

Content and Target in Pictorial Representation

Gabriel Greenberg August 28, 2017 Department of Philosophy, UCLA [DRAFT]

This paper argues for a model of pictorial representation which aims to explain the relationship between pictorial content and pictorial accuracy. The model distinguishes between two fundamental representational relations: on one hand, a picture *expresses a content*; on the other, it *aims at a target scene*. A picture is *accurate* when the content it expresses fits the target scene it aims at. In addition, content in this model has two aspects: *singular content* specifies the particular individuals which a picture is *of*, and *attributive content* specifies the properties and relations which the picture ascribes to those individuals. For a picture to be accurate, both aspects must be matched in the target. I call this the *Three-Part Model* because it distinguishes between the triad of factors, singular content, attributive content, and target, which together determine pictorial accuracy. Previous work on depiction has recognized a distinction between singular and attributive content, but it has failed to recognize the additional role played by target. I will argue that the concept of target is essential in order to make sense of accuracy judgements across a range of central cases.

In Section 1, I introduce the the concepts of content and target, and situate them in relation the literature on depiction. In Section 2, I describe the Three-Part Model in outline, with attention to its component parts— picture, context, content, target, and accuracy. The next two sections defend the model’s central claims in stages. First, in Section 3, I’ll argue that a picture is only accurate or inaccurate relative to a contextually selected target scene. Then, in Section 4, I’ll argue from cases that target scenes, so construed, must be independent of pictorial contents. Section 5 goes on to shows how the Three-Part Model may be adapted to handle the phenomena of counterfactual and generic depiction. In Section 6, the conclusion, I suggest that the same two-part representational architectural extends to language, vision, and mental imagery.

1 Content and Target

Consider the following print, published in the early 1800’s, as part of a project by French scholars to document what was known to them of ancient and contemporary Egypt.

Acknowledgements: This paper has grown directly from years of conversations with friends, colleagues, teachers, and students, to whom I am truly grateful. Special thanks to Josh Armstrong, Sam Cumming, Katie Elliott, John Kulvicki, Sun-Joo Shin, and two anonymous referees for comments on the current draft, and to audiences at the 2016 Central APA, Arizona State University, University of Texas at Austin, and Washington University in St. Louis.



Picture E

Picture E is first of all a picture of the Great Sphinx from the Giza Plateau on the bank of the Nile in Giza, Egypt. It is also a picture of the Pyramid of Khafre (on the left) and the Pyramid of Khufu (just visible on the right). We may also suppose that the horse and rider pictured in front of the Sphinx were drawn from life; it is a picture of them as well. Not only is Picture E *of* various objects, it also depicts them *as* having a variety of features. Thus it depicts the Sphinx as having a certain shape (e.g. with no nose), as sitting in a certain position relative to the pyramids, as catching the light at a certain angle. It depicts the pyramid Khafre as having a certain shape, as sitting in a certain position relative to the Sphinx. And so on.

What a picture is *of*, and what it depicts its subjects *as*— these are reflections of a picture's CONTENT. Content corresponds roughly to *what's happening* in the picture, or how the picture *depicts* the world. The situation depicted by a picture in this way is categorized as *content*, and not as a mere informational trace carried by a signal, because it determines a set of substantive accuracy conditions for the picture. When these conditions are met, the picture is accurate, and when they are not, it is inaccurate: it misrepresents. If it turned out, for example, that the Sphinx did not have the shape which Picture E attributes to it, the content of the picture would be no different, but the picture would be inaccurate. Thus the content of a picture corresponds to how the picture depicts the world— independent of whether the world fits this construal.

Over the last fifty years, following Goodman (1968), scholars of depiction have distinguished

between two aspects of pictorial content, which, for the sake of standardization, I'll call SINGULAR CONTENT and ATTRIBUTIVE CONTENT, though they have gone by many names.¹ Singular content includes all the individuals a picture is *of*— in the case of Picture E, the Sphinx, the pyramids, the horse, the rider, and so on.² Attributive content includes all the properties and relations ascribed to those individuals— here, their shape, orientation, illumination, and so on, and possibly also high-level properties like *being a statue* or *being a person*. To a first approximation, when a picture depicts some object *X* as having some property *F*, then the object in the *X* position is part of the singular content, while the property in the *F* position is part of the attributive content.

In addition to this two-part factorization of content, a second assumption, less well defined, pervades contemporary philosophy of depiction, and it is this claim that I will contest below. Very roughly, this is the idea that whether a picture is an *accurate* depiction or not is determined entirely by the degree of fit between a picture's singular content and its attributive content, perhaps measured against the backdrop of the state of the actual world. And the intentions of an artist setting out to make an accurate depiction can be characterized entirely by their intentions to produce a picture whose attributive content fits its singular content. It is in this spirit that Goodman (1968, p. 38) writes, "for a picture to be faithful [\approx accurate] is simply for the object represented to have the properties that the picture in effect ascribes to it." Or Kjørup (1978, pp. 62-64), characterizing the artist's act of accurate depiction: "the producer of the picture must apply the picture to some referent and predicate something about the referent through the picture."

The few formal treatments of pictorial accuracy in the literature on depiction reify these ideas. For example, Greenberg (2013) holds that a picture *accurately depicts* a scene just in case the picture bears a certain geometric relation to the scene and in addition bears a reference relation to it. Here the "scene" plays the role of picture's accumulated singular contents, while the geometric relation encodes the picture's attributive content. Accurate depiction, then, is determined entirely by these two aspects of pictorial representation.³ The same approach shows up in literature on the semantics of maps, where scholars have relied on the semantic predicate *truth* rather than *accuracy*. Thus Casati and Varzi (1999, pp. 194-5) and Rescorla (2008, p. 180), for example, define a map's truth conditions exclusively in terms of whether the particular regions it denotes (\approx singular

¹Goodman (1968, pp. 27-8) distinguishes the *kind* of a picture from its *denotation*. Kaplan (1968, pp. 197-8) differentiates between a picture's *descriptive content* and its *genetic character*. Kjørup (1974, p. 220) marks a split between *predication* and *reference* in depiction. Schier (1986) distinguishes *iconic predication* from *iconic reference*. Hyman (2012, p. 136) separates out a picture's *sense* and its *reference*. And Greenberg (2013, p. 222) marks a distinction between a picture's *content* and its *referent*. Note that not all authors have treated (what I am calling) *singular content* as part of (what they call) "content." Instead, "content" (or cognates, like "sense") are sometimes reserved for what I call *attributive* content, while an independent semantic relation ("denotation", "reference") is posited for the expression of what I call *singular* content (Greenberg 2013; Hyman 2012). Indeed the characterization of singular content in Greenberg 2013, for example, may have more in common with my construal of *target*, a point discussed in Section 4.

²Singular content likely includes other elements of the depicted landscape, such as individual dunes and ridges, as well as *parts* of the objects already mentioned— parts of the Sphinx, pyramids, dunes, and so on. The full extent of singular content is an issue I postpone for further inquiry.

³An alternative reading of Greenberg casts the reference scene in the role of *target*, while voiding any role for singular content. I discuss views of this kind in Section 4.

content) satisfies the complex of properties it denotes (\approx attributive content).⁴

Ultimately there are substantially different ways of articulating this two-part approach to pictorial representation, and I'll eventually engage each these variants directly. But in all of this, the common claim, and the one that I reject, is that a picture's accuracy is determined solely by its singular and attributive content, along with the state of the actual world. As I will argue in Section 3 and 4, this view has an impoverished conception of the relationship between pictorial content and accuracy, and as a consequence cannot account for a wide range of critical accuracy judgements. The novelty of the Three-Part Model is the counter-claim that accuracy is a function only in part of singular and attributive content— but also a further contextually selected index, the TARGET of the picture, which is distinct from the content and cannot be assimilated to the state of the actual world. The addition of this new index is a necessity because, as I'll argue later, content and target are genuinely independent. Thus the workings of the Three-Part Model cannot be reduced to the limited ingredients traditionally thought to ground pictorial accuracy.

The terminology of “content” and “target” deployed here is adapted from Cummins' (1996) theory of mental representation, though the correspondence with the present account is only approximate.⁵ More direct antecedents of the Three-Part Model can be found in contemporary philosophy of language. There, any number of authors distinguish the *proposition* expressed by a sentence (\approx content)— containing both predicative and referential constituents— from an element relative to which the proposition is true or false (\approx target); this element is variously characterized as the *index* (Lewis 1980), *world* (Kripke 1972), or *circumstance of evaluation* (Kaplan 1989) for a sentence, in context.

In the case of pictures, this index of evaluation takes the form of a possible spatio-temporal location anchored at a particular viewpoint, what I will call a SCENE. (Scenes are later defined as viewpoint-centered worlds.) The target of a picture is that scene which it is the picture's *purpose* or *function* to be accurate at, as determined by the artist's representational intentions. According to the Three-Part Model then, for a picture to be accurate, its singular content must realize its attributive content, *in the target scene*. In effect, the *target* is the subject matter of a picture, the *content* of the picture provides a kind of comment, and it is the function of the picture to offer that content as a comment on the target. Only, in the Three-Part Model, the “comment” involves both attributive and singular elements. An artist setting out to accurately depict the world not only intends to create a picture that expresses a certain content, she also intends that this content will match a particular target.

To see how the Three-Part Model plays out in specific situations, consider again the case of

⁴These authors grapple with the added complexity that if a map does not *explicitly* attribute a property to an individual it denotes, it may thereby attribute the absence of that property to that individual. This makes the truth-rule more difficult to state, but is consistent with the gloss given above.

⁵While the ideas here are directly inspired by Cummins presentation, they also diverge from his in important ways. Cummins' notion of a target seems to be that of the *content* which a computational system is supposed to express in context; by contrast, my notion is that of the *index* relative to which content is supposed to be evaluated for accuracy.

Picture E. The content of E is apparent—it describes a certain pictorial space, populated, as I have suggested, with objects like the Sphinx, the pyramids, and the rider, which are in turn attributed properties of shape, distance, texture, illumination, and so on. Because E was designed to document a particular scene, we know that its target consists of a particular time (in the 1820's), a particular location (on the Giza Plateau in Egypt), and a particular oriented viewpoint within that location. Picture E is accurate to the extent which its target happens to instantiate the singular and attributive content which it expresses. Assuming that Picture E is in fact perfectly accurate, then, had visible features of the target been otherwise, it would have been inaccurate.

What target a picture aims at is determined by the representational purpose of the picture, itself normally a function of the artist's intentions. Life drawings are intended to accurately represent the scene before the eyes of the artist; similarly, photographs are intended to accurately represent the scene before the lens of the camera. In both cases, the target is thereby fixed as the scene immediately before the eye or lens. In other cases, the causal relation between picture and target is less direct. In drawing from memory, the artist intends her picture to be accurate at some past scene which was *previously* before her eyes; then that past scene is the picture's target. In still other cases, the fixation of target need not be mediated by the artist herself seeing the target scene at all: in a police sketch, for example, it is the witness who sees the original scene and only verbally reports it to the artist; but since the picture is intended to be accurate at the originally witnessed scene, that is the target.

