and B. McGuinness (trans.) Routledge.

Wolff, P.H. (1963). "Observations on the early development of smiling". In: *Determinants of infant behavior* 2. Pp. 113–138.

Wollheim, R. (1987). *Painting as an Art*. Thames and Hudson.

— (1998). "On pictorial representation". In: *Journal of Aesthetics and Art Criticism* 56.3. Pp. 217–226.

Curriculum Vitae

Gabriel Greenberg

Assistant Professor
Department of Philosophy, UCLA

Education

2000-2005   BA Philosophy and Art (double major), Phi Beta Kappa, *Magna Cum Laude*
             Yale University

2005-2011   PhD in Philosophy
             Rutgers University

Positions

2005-2007   Graduate Fellow

2007-2008   Center for Cultural Analysis Fellow

2008-2009   NSF IGERT Perceptual Science Fellow

2009-2010   Graduate Fellow / Teaching Fellow

2010-2011   NSF IGERT Perceptual Science Fellow