It is not even necessary that the target of a picture be seen by *anyone*. Thus, I may set out to draw the Sphinx from a birds-eye-view, based on my background knowledge of the terrain; though I have never occupied that viewpoint (or talked to someone who did), it still defines the target scene for my picture. Even more extremely, targets may be located in the future, as when I set out to draw (what I expect will be) the state of the Sphinx in 100 years; then the target is located at that future time. Or, if you ask me what it *would* look like if an asteroid were to collide with the Pyramid of Khafre, and I draw a picture in response, then the target of my picture is a counterfactual scene—one that has been only imagined, but cannot be viewed. As these cases show, just as representational intentions can range across time, space, and possibility, so too can targets. Perhaps the further a target is from the artist's immediate visual context, the more likely the picture is to be inaccurate. But no matter for target, whose role is simply to fix the standard by which accuracy is measured.

In all of these cases, the factors which determine content and target are independent. A picture's content is always grounded in the spatial and chromatic organization of the picture itself; while intentions and other contextual factors play a role in determining content, they only do so via features of the picture plane. By contrast, the target of a picture is unconstrained by the format of the picture itself, and may be entirely fixed even before the artist begins work. As a consequence, there is no guarantee that a picture's content will fit its target, no matter how earnestly an

artist intends to depict accurately. When such intentions are not realized, and content and target come apart, inaccuracy results.

The same independence of content and target mean that two pictures may have *different* contents, but the *same* target: perhaps *two* artists attempted to capture the scene in front of the Sphinx that day in the 1820's. One produced Picture E; the other artist, much less skilled, produced something wildly inaccurate, call it Picture E'. Clearly, Picture E and E' have different spatial contents; but in another sense, they seem to depict the same thing; this is the sense in which have they have the same target— they aim at the same location, time, and viewpoint in Giza, before the Sphinx. One is accurate because its content fits this target, and the other inaccurate, because its content and does not fit with the same target.

In the same way, two pictures may also have the *same* content, but *different* targets. For example, suppose one person drew Picture E₁ from life at time 1, and another person drew Picture E₂ from life at time 2. Supposing that both E₁ and E₂ are qualitatively identical to the original E, and supposing that all the same individuals appeared in the scene at time 1 and time 2 (the same dunes, the same rider, the same horse), then E₁ and E₂ may have the exact same content. In the Three-Part Model, the sense in which E₁ and E₂ are representations of *different* times is not be traced to content; rather, it is solely matter of target: the target of E₂ is located at a time after the time of the target of E₁.⁶

Though content and target are distinct, they can be confused. This is due in part to an ambiguity in talk of “depiction” and its cognates. Goodman (1968, p. 22) has already observed that speaking about what a picture “depicts” or what it is “of” is ambiguous between ascriptions of attributive and singular content. To add to the interpretive possibilities, one may use the same expressions to refer to a picture’s target as well. Thus, for example, a picture may said to be *of* or to *depict* the objects in its *singular content*; this is the sense in which Picture E is *of* the Sphinx. But it is also possible to describe a picture as being *of* a scene by way of picking out its *target*; it is in this sense that Picture E is *of* a particular view of the Giza Plateau at a certain moment in the 1820's. Both correspond loosely to “what a picture depicts”. Still, for the most part, the philosophical literature which seeks to give necessary and sufficient conditions on “depiction” is directed at one or another aspect of pictorial content, and has not dealt directly with the issue of target.⁷

The target of a picture is the scene relative to which it functions to be accurate. This function, in turn, is a product of how the picture is *used* by its artist— allowing that such “use” may be

⁶Beyond the core tenets of the Three-Part Model, these facts reflect a commitment to “viewpoint neutrality”— the position that pictorial content doesn't itself contain a specific central viewpoint, but is compatible with different viewpoints. I assume viewpoint neutrality throughout; one motivation, highlighted in the section to follow, is that it makes room for the intuition that the accuracy value of the content of a given picture may vary when evaluated at different viewpoints.

⁷An exception is Kjørup (1978) who takes “depiction” to be a special pictorial act which aims at accuracy, roughly equivalent to what I call “assertoric depiction” below. My own method is to use the term “depiction” loosely, identifying and debating more specific representational relations as required, with the caveat that I normally use “depiction of” in the singular content sense.

simultaneous with the creation of the picture, and that it is generally determined by the artist's intentions, rather than some pattern of behavior. The concept of target, then, most nearly falls within the domain of the THEORY OF PICTORIAL ACTS, the pictorial analogue of *speech act theory* (Novitz 1975; Kjørup 1974; Kjørup 1978). Within the domain of pictorial acts, a core class of acts use pictures to accurately depict the world— as opposed to issuing warnings with pictures, or introducing plans with them (Kjørup 1978; Eaton 1980; Korsmeyer 1985). I shall call these ASSERTORIC pictures; it is all and only assertoric pictures to which the notion of target applies, since it is all and only these picture which function to be accurate at a target. Like linguistic assertions, assertoric pictures are used to provide information about the state of the world, and constitute one of the central uses to which pictures are applied.

By contrast, for example, IMPERATIVAL PICTURES, like Ikea instructional diagrams, architectural plans, or road-side warning signs, function to convey instructions or plans, but not to be accurate; hence they lack targets. Certain situations may or may not *satisfy* an imperatival picture, but the picture itself is neither accurate nor inaccurate.⁸ Still other kinds of images are neither assertoric nor imperatival; their central function is to please, inspire, stir the imagination, or trick the perceptual system. Examples include many doodles, some kinds of artworks and childrens' drawings, many optical illusions, and most decorative images, like wallpaper designs or fabric prints. Such images, when they are representational, do express spatial content, but there is no actual or possible scene which the artist in these cases has set out to illustrate, hence they too lack targets.⁹

The Three-Part Model, then, is not a model of *all* forms of pictorial representation, but is directed instead only at one of its prime manifestations, that of assertoric uses of depiction.¹⁰ Just as assertion plays a central role in speech act theory and the study of language generally, assertoric depiction is a core case of pictorial representation, and the natural domain for a theory of pictorial accuracy. Henceforth, when I refer to "pictures" or "pictorial representation" without further specification, I mean to restrict my attention to the case of assertoric depiction.

In all this, the theory of pictorial acts should be contrasted with the parallel fields of pictorial

⁸In the background here is an assumption of Frege's classical distinction between *force* and *content* (Frege 1879; Geach 1965). Following Kjørup (1978, p. 220), I conceive pictorial content as neutral with respect to force. The same spatial content can be used assertorically, to describe the world, or imperatively, to issue an instruction. Hence, when I make reference to the "accuracy conditions" determined by a picture's content, I mean this to be neutral as to whether the picture functions to be accurate or not. These ideas invite a fuller development of the theory of pictorial acts, but lie outside the scope of this paper.

⁹Adjudicating this question in particular cases is delicate. While some pictures undoubtedly lack targets, others (including many artworks and childrens' drawings) function to *introduce* or *specify* an imagined scenario. Such pictures are, typically, *trivially* accurate, for they *have* targets— the very scenes they served to introduce or specify. In general, it is more natural to assess pictorial contents for accuracy relative to scenes which have already been specified in prior discourse or mental activity.

¹⁰That said, the formal scaffolding of the Three-Part Model applies to *any* form of pictorial representation. Even when a picture is not assertoric, its content determines accuracy conditions, so can be assessed, more or less by fiat, for accuracy. Yet such assessments of accuracy must advert to a here-to-for unrecognized index of evaluation— what takes the role of target in assertoric depiction. I return to this theme in Section 3.


semantics and pictorial *pragmatics*. The project of pictorial semantics and pragmatics is to determine the mechanisms by which pictures, in context, are associated with their contents. Whereas a theory of pictorial acts takes the assignment of content as given, and offers an account of what the artist attempts to *do* by expressing such content.¹¹

There is little precedent for the notion of target within the philosophy of depiction, simply because most contemporary work on the topic focuses instead on the (semantic and pragmatic) question of how pictures, in context, are associated with their contents. Recent work on the resemblance theory of depiction (Hyman 2006; Abell 2009; Blumson 2014), for example, as well as on structural approaches (Kulvicki 2006; Greenberg 2013) and perceptual approaches (Lopes 1996; Hopkins 1998; Newall 2011) all take aim at the same problem: what constraints guide the mapping from picture to the content it expresses? In this paper I sidestep the important debates between resemblance, structural, and perceptual accounts by simply assuming *some* way of associating pictures with contents. The Three-Part Model itself is neutral about the nature of this mapping. Instead, the primary contribution of the model is to show how the addition of target— independent of content, hence outside the concerns of semantics and pragmatics— interacts with content to determine pictorial accuracy.

2 The Three-Part Model

Having outlined the core contribution of the Three-Part Model I now turn to provide a more systematic summary of its constituent parts, first describing a particular version of the standard conception of pictorial content, and then introducing the role of target in determinations of accuracy. For now, my goal is merely to describe the proposed model, leaving evidence and argumentation to subsequent sections. The schematic relations between the elements of the Three-Part Model are illustrated below:

¹¹The term “pragmatics” employed here is used at different times to mean different things. The conception of pragmatics as a determiner of the content conveyed or expressed by an utterance is derived from the project of *Gricean pragmatics* (Grice 1975; Lepore and Stone 2015). By contrast, more recent approaches to *formal pragmatics* overlap directly with speech act theory (e.g. Stalnaker 1978; Roberts 1996; Murray and Starr 2017). In this restricted sense I am happy to call the Three-Part Model a “pragmatic” one.



The Three-Part Model

To begin, PICTURES, as I shall use the term, are a species of public signal whose form and content are in some sense distinctively visual. Canonical examples include life drawings, engineering plans, architectural renderings, courtroom sketches, medical illustrations, and photographs. Pictures include many works of visual *art*, but are not limited to them. Officially, I will treat pictures as abstract types, specified in terms of their 2-dimensional geometrical and chromatic organization. Particular picture tokens correspond to the instantiation of pictures in context, though for convenience, I'll often call these "pictures" as well.

A picture is created in a particular historical, causal, and psychological setting, which I will call its CONTEXT. In the Three-Part Model, context plays two roles, just as it does in a traditional model of linguistic meaning (Kaplan 1989). On one hand, context determines the factors necessary for a picture to express its content; this is what MacFarlane (2009) calls *indexical context-sensitivity*. On the other hand, context fixes the target which the picture aims at, thereby securing the scene relative to which the picture is accurate or inaccurate; such context-sensitivity effects the accuracy value of a picture, but not its content; this is what MacFarlane calls *non-indexical context-sensitivity*. Like sentences, pictures exhibit both kinds of sensitivity to context.

Starting with content, I'll officially say that a picture EXPRESSES its content relative to a context. In determining a picture's content, the picture's own spatial and chromatic organization plays the primary role; no content, attributive or singular, is expressed except through some particular region and pattern of marks on the picture. Still, context plays an important role in guiding pictorial expression. While the role of context in determining content isn't the focus of this paper, two contributions of context are worth noting. First, a central task for context is the selection of an operative SYSTEM OF DEPICTION, the pictorial analogue of a language. Systems of depiction play an essential role in associating the geometrical and chromatic surface features of pictures with elements of attributive content (Giardino and Greenberg 2015). Second, as Kaplan (1968, pp. 197-8) and Lopes (1996, ch. 5) have argued, a picture's singular content is largely a function of the causal

context in which the picture was created. In this way, pictorial singular content seems to be fixed in a manner at least analogous to that envisioned by the causal theory of reference for names.¹²

A picture's content determines the conditions under which the picture is accurate, just as the contents of sentences determine truth conditions. Here and throughout, I set aside hypothesized aspects of content that are not directly reflected in accuracy conditions (understood as sets of centered-worlds), such as modes of presentation or other hyper-intensional meanings. Peacocke (1992) and Burge (1991; 2014), among others, have argued for such elements in visual perception. Whether these conclusions carry over to depiction, I leave to future investigation, though I suspect they will.

While pictorial content is reflected in facts about what objects a picture is of, and what properties it depicts these objects as having, no single such fact captures the full content of a picture. A picture never merely depicts a plant as having green leaves, for instance; it also depicts it as having a certain shape, orientation, location, all relative to a general viewpoint. Thus it has proven natural to conceive of the contents of a picture as a PICTORIAL SPACE centered at a general viewpoint: a three-dimensional region populated with individuals, properties, and relations, where the locations of these elements are always specified relative to a central perspective position. (Panofsky 1924 (1991); Howell 1974; Rogers 1995; Koenderink and Doorn 2003) Intuitively, a space like this may be accurate or inaccurate at a target scene to the extent that the scene is a realization of that space. The characterization of pictorial space offered here is inspired in part by Peacocke's (1992, ch. 3) concept of *scenario content*, which formalizes the idea of perceptual content as a type of viewpoint-centered space. Such regimentation won't be necessary for my purposes, but readers may find in scenario content a useful way to precisify the structure of pictorial space described above.

In general, on the Three-Part Model, content is neutral with respect to any particular world, time, or viewpoint. Yes, pictures express pictorial spaces, and pictorial spaces are accurate *at* particular worlds, times, and viewpoints. But such spaces themselves contains no particular information about what world, time, or viewpoint they are to be evaluated at. Thus, what's depicted by a given picture— its content— may be accurate at some worlds and not at others; accurate at some times and not at others; and even accurate relative to some viewpoints, and not relative to others. If the content itself contained these elements, such evaluations would be impossible.

In characterizing pictorial space, I invoked the concept of a GENERAL viewpoint, or viewpoint *type*, in contradistinction with the notion of a PARTICULAR viewpoint, or viewpoint *token*. A picture depicts objects and properties relative to a *general* viewpoint, because the picture's content cannot include any *particular* token viewpoint. At the same time, a given pictorial space may be *realized* by some particular viewpoint in a particular space. So while a pictorial content is general with

¹²See Goodman (1968, ch. 1) Kjørup (1978, p. 57) and Hyman (2012, pp. 138-40) for amplification of the theme that attributive content is determined by general and systematic interpretive mechanisms, while singular content traces more directly back to local and causal features of the context.

respect to viewpoint, content is accurate or inaccurate relative to a particular viewpoint. In what follows, I normally reserve the term “viewpoint” for particular token viewpoints, disambiguating when necessary.

For many purposes, it is useful to model pictorial content within the framework of possible-world semantics.¹³ Such formalism isn’t essential to the Three-Part Model, but helps to fix ideas. To this end, I introduce the concept of a “viewpoint-centered world” or SCENE, defined as the pair of a possible world and a (particular) viewpoint. A VIEWPOINT here is conceived of as a particular oriented location in space and time, and carries no implication of a real or metaphorical viewer. Then, following Ross (1997), Blumson (2009), and Greenberg (2016) we may model a picture’s content as a *set* of scenes. Formally, such a set may be thought as a relation between worlds and viewpoints: the content of a picture imposes a simultaneous constraint on the way the world is relative to a viewpoint, and the location of that viewpoint relative to the world.

This model of a picture’s content fits smoothly with the characterization of a picture’s content as a pictorial space centered on a general viewpoint, populated with objects and properties. This is because every such space determines a set of scenes— those pairs of worlds and particular viewpoints for which, when that particular viewpoint occupies the center of the pictorial space, the world is an instantiation of that space. Of course, an arbitrarily selected set of scenes won’t correspond to a pictorial space in this way, as its members may be entirely disjoint. But those sets of scenes which *are* the contents of pictures will all coincide precisely insofar as the picture’s represented space is determinate, and diverge insofar as it is indeterminate.

As advertised, pictorial content can be further subdivided into singular and attributive aspects. SINGULAR CONTENT is that aspect that is composed of (or determines) the particular individuals which a picture is *of*. ATTRIBUTIVE CONTENT is that aspect of a picture’s content that is composed of (or intrinsically determines) an array of properties and relations. Taken as a whole, a picture’s content ascribes the properties and relations in its attributive content to the objects in its singular content. In terms of pictorial spaces, the objects that populate a given pictorial space belong to a picture’s singular content, while the properties and relations located in that space, along with the structure of the space itself, belong to its attributive content.

In fact, singular content belongs to the more general category of OBJECTUAL CONTENT; such content may be an INDEFINITE, as is the case when a picture merely depicts *some* cube as being located in a given direction— or it may be SINGULAR, as when a picture depicts a *particular* cube as being located in a given direction. While indefinite objectual content is both theoretically important and common in depiction, it is the treatment of singular objectual content— in particular,

¹³Inevitably, this approach also has its drawbacks. Possible-worlds models of content, for example, are generally ill-equipped to give a plausible account of the contents of *impossible* images, since they seem to imply that all impossible images have the same null content. (But see Blumson 2009 for an attempt to rectify this.) Though I make use of the possible-worlds framework for the elucidation it offers, the flat-footed treatment of impossible content should be understood as a defect of this particular model of pictorial content, and not of the Three-Part Model itself.

the contrast between singular content and target— which is distinctive of the Three-Part Model, and it will be the focus of the discussion to follow.

Crucially, in the Three-Part Model, singular content is independent of the other two representational relata, in the sense that it cannot be derived from facts about attributive content or target alone. (I'll argue for this conclusion in Section 4.) Since the objects picked out by singular content cannot be derived from other aspects of content, it is natural to conceive of these objects as being genuine constituents of the content itself.¹⁴ In the language of Kaplan (1989), we might say that pictures are devices of *direct reference*, for the objects which they are *of* are not merely specified by an intermediary description (e.g. the attributive content), but are themselves parts of the content. Of course, pictures also express attributive content, setting them apart from standard cases of directly referring terms in language. In this sense, pictorial content is both singular *and* attributive, making pictures both directly referential *and* descriptive.¹⁵

I turn finally to elements of the Three-Part Model which distinguish it from previous work on depiction, having to do with pictorial accuracy. Thus far I have informally described pictorial content as “fitting” or “matching” a target scene; to fix ideas, I'll officially say that a picture's content is ACCURATE AT a scene. The idea is that the content of a picture can be evaluated as accurate or inaccurate at a scene— much the way linguistic propositions are thought to be evaluable as true or false relative to a world or index of evaluation. Since contents are modeled as *sets* of scenes, a content is accurate at a scene if and only if the scene is a *member of* that content. Note that, for now, it is *contents* which are accurate at scenes. I return to the question of what makes a *picture* accurate in a context shortly.

A picture's content is accurate at a scene when *both* the singular and attributive components of the content correspond in the appropriate way with the objects, properties, and relations which populate that scene. Thus, in the Three-Part Model, accuracy at a scene requires (i) that the objects in the singular content actually exist in the scene; and (ii) that the objects so depicted have the properties and stand in the relations, in the scene, that they are associated with by the picture's attributive content.

An important wrinkle here is that accuracy, unlike truth, comes in degrees; what degree of accuracy defines a pictorial content's accuracy conditions? Here I help myself to the notion of maximal or PERFECT accuracy: the content of a picture corresponds to the set of scenes relative to which that content is perfectly accurate. Henceforth, by “accuracy” I mean perfect accuracy (or very near it); by “inaccuracy” I mean less than perfect accuracy. Providing an account of graded

¹⁴Alternatively, one might think of singular content as being made up of meaningful parts which intrinsically *determine* particular objects, in the manner of “*de re* senses” (McDowell 1984). For the sake of simplicity, I'll usually just talk of objects being “in” a picture's singular content, but officially I remain neutral between the constituent and determination formulations.

¹⁵Theorists of depiction have often assumed that the relationship between a picture and the objects in its singular content is one of *pictorial reference*, analogous in important respects to reference in language and thought. (See, e.g. Schier 1986, Lopes 1996.) Whether or not these amount to the *same* relation, there are clear commonalities.

accuracy is left to future investigation.¹⁶

Note that accuracy, in the sense I intend, does not imply realism or closeness to reality. The contents of a black and white drawing and of a color painting may each be perfectly accurate, albeit expressed in different systems, though the experiences these pictures elicit obviously differ in the proximity to normal perceptual experience. A full scale, working model might be closer to reality, in some sense, than a technical drawing, but the content of each may be perfectly accurate. What matters for accuracy is the absence of misrepresentation— not the quantity or type of information represented.¹⁷ A picture's content is accurate at a scene when, to the extent that it represents things as being a certain way in that scene, that is the way those things are. This notion of accuracy is not the exclusive denotation of the word "accuracy" in colloquial English, but it is arguably one of them.

Thus far I've discussed the ways in which a picture's content may be evaluated for accuracy relative to an arbitrary scene. The concept of TARGET serves to isolate one such scene as having special status: a picture's target corresponds to the scene relative to which the picture is *intended* or *functions* to be accurate. I'll say that a picture AIMS AT its target scene. A picture which is accurate of its *target* has achieved an important standard of representational success not conferred by mere accuracy at *some* scene. If I set out to draw a picture of my office from the viewpoint of the doorway, thereby fixing my target, and if the picture is accurate at this scene, it succeeds; by contrast, if it is inaccurate at this scene, but happens to be accurate of some other scene, say in some other office, the picture has not succeeded as a representation.

Using the concept of target, I can finally define accuracy for pictures in context. The concept of accuracy for *pictures* is derived from that of accuracy for *contents*, much the way the concept of truth for utterances is related to that of truth for propositions. They are distinct notions, but clearly inter-related. In context, a picture both expresses a content and aims at a target. Given a context, when the content a picture expresses (in that context) is accurate at the target scene the picture aims at (in that context), then we may say that picture itself is ACCURATE IN CONTEXT. Combining this formulation with the definition of accuracy at a scene for content, we may derive the following statement of accuracy conditions for pictures. For any assertoric picture *P* and context *C*:

¹⁶In this characterization we risk losing sight of that aspect of pictorial content which, relative to a scene, determines intermediate degrees of accuracy. Still, it should be noted that those scenes at which the content of a picture is perfectly accurate are, intuitively, exactly those which reflect that content. For example, given a drawing of my plant, it is part of the content of the drawing that it has precisely *that* shape of leaf— not big spiky leaves, not even leaves which are slightly more spiky, nor leaves which are slightly less spiky. For any leaf-shape not perfectly accurately represented by the picture, that shape is not part of the content of the picture. For these reasons, perfect accuracy rather than graded accuracy seems to play the foundational role in pictorial semantics.

¹⁷This does not mean that pictures may be automatically accurate in virtue of being blank or omitting marks— depending on the operative system of depiction, blankness itself carries content. (See Rescorla (2008) and Camp (2007) for discussion.) The requirement for a picture's accuracy is that its *content* not misrepresent its target.

P is accurate in *C* iff the attributive content expressed by *P* in *C*
is instantiated by the singular content expressed by *P* in *C*
in the target scene selected by *C*.

It is this statement which is in some sense the hallmark of the Three-Part Model. In subsequent sections I'll show how alternative approaches, derived from the traditional two-part conception of pictorial accuracy, depart from this definition.

In Section 1, I characterized the target of a picture as the scene relative to which it is the picture's function to be accurate or inaccurate.¹⁸ The notion of *function* here is broad. For many kinds of pictures, it is fixed by the artist's intentions or purposes.¹⁹ I've already discussed the ways in which such intentions may pick out characteristically different scenes in life drawing, drawing from memory, or drawing from description. In the case of mechanically produced images, like digital photographs, it may be the function of the picture-taking device, rather than the intentions of the artist, which fix the target.²⁰ Typically, in these cases, the target is the scene before the lens of a camera. In Section 3, I'll discuss further cases in which pictures aim at *counterfactual* and *generic* targets.

For a scene to be a picture's target, it need not be the case that the picture functions to be *accurate* at that scene, merely that it functions to be *accurate or inaccurate*. A pictorial liar may intend to misrepresent my office as larger than it actually is. This makes my office the target of the picture, even if the artist's goals are satisfied only when the picture is inaccurate. Instead, it seems that it is a norm for sincerity in pictorial communication— but not a condition on pictorial assertion— that a picture's content be accurate at its target.²¹

It is important to distinguish the target, conceived as the intended index of evaluation for a picture, from *the content the artist intended to express* with the picture. Both are, in some sense, ideals of picture production— one is an ideal of expression, the other of evaluation— but the two are independent. Consider an artist with prodigious artistic skill, but whose memory is unreliable, and who sets out to draw a particular scene from memory. Because of her artistic skill, the content her picture expresses is exactly what she intended to express by it; but because of her faulty memory, that content may be highly inaccurate at the scene she intended to draw. Thus the target of a

¹⁸The concepts of representational function and success here derive from Burge's (2010, pp. 308-315) discussion of representational function in perception.

¹⁹Where representational intentions are confused (in the style of examples in Kripke 1977), there may be more than one scene at which the artist intends her picture to be accurate. Such cases may require the recognition of two targets, or none; I leave this matter to further investigation.

²⁰Of course, cameras are artifacts, so their function itself is the product of human intention. But the relevant intentions in this case are directed (*de re*) to the picture making device, and not (*de re*) to each picture it produces. By contrast, the artist of a drawing or painting has representational intentions about the individual work she is producing. Still, as Megan Hyska pointed out to me, cameras themselves can be used for different functions— e.g. as devices of duplication (as with a Xerox) or as devices of photography; selecting one such function might indeed turn on the *de re* intentions of the device's controller with respect to its output.

²¹Kjørup (1978) articulates a cognate norm of "seriousness" for pictorial speech acts, though he does not allow that an intentionally lying artist could be using a picture in the same assertoric sense as a sincere one.

picture and the intended content of a picture come apart in characteristic ways. Given an artist's expressive intention, whether a picture expresses the content it was intended to is largely a matter of artistic ability, and is independent of how the world is. But given a picture's content, whether that content is accurate at its target is wholly a matter of how the world is (at the target), and is independent of artistic skill.

Thus when an artist sets out to create an (assertoric) drawing, she comes to the table with at least two sets of intentions. One is an expressive intention; this is the intention to create a picture which expresses a particular content, regardless of its accuracy. The other is an "aiming" intention; this is the intention to create a picture whose content is accurate at a particular scene, whatever that content happens to be. Assertoric drawing always involves both kinds of intentions, and both play a role in the psychological explanation of a given picture's production.

This concludes my overview of the Three-Part Model. In the next two sections, I'll present the evidence and arguments that favor it over alternative views. Section 3 is devoted to the claim that accurate depiction depends on a variable target scene; Section 4 supports the further claim that this target scene is independent of a picture's content.

3 Target and Accuracy

What distinguishes the Three-Part Model is its recognition that in order to determine pictorial accuracy, a picture's content must be compared to a contextually selected target scene. This stands in contrast with the view that singular and attributive content alone, together with the state of the actual world, is sufficient to determine accuracy. In this section and the next, I argue for the Three-Part Model's key claim in stages. First, in this section, I argue that whether a picture is accurate depends in part on a contextually selected scene. In the next section, I'll argue that such a scene cannot be part of, or derived from, a picture's content. As will become clear, the arguments for the first claim flow from relatively straightforward observations about variation in judgements of pictorial accuracy, while the second will demand a more involved examination of alternative proposals and specific cases.

For ease of reference, I set out both theses here, assuming a restriction to cases of assertoric depiction:

Thesis 1. A picture's accuracy depends in part on a contextually selected target scene.

Thesis 2. The target scene for a picture is independent of its content.

Together, Theses 1 and 2 support the Three-Part Model's definition of accuracy advanced in the previous section. That is, for any assertoric picture P and context C :

P is accurate in C iff the attributive content expressed by P in C
is instantiated by the singular content expressed by P in C

in the target scene selected by *C*.

Thesis 1 establishes the relevance of a target scene, or world and viewpoint, to determinations of accuracy, but falls short of claiming that this index is not a part of, or derived from, a picture's content (or vice versa). Thesis 2 goes on to separate the target scene, so construed, from content. Putting these together yields a definition of accuracy which necessarily adverts to both content and target as separate parameters, as in the formulation above.

Turning now to Thesis 1: Thesis 1 distinguishes the Three-Part Model from a prominent interpretation two-part approach to pictorial accuracy to be found in the traditional literature on depiction. According to this view, a picture is accurate if and only if its singular content instantiates the properties and relations ascribed by its attributive content in the actual world. Such theories pursue what I will call the ACTUALITY APPROACH, articulated by the accuracy-conditions below.²²

P is accurate in *C* iff the attributive content expressed by *P*
is instantiated by the singular content expressed by *P*
in the actual world (of *C*).

This proposal itself can be given different metaphysical interpretations. On one, it requires that the singular content be *actually instantiated* by the attributive content— where the relation of actual instantiation is a basic concept. On another, it requires that singular content be instantiated by the attributive content *at the actual world*— here, instantiation itself is relativized to worlds, but only the actual world is allowed to play the relevant role. On still another, the singular content itself is rigidified to the actual world, picking out world-bound tokens, and the definition of accuracy merely requires that these tokens instantiate the relevant properties in the attributive content.

On all interpretations, the actuality approach contrasts with Thesis 1. For one, the actuality approach fixedly ties accuracy to the state of the actual world, whereas according to Thesis 1, the world of the target must be allowed to vary by artist intention— this is the sense in which the thesis requires that the target be *contextually selected*. Second, the actuality approach makes no room for variation in viewpoint in the index of evaluation, but Thesis 1 makes this an essential feature of the target, in virtue of characterizing it as a *scene* rather than a world.

The first step in arguing for Thesis 1, then, is simply establishing that assessments of pictorial accuracy depend on an index consisting of a world and temporally-located viewpoint; from there I'll make the case that this index is selected by context in the manner of target.

²²In the final clause of the definition below, I employ the ambiguous phrase “the actual world (of *C*)”; following Kaplan (1989), the most general way to understand this clause is as *the world of the context C*, rather than tying the condition rigidly to the actual world. This means that we may reasonably ask after the accuracy of pictures created in hypothetical situations. This still restricts variation in the index of evaluation considerably, since, relative to a world in which the picture is created, only that world can be the index. For the Three-Part Model, even once the world of creation is fixed, the world of evaluation is allowed to vary according to the artist's intentions.

A basic motivation for the scene dependence of pictorial accuracy is the observation that it is possible, simply by fiat, to evaluate the content of a given picture relative to arbitrary scenes. Thus I may ask of Picture E whether what it depicts—its content—*would* be accurate at various alternative points in time and space. Intuitively, there is variation here. Variation with respect to *world*: as things actually are, the content of the picture is accurate; but in a counterfactual situation in which the Pyramid of Khafre had never been built, the same content would be *inaccurate*. Variation with respect to *time*: at some points in time the Sphinx lacked a nose, and the content of Picture E could be accurate at these times; at other times, it had a nose, and the content would be inaccurate at those. And variation with respect to *viewpoint*: though the content of Picture E may be accurate from some viewpoints, at others, where the Pyramid of Khafre is not visible for instance, that content would not be accurate. The most direct explanation for these judgements is that content alone does not determine an accuracy value. Instead, the accuracy value of pictorial content can vary relative to world, time, and viewpoint—in other words, relative to a scene.²³

The same lesson emerges from variation in the factors which determine whether pictures produced in different kinds of context are accurate. Consider again the difference between actual and hypothetical depiction. On one hand, if you ask me to draw a picture of Saturn in orbit, and I produce a picture in response, whether my picture is accurate would depend on the state of the actual world. On the other, if you ask me to draw what it *would* look like if Saturn and Jupiter collided, whether my picture is accurate depends on the state of a counterfactually specified situation, and not directly on the actual world. The same type of variation emerges for viewpoint: if you are asked to draw a picture of what is in front of you, the accuracy of the picture depends on your current viewpoint; if you attempt to draw a familiar landmark from a bird's eye view, the accuracy of the picture now depends on the imagined viewpoint suspended in air. And parallel reasoning applies to variation in times. Thus, even when target scenes are not explicitly specified as indices of evaluation, different contexts have the effect of making pictorial accuracy depend on different worlds and viewpoints, just as Thesis 1 predicts.

This conclusion, that picture can only be assigned an accuracy value relative to an viewpoint-centered world, is all but inevitable from the perspective of possible world semantics. For, according to a standard possible-worlds framework, objects do not have their properties absolutely, but only relative to a possible world. Since the Sphinx could have had a different size and shape than it actually does, its properties vary by world. Thus pictorial accuracy must be relativized at least to world, for the properties it attributes to the objects in its singular content are only instantiated at some worlds, and not others.

²³These considerations can be used to argue for the strong claim that the target is not part of the content, since, as I've described the situation, the cases hold fixed content while varying the scene of evaluation. An alternative reading, however, takes "what a picture depicts" to be a *component* of content, but not the entirety; this allows that the observed variation is in fact variation in content. For now, I'll settle for the common conclusion that evaluations of accuracy depends in one way or another on a variable scene.

The same types of considerations carry over to times and viewpoints. What properties an object instantiates varies by time; the Sphinx had a nose at one time, and not at another. Thus the accuracy of Picture E's content can only be assessed relative to a time. In the same way, the distinctively perspectival properties and relations attributed by a picture are only instantiated relative to a viewpoint. For example, the Sphinx occludes the Pyramid of Khafre (and not Khufu) relative to some viewpoints; the reverse is true relative to others. Since perspectival relations like distance, visual angle, relative position, and occlusion are expressed by pictures, their contents again can only be evaluated for accuracy relative to a viewpoint. Putting these elements together—world, time, and viewpoint— one arrives once again at the notion target scene which I have in mind.

Such considerations support the claim that pictorial contents can only be accurate or inaccurate relative to a world and viewpoint. But what determines whether a given picture, in context, is accurate *simpliciter*? The view of the Three-Part Model is that this index is selected by context, in a manner heavily dependent on (or determined by) the artist's intentions.

In some respects, this follows the conception of the relation between context, content, and circumstance of evaluation traditional from the philosophy of language. Just as in linguistic semantics, pictures are tokened in context; relative to a context, they express a content; and relative to context, they are assigned a circumstance of evaluation; the picture itself is accurate when its content is accurate at its context-selected circumstance. The chief difference between the Three-Part Model and this standard configuration has to do with the *way* in which the circumstance of evaluation (target) is determined by context. In a traditional linguistic semantics, it is assumed that circumstance of evaluation is simply a world (or a world and a time), and that this world is *the same as* the world (and time) of the context of utterance. Thus the determination of circumstance of evaluation (for unembedded sentences) is covered by a simple default rule which privileges the actual world over others: ϕ is true in C iff the proposition expressed by ϕ is true at the world of C .

But this simple rule cannot hold for the pictorial case, for two reasons. First, the relevant world of evaluation is not always the actual world. We saw this already for cases of hypothetical depiction: in some cases accuracy doesn't depend on the state of the *actual* world, but rather the state of some *counterfactual* scenario. In general, as the cases of future, past, and hypothetical depiction illustrate, the world and time of evaluation for a picture cannot be identified with the world and time of the context of creation; instead, they correlate directly with the representational intentions of the artist.

Second, there is the matter of viewpoint. There is no privileged viewpoint relative to which pictures could be evaluated in the same way that the actual world might be considered a privileged world of evaluation. Viewpoint must either be selected for by context, or quantified over in the definition of accuracy. But consideration of cases suggests that it must be the former. Suppose

I set out to draw the Sphinx from an angle where it in fact occludes the Pyramid of Khafre (as in Picture E). But what results, Picture F, is a picture in which the Sphinx does *not* occlude the Pyramid of Khafre. We may also suppose that picture F is accurate from some *other* viewpoint besides the one intended. While there may be a secondary sense in which Picture F correctly depicts the target scene, it is not in the first place a successful representation, for it fails to meet the representational standard set for it by the artist. It is not accurate because its content is not accurate at the intended viewpoint. Generalizing, it seems that, like the world of evaluation, the viewpoint of evaluation is also picked out by context largely through the intentions of the artist.

Putting these conclusions together, it appears that the accuracy of a picture in context depends on an contextually-selected index scene. This is Thesis 1. In the next section, I complete the argument for the Three-Part Model by showing that this target scene neither includes nor is part of a picture's content.

4 Target and Singular Content

In the last section, I argued that pictorial accuracy depends on a target scene. This in itself leaves open the question of what relation holds between the target scene and content. In this section, I argue for Thesis 2, the claim that target is independent of singular content and attributive content. Just as Thesis 1 contrasted with the actuality approach to pictorial accuracy, Thesis 2 contrasts with a different set of strategies for reviving the traditional, two-part conception of accurate depiction. These views accept that a contextually selected scene plays an essential role in determining accuracy, but attempt either to assimilate that scene into singular content (and do away with target), or assimilate the objects a picture is *of* into target (and do away with singular content). Either way, these theories countenance only *two* representational relata as the determiners of a picture's accuracy, as against the central posit of the Three-Part Model.

Such two-part views are fueled by the fact that singular content and target are easily confused. For, according to the Three-Part Model, pictures now have two relata which are both in some sense "representational" and both in some sense "singular"—singular content on one hand, and target on the other. This conceptual homophony between singular content and target is, I believe, a primary reason that the latter has not been clearly distinguished from the former.

In prising singular content from target, I clarify the contrast between the Three-Part Model and previous accounts of depiction. At the outset, I characterized the depiction literature as defining accuracy exclusively in terms of singular and attributive content, to which the Three-Part Model added target. A more nuanced description might be that the literature recognizes a distinction between attributive content and *some* singular element, but conflates aspects of singular content and target. For instance, this singular element is supposed to reflect which particular objects a picture is of (like singular content), but it is also often thought that the singular element, when combined

with attributive content, determines an accuracy value (like target). For these reasons, a more careful statement of this contribution of the Three-Part Model is that, unlike previous accounts, it not only recognizes a role for target (as argued in the previous section), but distinguishes it from the role played by singular content.

This position can be characterized by a pair of sub-theses, each of which is defended separately below. Together they entail Thesis 2 set out above.

Thesis 2A. A picture's singular content is not determined by its attributive content or target.

Thesis 2B. A picture's target is not determined by its attributive content or singular content.

Here, saying that one representational relatum is not determined by the others does not imply that it is unconstrained by them. The point for Thesis 2A is to hold that pictures have genuine singular content, in a sense which is not merely derivative in some way of the picture's target and attributive content. Likewise for Thesis 2B, the point is that pictures have genuine targets, which are not merely derivative of singular and attributive content. If both theses are right, then singular content and target can vary independently of one another. This fact will play a central role in the argument for the Three-Part Model developed below.

In what follows I'll argue for Thesis 2 by appeal to cases which appear to dissociate singular content and target. The Three-Part Model is well-suited to handle such cases, since it treats these as independent relata. I'll show that theories which fail to treat singular content and target as independent elements cannot explain key semantic features of these cases.

4.1 Thesis 2A: Singular Content

I first argue for for the Three-Part Model's Thesis 2A, that pictures have singular contents, independent of their attributive contents or targets.

I begin with the assumption that pictures are *of* or *about* particular objects. It is in this sense that Picture E is a depiction *of* the Sphinx, or *depicts* the Sphinx *as* having certain shape and texture properties. As I've noted, in the Three-Part Model, facts about which objects a picture is *of* are direct reflections of the fact that those objects are constituents of the picture's singular content.

The skeptical alternative is that objects play no role in content, and pictorial content is instead PURELY ATTRIBUTIVE: assertoric pictures merely attribute properties to their targets, but they do not express singular content. Defenders of a purely attributive approach to pictorial content need not deny that pictures are *of* or *about* particular objects, but they must offer an alternative method for accounting for these facts. Two strategies in particular present themselves. The first attempts to derive *of*-facts from features of the picture's target scene, while the second attempts to derive such facts primarily from features of the picture's attributive content.

The SYNECDOCHIC STRATEGY trades on the idea that talk of what a picture is *of* is just a way

of picking out its *target* via its *visible parts*.²⁴ Thus, roughly, to claim that Picture E is a depiction of *the Sphinx* is just to point out that E's target is a particular scene in 1820's Egypt, and the Sphinx is visible in that scene. Schematically, for any picture *P* and object *O*:

- P* is of *O* iff (i) *P* aims at a target scene *S*;
(ii) *O* is a part of *S* which is visible from the viewpoint of *S*.

The concept of *visibility* here is not quite the familiar one of perceivability, but a geometrical adaptation for which no viewer or optics need be involved. An object *O* is *VISIBLE* in a scene *S*, whose viewpoint is *V*, just in case, in a geometrical projection of *S* relative to *V*, some parts of *O* are projected onto the picture plane. Thus, if an object is wholly occluded by some surface (relative to *V*) it won't project to the picture plane, so won't be visible. Otherwise, if it is still within the picture frame, it will count as visible.²⁵

The synecdochic strategy has considerable appeal, for it delivers correct verdicts for large swaths of cases, including all cases in which pictures are fully accurate, and all cases where the inaccuracy is limited to attributive inaccuracy. In this way, perhaps all facts about what objects a picture depicts can be reduced to facts about what target it aims at.

A second approach to a purely attributive view of pictorial content seeks to derive *of*-facts primarily from a picture's attributive content (as opposed to its target). This *DESCRIPTIVIST STRATEGY* was originally outlined and criticized by Lopes (1996, pp. 93-97). The idea is to think of pictures as akin to names on a descriptivist analysis: in the first instance they express sets of properties, and in the second, derivatively pick out individuals which satisfy these properties (Kripke 1972). In particular, this strategy holds that the attributive content of a picture, together with its target, pick out various objects, and these are the objects which the picture is of. Schematically:²⁶

- P* is of *O* iff (i) some region of *P* expresses the attributive content *A*;
(ii) *P* aims at target scene *S*;
(iii) *A* and *S* uniquely specify *O*.

In clause (iii), the manner in which the content *A* and the scene *S* uniquely specify *O* may be fleshed out in different ways. But the general idea is easily illustrated by the case of Picture E: E's target is culled from the actual world; the picture's unusual attributive content is satisfied by only one thing in the actual world—the Sphinx; hence, the picture is *of* the Sphinx. This approach differs from the synecdochic strategy because the object in question need not be *visible* in the target

²⁴This is arguably the tack taken by Greenberg (2013, p. 221) when he writes: "Informally, I will often talk of pictures representing objects, but only insofar as those objects are parts of scenes." This is consistent with the fact that he describes pictorial "content" in purely attributive terms, saving "reference" for scenes.

²⁵The concept of the geometrical projection at work here will vary by system of depiction; some systems rely on linear perspective projection, others isometric projection, and so on (Greenberg 2013; Giardino and Greenberg 2015). As a consequence, what counts as visible may also vary with system of depiction.

²⁶Clause (i) refers to "some region" of *P*, rather than *P* itself, since a given picture may be of multiple objects, corresponding to different regions of the picture plane.

scene, so long as it is appropriately picked out by the attributive content.²⁷ Thus, according to the descriptivist, even if the Sphinx *weren't* visible in the target scene of Picture E— located perhaps in another part of the world— it could figure in the singular content of E, so long as it uniquely satisfied E's attributive content.

Yet I hold that neither the synecdochic, nor descriptivist, nor any other purely attributive strategy can successfully account for facts about what pictures are of. My approach is to identify cases where the picture is clearly *of* some object, but this fact cannot be derived from facts about the target scene. The only way to account for what a picture is of, in such cases, is to independently posit objects in the singular content, as per Thesis 1.

But cases like this are not trivial to come by. Normal cases of accurate depiction won't do, because these are situations where the objects a picture is of are present and visible in the target. Both the synecdochic and descriptivist strategy were designed for such situations. In addition, the most familiar forms of inaccurate depiction also fail to make the point. These are cases where the objects a picture is of are visible in the target scene, but the picture misrepresents these objects, say, with respect to color, or shape, or position. In these situations, so long as the misrepresented objects are visible in the target, they will be accounted for by the synecdochic strategy; and so long as the misrepresentation is not too great, they may be explained by the descriptivist strategy as well. Thus, for a great range of cases, the Three-Part Model does not make obviously different predictions than these purely attributive strategies.


Instead, the most compelling argument for the independence of singular content derives from cases where a picture is *of* some object, but that object is not present in the target scene. Such cases effectively dissociate singular content and target to such an extreme degree that attempts to derive one from the other are undermined.

In what follows, I'll describe a case like this, involving an artist who undergoes a hallucination while drawing from life. Note that hallucination itself is inessential to the argument; intent to misrepresent, false background beliefs, or simple fantasy could deliver the same result. It is also inessential that the picture be drawn from life; drawing from memory or description would work just as well. Here is the case:

²⁷In addition to the kinds of objections pursued below, I have serious misgivings about the descriptivist strategy which echo Kripke's (1972) own objections to descriptivism in the nominal realm. It seems to me that it is possible for a picture to be *of* a given object but misrepresent it so thoroughly that virtually any other (or no) object satisfies the attributive content of the picture. (I have in mind the likes of children's scribbles, or drawings made while blindfolded. See Greenberg (2013, p. 225) for such an example.) If this is true, there would seem to be no recourse for the descriptivist. Yet my claim that such cases are possible is contentious, and I don't intend to argue the point here. For opposing views, see e.g. Hopkins (1998, p. 30) and Abell (2009, p. 212).

Object Hallucination

Yesterday, I see for the first time a cube, named Cubey, sitting in my garden. (Call the scene defined by this world, time, and viewpoint *Garden*.) Today my desk is empty. I sit down at my desk, and set out to draw what I see. (Call this scene *Empty-Desk*.) But at that point I suffer from a partial hallucination, in which it seems to me that Cubey is sitting on the desk before me. I proceed to draw the situation I take myself to be seeing, producing Picture A:



Picture A

The following two facts seem to be implied by my description of the case, and I assume them in what follows:

Assumption 1: Picture A is of Cubey.

Assumption 2: Picture A is not accurate.

Assumption 1 corresponds to the intuition that the picture is of Cubey, and that it depicts Cubey *as* sitting on the desk. Note that this is *not* the theoretical assumption that Cubey is part of the picture's singular content, only that the picture is of Cubey; the stronger claim I will have to argue for. Assumption 2 reflects the fact that although the picture depicts a particular object as located at a certain position, that object is not in fact located at that position— hence the picture must be inaccurate.

The Three-Part Model straightforwardly accommodates these two assumptions. For the first: Picture A is of Cubey because Cubey is part of the singular content of the picture. For the second: Picture A is inaccurate because the target of the picture is *Empty-Desk*, and the content of the picture isn't accurate at *Empty-Desk*. In the Three-Part Model, no conceptual tension is created by a picture whose singular content isn't present (or visible) in its target; such content is simply inaccurate. The same cannot be said for purely attributive approaches to pictorial content, as I'll now demonstrate.

To begin, I note that *Empty-Desk* is naturally construed as the target of Picture A. This was the scene before my eyes at the time of drawing, and what I set out to draw. And if *Empty-Desk* is the target of Picture A, that straightforwardly explains the second assumption, that Picture A is

inaccurate. There is some temptation to think that the target of A might instead be the non-actual scene which I *took* myself to be seeing, which I in fact rendered accurately. But this is to confuse the content I intended to express with the target I intended to accurately depict. It's true that I was in a perceptual state with a certain content, and I did set out to express that content through my picture; that was my expressive intention, and in this case, the intention was fulfilled. But in addition, I set out to make an accurate depiction of a scene in the world, regardless of whether my perceptual content was accurate—and in this respect, my intentions were thwarted. It is only the scene in the world which the picture is intended to be accurate at which counts as the target. In this case, it must be *Empty Desk*.²⁸

The challenge now facing both the synecdochic and descriptivist strategies is to explain how Picture A can be *of* Cubey (Assumption 1) without allowing that Cubey is in Picture A's singular content. But now we can see that the synecdochic strategy is directly counter-exemplified by the Object Hallucination case. For there, the picture is clearly *of* Cubey. But Cubey is not a visible part of *Empty-Desk*, relative to its viewpoint. Thus, in general, what a picture is of cannot be assimilated to the visible parts of its target, and, for all its initial plausibility, the synecdochic strategy quickly unravels.

The descriptivist strategy also fails, for related reasons. If the manner in which the attributive content and target pick out an object is too unrestrictive, then too many objects will satisfy the attributive content of a given picture, and none will be uniquely specified. For example, it is natural to think that an object *O* might be specified as the unique object which satisfies the attributive content *A* in the *time* and *world* of the target scene *S*. (By analogy, as Kripke (1972) construes descriptivism, it holds that a single definite description picks out different objects at different possible worlds.) But as Lopes (1996, p. 97) has noted, this cannot work.²⁹ Applied to the case above, Lopes' point is that there may be indefinitely many other indiscernible cubes at the world and time of the picture's target scene, which satisfy the picture's attributive content. Yet Picture A is of Cubey only—not the many other possible but spurious cubes.

So some more restrictive approach to picking out the relevant object is required. But the same problem of spurious cubes will continue to arise unless the range of available objects is restricted to those which are *visible in the target of the picture*.³⁰ But then, note, we have just recreated the synecdochic strategy, along with its characteristic defects. For as we saw, such a restriction is too

²⁸One might wonder if *Garden* is in fact the target of the picture; but there is little to support this, since that is clearly not the scene I set out to draw, even if it is part of the causal source of my hallucination. Indeed, we could modify the case so that Cubey was originally spotted on my desk, rather than my garden, in such a way that Picture A would have been *accurate* at that scene. Then Assumption 2 would rule this out as the target, leaving *Empty-Desk* as the only available option to explain the sense of inaccuracy.

²⁹Lopes (1996, p. 94) loosely characterizes the descriptivist view as holding that a picture is *of* an object when its attributive content uniquely specifies that object. He doesn't make reference to times or worlds explicit, but these are clearly necessary to the viability of the descriptivist strategy.

³⁰Though Lopes does not discuss this option, it would answer his original objection to descriptivism, since it adequately distinguishes the accurate, qualitatively identical portraits of indiscernible twins.

narrow; in the Object Hallucination case, Picture *A* is of Cubey, but Cubey is not visible in the target scene.

In sum, there is no valid way to use the attributive content, even together with a picture's target scene, to specify the objects the picture is of. If the mode of specification is too relaxed, allowing for objects beyond those visible in the target, then the picture is predicted to be of objects which in fact it is not. If the specification of objects is limited to those visible in the target, then the picture is predicted not to be of objects which in fact it is.

Lopes (1996, pp. 96-97) highlights an additional problem facing the descriptivist, stemming from cases in which a picture misrepresents an object by misattributing certain properties to it. Such an object would not satisfy the picture's attributive content in any straightforward way, raising doubts about the descriptivist's basic strategy of picking out objects via attributive content. The only feasible response is to hold that for a picture to be of an object is for that object to satisfy the picture's attributive content *to some limited degree*. But then all the problems highlighted above come back, only with more force. For now one must attempt to uniquely pick out Cubey from all the objects at the world and time of the scene that fit the attributive content *to some degree*. There seems to be no chance of picking out Cubey, not only from other cubes, but also from other objects that are sufficiently similar.

Together, these considerations rule out views according to which what a picture is of can be derived purely from its attributive content, its target, or some combination of these elements. It follows that, for at least some of the objects a picture is of, these objects directly constitute its singular content.³¹

4.2 Thesis 2B: Target

I turn now to Thesis 2B: that a picture's target does not depend on that picture's attributive content or singular content.

Skepticism about the existence of independent targets may take different forms. One approach holds that pictures are evaluated for accuracy relative to target scenes, but these scenes are derived from facts about singular content and attributive content. In essence, on this view, pictorial content

³¹One might take the arguments here only to support the conclusion that pictures have referents, independent of their attributive contents and targets. But why think that these referents are also part of content? Two reasons prevail.

First, facts about what objects pictures are *of* and about what those objects are represented *as* seem to be thoroughly intermingled. It's not as if a picture is *of* Cubey, and then, separately, depicts something *as* a cube. Instead, it depicts *Cubey as a cube*. In so far as these facts are reflections of content, as I have assumed, they indicate that content contains and integrates both singular and attributive content.

Second, judgements of accuracy support the conclusion that there is at least some level of pictorial content in which what a picture is of determines accuracy conditions. Suppose I draw an accurate picture of Cubey sitting on my desk. My other favorite cube, Yebuc, might be sitting on another desk in the same configuration. An accurate picture of Yebuc might come out qualitatively identical to my picture of Cubey. Still there is a sense in which what is depicted by the original picture—that picture's content—would *not* be accurate at the scene in which Yebuc is visible. This is simply because the content contains the object Cubey, so Cubey helps determine the picture's accuracy conditions. (I don't deny that we may also, more or less at will, track purely attributive content as well. The point is that there is a level of content which incorporates singular content.)

fixes a picture's target; for instance, the target might be defined as any one of the scenes which satisfies a picture's content. But this hypothesis proves unworkable. The problem is that, if content fixes target, it is hard to see how a picture's content could ever be inaccurate, since accuracy is defined in terms of the match between content and target.

A more plausible and challenging view holds that pictures simply have no target scenes relative to which they are evaluated for accuracy. I'll call this the INDEX-FREE APPROACH. According to the index-free approach, pictures have attributive and singular contents, but no targets. But since pictures are accurate or inaccurate, content itself must intrinsically determine an accuracy value, without recourse to a further parameter of evaluation. This presents a challenge, since the objects which are normally thought to make up singular content do not themselves instantiate their properties and relations absolutely, but only relative to a world and time. Here the index-free theorist proposes to build worlds and times *into* the objects that make up the singular content. Singular content, on this account, is not made up of "standard" objects, but world- and time-bound object "tropes". The idea is that by rooting the singular content of a picture to a particular time and world, ascriptions of attributive content may simply be accurate or inaccurate, obviating the need for an additional index.

To develop this strategy more carefully, I introduce the notion of time- and world-bound objects. Such objects correspond to temporal and modal *slices* or *parts* of ordinary objects. For example, Cubey persists through time in the actual world, and has a variety of exciting careers in other possible worlds and times. But for each such world w and time t we may derive from Cubey a time- and world-bound object C , which exists only at w and t . At w and t , C completely overlaps Cubey and has all the same intrinsic properties as Cubey at w and t . To fix ideas, let us say that " $X@wt$ " denotes that object which (i) exists only at w and t ; (ii) completely overlaps X at w and t ; and (iii) has all the same intrinsic properties as X at w and t . If X does not exist at w and t , then the denoting expression is empty. Let us say that $X@wt$ is the INSTANCE of X at w and t . In addition, since scenes determine both times and worlds, let us say that, for a scene S and object X , " $X@S$ " (the instance of X at S) denotes $X@wt$ where w and t are the world and time of S .

The index-free approach proposes that it is *instances* of objects (and not their transworld, transtemporal parents) which make up the singular content of pictures. With this assumption in place, the theorist can now advance a plausible (albeit rough) account of accuracy that is not relativized to a target:

The content of P is accurate iff for every world- and time-bound object O in P 's singular content, O instantiates the properties associated with it by P 's attributive content.

A picture's content is accurate when all the object instances in its singular content instantiate the properties attributed to them. Here it makes sense to talk of properties being instantiated without reference to a world or time, because instances of objects, in virtue of existing only at single times

and worlds, have their properties absolutely.³²

Further, the index-free theorist holds that ordinary judgements about what a picture is *of* can be derived from facts about singular content. A picture is of an object just in case the picture has one of its instances in its singular content:

P is of *O* iff there is some *w* and *t* such that *P* has *O@wt* in its singular content.

Finally, which object instances make up a picture's singular content are thought to be determined by something like a relation of pictorial reference (just as they are in index-based accounts). At the very least, for an object instance to be part of a picture's singular content, there must be some causal connection between the instance and the creation of the picture; and in many cases this causal connection must be mediated by the representational intentions of the artist.

We can now see how the Object Hallucination case might be accounted for under the index-free approach. Recall that the scene before the artist is *Empty-Desk*, but the scene at which the artist originally viewed Cubey, and with which the hallucination is causally linked, was *Garden*. A natural thought is that, because of this link, the region of the picture that depicts Cubey has as its singular content *Cubey@Garden*. Other parts of the picture, for example the regions which depict the desk, will have instances of objects in *Empty-Desk* as their singular contents.


The first assumption about the case is that Picture A is of Cubey. Following the formulation above, the theorist holds that this is so because *Cubey@Garden* is in the picture's singular content. The second assumption is that the picture is inaccurate. This can now be accounted for directly: the picture attributes to *Cubey@Garden* the property of sitting on a desk; but in fact, in *Garden*, Cubey was sitting in a pot of flowers. Thus the picture is not accurate. Thus it seems that the index-free theorist can avoid the challenge posed by the Object Hallucination case.

Although the index-free approach has an answer to the Object Hallucination case, the deeper problem posed by cases *like* it is merely postponed. Ultimately, the index-free approach is committed to a very close connection between the objects a picture is of, and the world and time relative to which its accuracy is evaluated. Cases which dissociate these elements even more extremely than Object Hallucination therefore challenge its foundational assumption. In this spirit, consider the following:

³²In fact, viewpoint-relative relations, legion in pictorial content, still pose a challenge. Such relations are instantiated only relative to a viewpoint, but this appears to be the kind of parameter of evaluation which the index-free theorist sought to avoid. Perhaps there is a way around this. Perhaps the index-free theorist can claim that viewpoints themselves are part of pictures' singular contents. I am skeptical that such a proposal would work, but for the sake of argument, I give the index-free approach the benefit of the doubt.

Scene Hallucination

Yesterday, I see for the first time a cube, named Cubey, sitting on my desk. (Call the scene defined by this world, time, and viewpoint *Desk*.) Today, I visit my favorite forest. I sit down beneath the trees, and set out to draw what I see. (Call this scene *Forest*.) But at that point I suffer from a holistic hallucination, in which the scene I perceived the day before appears as if before me— I no longer perceive the forest. I proceed to draw the situation I take myself to be seeing, producing Picture B.



Picture B

Here, as before, I assume that two facts follow from my description of the case: first, Picture B is *of* Cubey; and second, Picture B is not accurate. But now, the index-free strategy of binding objects to worlds and times no longer helps. The origin of the hallucination is the scene *Desk* perceived on the first day, so it is natural to think that the singular content of Picture B includes Cubey@*Desk*. But if so, then it seems the content of Picture B would have to be *accurate*, since, by stipulation, Cubey@*Desk* has all of the properties the picture ascribes to it. But this is the wrong prediction.

Alternatively, perhaps the singular content of Picture 2 includes Cubey@*Forest*. (Recall that “Cubey@*Forest*” refers to Cubey at the time and world of the scene *Forest*, even if, as is the case here, Cubey is not actually visible in that scene.) But we may extend the case by stipulating that by the time of *Forest*, Cubey’s shape and position have not changed since the time of *Desk*. In that case, the attributive content of the picture (that of a cube sitting on a desk) would still be accurate of Cubey@*Forest*. But this again is the wrong prediction. Attempting to further restrict Cubey to objects *visible* in *Forest* yields no gains, since Cubey is not in fact visible in that scene.

The moral here is that invoking time- and world-bound instances of individuals is not enough. Even when *all* of the objects in a picture’s singular content are bound to the same scene, it may be that the picture must *still* be evaluated for accuracy at yet another scene. This corresponds to the additional index of evaluation championed by the Three-Part Model.³³

³³In a personal communication, John Kulvicki suggests another tack. Instead of thinking of the singular content of pictures as scene-bound objects, model them instead as *(object, scene)* pairs. A picture is accurate when, for each such

The lesson is driven home if we now consider the Three-Part Model's treatment of the Scene Hallucination case. The Three-Part Model holds that the content of Picture B is straightforward—it depicts Cubey (its singular content), as being located in a certain position, sitting on a desk, and so on (its attributive content). What is unusual is that Picture B aims at the target scene *Forest*, rather than *Desk*, and is therefore inaccurate; it is inaccurate both because of the properties it attributes to the scene and because of the objects whose existence it posits.

The Scene Hallucination case shows unambiguously that content cannot determine target. As advertised in Section 1, this implies that two pictures with the same content can have different targets. We can extend the Scene Hallucination case to make this point vivid. Compare Picture B, which was drawn at *Forest*, with a new Picture C, qualitatively identical to B, accurately drawn from life in *Desk*. According to the Three-Part Model, B and C have identical contents, both singular and attributive. Nevertheless, B is inaccurate, because its target is *Forest*, while C is accurate because its target is *Desk*. Such comparisons illustrate the basic dimensions of freedom between content and target defined by the Three-Part Model.

In sum, only the Three-Part Model can explain our judgements about cases in which singular object and target dissociate from one another, as in the two cases above. Theories which collapse or reduce these two elements lack the resources to account for such variation.

5 Non-Factual Targets

Thus far I have articulated the Three-Part Model and argued for its core tenets. In this section I turn to *defend* the model by showing how it can handle a range of potential counter-examples. The examples are culled from the phenomena of *counterfactual* and *generic* drawing. They involve pictures that are used in a manner which is intuitively assertoric, but which seem to have no particular or actual target. Such examples are prima facie challenges to the Three-Part Model's claim that all assertoric pictures have particular targets. The solution, I propose, is to allow for targets which are, on one hand, not maximally specific, and, on the other, not sampled from the actual world. With this kind of NON-FACTUAL target, the Three-Part Model can be extended to a range of new cases.

I begin with the class of what I will call COUNTERFACTUAL pictures. These are drawings which are specifically produced to illustrate counterfactual scenarios. Suppose I ask you, as my architect, what my house *would* look like if such-and-such modifications were made. If you reply with a drawing, this seems to provide an informative answer to my query, and we may judge it

pair, the object in the pair instantiates its associated properties, at the scene in the pair. Allowing for suitable freedom between the elements of the pair, the key assumptions in the Scene Hallucination case are straightforwardly accounted for. But I count this as a formal variant of the Three-Part Model, for it crucially recognizes the same three semantic relata for pictures, and allows that each is independent of the others. The choice between this formalization of the Three-Part Model, and the one I have offered in the text, is guided by methodological considerations such as ease of use and elegance. My own formalization has the virtue of making the structure of pictorial representation commensurate with that of language.

accurate or not, according to whether the house really would look like that under the envisioned conditions. Such pictures may be deemed *assertoric*, for they directly convey information, and they are naturally evaluated for accuracy.³⁴ According to the Three-Part Model, assertoric pictures derive their accuracy values by comparison with a target. But what if anything fulfills the role of target here?

Cases like these present two challenges. The first concerns the modal status of the target. It is clear that, in the case described, the picture is not accurate at any *actual* target; if it is accurate at a target, that target must be counterfactual. So the solution here is to allow that targets need not be culled from the actual world. Of course, one cannot *ostensively* pick out a counterfactual target in the same way one may pick out an actual target in the course of, say, life drawing. But we have already made allowances for descriptive selection of targets— after all, this is what happens when I aim at an (actual) scene which is relayed to me merely via description, as in the case of a police sketch. In the case of counterfactual depiction, the artist’s descriptive intentions simply pick out a target culled from a merely possible world.

The operative description must specify the modal relationship between the actual world and the counterfactual target, in order to capture the assertoric force of the original image. It is not enough to say that the picture is merely accurate at *some* possible world, but rather, that it is accurate at a possible world which is counterfactually related to the actual world in a specific way. Here I appeal to the extensive literature on the semantics of counterfactuals in *language*, such as that of Stalnaker (1968), Lewis (1973), or Kratzer (2012). Authors in this tradition have attempted to give precise conditions whereby a counterfactual supposition may pick out a set of worlds that are suitably “nearby” the world of evaluation. We may assume that much the same mechanisms are in play, albeit at the level of thought, when artists’ intentions pick out a set of possible worlds as their targets in counterfactual depiction.³⁵

The second challenge raised by counterfactual depiction has to do with the *specificity* of the target. Thus far I have proceeded on the assumption that targets are modeled as individual scenes. But my comments above suggest that we may need to reconceive pictorial targets as *sets* of scenes instead. For life drawing, the assumption of a singleton target scene seems reasonable: I intend to draw a bit of the *actual* world, and I intend to draw it from *this* viewpoint. These intentions are sufficient to pick out a particular scene as target.³⁶ But for counterfactual representations, it

³⁴It is important to distinguish architectural *plans* from the kind of drawing here, which (for certain style of depiction) is sometimes called a *rendering*. The former are plausibly expressions of intention, or instructions to builders, and are not assertoric. The latter are illustrations of established plans, and *are* assertoric, though not typically aimed at the actual world.

³⁵An interesting possibility arises here. *Counterpossible* counterfactuals are those whose antecedents specify impossible situations, as in “if there were a square circle, then...” Whatever the semantics of these counterfactuals, itself an open question, they might be deployed to specify *impossible targets* relative to which pictures with impossible content could be accurate. Presumably this is the only kind of situation in which impossible pictures may be accurate.

³⁶Even here it is plausible that artist intentions allow some leeway in the particular viewpoint selected. For example, do intentions select the viewpoint located at one or the other of the artist’s eyes? Or somewhere in between them? If there

isn't plausible that intentions are so specific. I intend to draw the counterfactual scenario as we discussed it, but our discussions did not determine a *unique* possible world, they specified only a *partial* possible situation. And partial situations like this are modeled by sets of worlds.

If targets are understood as sets of scenes, we must adjust the definition of accuracy accordingly. Here I assume that we maintain the core notion of accuracy at an arbitrary scene, and define accuracy at a *target set* of scenes in terms of the former notion:

The content of *P* is accurate at a target set of scenes *S* iff
there is at least one scene *s* in *S* such that the content of *P* is accurate at *s*.

Recall that for a *picture* to be accurate in context, its content must be accurate at its target. Thus, when a picture's target is non-factual, the picture is accurate when there were is *some* scene in its target such that the content is accurate relative to that scene.

Here it is significant that accuracy at a set requires accuracy at *some* member of the set, rather than accuracy at *every* member. Consider again the hypothetical modification to my house. The resulting picture could be quite specific—showing *one* way the house might look if the modifications were made—even though the preceding discussion may have been comparatively open-ended. (For example, the picture might depict the light on the house coming from a specific angle, even though this was not antecedently specified.) Typically, no one picture will be accurate at *every* scene compatible with a counterfactual description of a scenario, but it may very well be accurate at *some* such scene. Our intuitions of accuracy seem to track the latter condition.³⁷

The case of GENERIC pictures presents a related and in some ways more demanding challenge to the Three-Part Model.³⁸ Commonly found in textbooks and encyclopedias, generic depictions are pictures which depict an individual (e.g. Obama) or a kind of thing (e.g. eagles), but not at any particular time or place. As it were, such pictures present a “normal view” of the object in question. Generic depictions are clearly assertoric, in the sense that they are meant to convey accurate information about the actual world, but once again, there seems to be no particular time or place at which they are supposed to be accurate. In fact, in some cases, there may *be* no actual time or place at which they are accurate. For example, an encyclopedia illustration of an eagle may be naturally judged accurate even if, by happenstance, no actual eagle was ever positioned so as to perfectly realize that picture's spatial content.

It is tempting, in these cases, to eschew the framework of the Three-Part Model for a seemingly

is indeterminacy here then the move to model targets as sets may even apply to many cases of life drawing.

³⁷Though it deserves more discussion than I can afford here, I expect that the same basic strategy can be extended to the domain of fictional drawing. Fictions, like counterfactual discourses, can be thought of as specifying sets of possible worlds. Fictional pictures, in turn, can be used to illustrate or extend an existing fiction, or to specify a new one. Though a given fiction may not be real, there is nothing amiss in evaluating a picture for accuracy *relative* to that fiction (following Lewis 1978). Thus I may draw more and less accurate portraits of Sherlock Holmes. In the present framework, that would be mean evaluating the picture for accuracy at a set of scenes compatible with the possible worlds specified by that fiction.

³⁸Thanks to an anonymous reviewer for pressing the challenge posed by generic pictures, and to Susanna Siegel for discussion of possible solutions.

simpler treatment. Perhaps the “target” of a generic picture is in fact a particular *object* (in the case of the generic portrait of Obama) or a *kind* of object (in the case of the generic illustration of eagles)— rather than a scene. On this way of thinking, some element of the content is itself the index relative to which a picture’s accuracy is to be measured. There is no role here for the additional involvement of a target, as I have envisioned it.

But this suggestion cannot work, at least as stated. The problem is that a picture cannot be assessed for accuracy simply relative to an individual. Individuals have different properties at different worlds, and at different times within those worlds. Even a generic picture of Obama would be inaccurate relative to a possible Obama with a differently shaped face, or actual Obama as a child. So it is necessary to define at least a world (or set of worlds) and a time (or time span) relative to which the picture will be assessed. But at this point note that a world and time already make up two elements of the target as I conceive it.³⁹ At the very least, the idea that we might be able to dismiss an index of evaluation, and define accuracy for generic pictures in terms of content alone, cannot work.

How might the Three-Part Model handle the case of generics? Here I’ll focus on the case of generic pictures of individuals. I propose that the special difference between generics and other pictures lies in the selection of target, rather than in their content or definition of accuracy. And the targets of generics, like counterfactual depictions, are non-factual: they comprise sets of scenes, none of which need be actual. For generic pictures, however, targets are picked out in a distinctive way: by reference to part of an individual’s life. The idea is that the target of a generic picture is the sets of scenes which contains an individual in a manner which is *normal* for them— where the notion of normalcy, borrowed from the literature on generics in language, is something like “according to expectations,” or “according to an archetype,” with no requirement of actuality or statistical regularity (Nickel 2016). Roughly, the target of a picture relative to an individual *O*, and timespan *T* of *O*’s life, is the set of scenes *S* such that *O* appears as they normally would during *T*, in *S*. Then such a picture is accurate at its target in the standard way, just in case *some* scene in the target is compatible with the content. Consequently, for a generic *picture* of Obama, in context, to be accurate, its content must be accurate at some particular possible scene where Obama appears as he normally would.⁴⁰

³⁹We must also accommodate the fact that pictorial content is viewpoint relative. Viewpoint relative content can only be assessed for accuracy relative to a viewpoint, so somehow viewpoint must play a role in the definition of accuracy, even for a generic picture. This does not establish that viewpoint must be part of the *target* relative to which a generic picture is evaluated— but the only alternative is that viewpoint is quantified over in the definition of generic accuracy. Indeed, it is natural to think that, for a generic picture to be accurate, its content must accurate at *some* viewpoint located within a “generic” world. The only question is whether the quantification must be built-in as a special feature of the definition of accuracy for generics. In the alternative analysis offered below, I show how such quantification falls out naturally from the framework of non-factual targets, without modifying the definition of accuracy.

⁴⁰The case of generic *kind* pictures is analogous. Roughly, the target of a picture relative to a kind *K* is the set of scenes *S* such that a normal *K* is visible in *S*. A generic picture of an eagle is accurate because it is accurate at a scene where a normal eagle is visible. (Perhaps a temporal restriction is necessary here as well.) At any rate, this a first-pass analysis of this kind of depiction.

Note that, on this account, no particular viewpoint, or set of viewpoints need be picked out *de re* by the artist in selecting a target. The generic picture is accurate if it is accurate at *some* particular scene in the target, but the targets here may vary arbitrarily with respect to viewpoint. (Still, there is reason to believe that only certain kinds of “standard” viewpoints are appropriate for generic depiction.⁴¹ A drawing of a seated Eagle from *underneath*, for example, would plausibly *not* count as a generic picture of an eagle, except perhaps in a specialized context. So some constraints on viewpoint may need to be introduced as well.)

The account of generic pictures we’ve arrived at has the same general form as that offered for the analysis of counterfactual depiction, relying as it does on the notion of a non-factual target, and the accompanying definition of accuracy. For, on one hand, generic pictures aim at targets that are not (necessarily) culled from the actual world, but are related to the actual in reliable ways. And on the other, the targets of generic pictures are *sets* of scenes; a generic picture is true at such a set when its content is accurate at some member of the set.

Although the phenomena of counterfactual and generic depiction were initially presented as challenges to the Three-Part Model, ultimately, I think the considerations here reveal the flexibility of the model to illuminate a wide range of cases— everything from photography and drawing from life to counterfactual and generic depiction. It is clear that much remains to be said, but I hope the sketches offered here suggest that the resources of the Three-Part Model are well-suited to the task.

6 Conclusion

In this paper I’ve argued for the Three-Part Model, focusing exclusively on the case of pictorial representation. But the considerations raised suggest that the same basic semantic architecture may arise in representational systems beyond depiction. As I noted in Section 1, the key elements of (i) singular content, (ii) predicative content, and (iii) an independent index of evaluation are already part of the standard account of the semantics of sentences within philosophy of language and linguistics (Kripke 1972; Lewis 1980; Kaplan 1989). These bear clear analogies to the notions of singular content, attributive content, and target in the Three-Part Model.

Similar comparisons may also apply to the cases of visual perception and mental imagery. In the philosophy of perception, it is commonly (though not universally) thought that perceptual states have both singular and attributive content (Burge 2010; Siegel 2011). Following the spirit of Cummins (1996), I suspect the notion of target also finds a home here. For example, its plausible that much of the argument of this paper which stems from the Scene Hallucination case can be carried over to the domain of perception with little loss. If it can, then a distinction between

⁴¹See Huffman (1971, p. 298) for the concept of a “general” viewpoint, and Willats (1997, pp. 23-24) for discussion. Hagen (1986, ch. 6) describes a number of viewpoint “standards” in non-Western depiction.

content and target would be necessary for perception as well. Further, the mere fact that mental images, like drawings, can represent both past, present, actual and counterfactual scenarios, strongly suggests the work of a counterpart notion of target for imagery. Meanwhile, perception, like photography, seems always to target the scene before the eye. Thus broad analogies in the interplay of content and target emerge between mental visual representation and depiction.

If these speculations are on track, then, across a range of cases, from pictures, to perception, to language, we find the same basic elements that make-up the Three-Part Model: a representation whose function it is to be accurate expresses both (i) singular content, and (ii) attributive or predicative content; and (iii) this content is evaluated for accuracy (or truth) at a target index. Though the rules that associate representations with contents will vary dramatically from case to case, I conjecture that the same basic architecture is characteristic of all sufficiently complex representational systems.

References

- Abell, C. (2009). "Canny Resemblance". In: *The Philosophical Review* 118.2, pp. 183–223.
- Blumson, B. (2009). "Pictures, Perspective and Possibility". In: *Philosophical Studies* 149.2, pp. 135–151.
- (2014). *Resemblance and Representation: An Essay in the Philosophy of Pictures*. Open Book Publishers.
- Burge, T. (1991). "Vision and intentional content". In: *John Searle and his Critics*. Ed. by E. Lepore and R. van Gulick. Blackwell.
- (2010). *Origins of Objectivity*. Oxford: Oxford University Press.
- (2014). "Reply to Rescorla and Peacocke: Perceptual Content in Light of Perceptual Constancies and Biological Constraints". In: *Philosophy and Phenomenological Research* 88.2, pp. 485–501.
- Camp, E. (2007). "Thinking with Maps". In: *Philosophical Perspectives* 21.1, pp. 145–182.
- Casati, R. and A.C. Varzi (1999). *Parts and Places: The Structures of Spatial Representation*. MIT Press.
- Cummins, R. (1996). *Representations, Targets, and Attitudes*. Cambridge, MA: MIT Press.
- Eaton, Marcia (1980). "Truth in pictures". In: *The Journal of Aesthetics and Art Criticism* 39.1, pp. 15–26.
- Frege, G. (1879). *Begriffsschrift, ein der arithmetischen nachgebildete Formelsprache des reinen Denkens*. Translated as *Concept Script, a formal language of pure thought modelled upon that of arithmetic*, by S. Bauer-Mengelberg in J. vanHeijenoort (ed.), *From Frege to Gödel: A Source Book in Mathematical Logic, 1879–1931*, Cambridge, MA: Harvard University Press, 1967. Halle/Saale: Verlag L. Nebert.
- Geach, P. (1965). "Assertion". In: *Philosophical Review* 74, pp. 449–65.
- Giardino, V. and G. Greenberg (2015). "Varieties of Iconicity". In: *Review of Philosophy and Psychology* 6.1, pp. 1–25.
- Goodman, N. (1968). *Languages of Art: An Approach to a Theory of Symbols*. Indianapolis: Bobbs-Merrill.
- Greenberg, G. (2013). "Beyond Resemblance". In: *Philosophical Review* 122.2, pp. 215–287.
- (2016). "The Geometry of Pictorial Representation". Manuscript.

- Grice, H. P. (1975). "Logic and Conversation". In: *Syntax and Semantics, Volume 3: Speech Acts*. Ed. by Peter Cole and Jerry Morgan. Academic Press, pp. 41–58.
- Hagen, M.A. (1986). *Varieties of Realism: Geometries of Representational Art*. Cambridge: Cambridge University Press.
- Hopkins, R. (1998). *Picture, Image and Experience: A Philosophical Inquiry*. 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. (2006). *The Objective Eye: Color, Form, and Reality in the Theory of Art*. University of Chicago Press.
- (2012). "Depiction". In: *Royal Institute of Philosophy Supplement* 71, pp. 129–150.
- Kaplan, D. (1968). "Quantifying In". In: *Synthese* 19.1-2, pp. 178–214.
- (1989). "Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals". In: *Themes from Kaplan*. Ed. by Joseph Almog, John Perry, and Howard Wettstein. Oxford University Press, pp. 481–563.
- Kjørup, S. (1974). "George Inness and the Battle at Hastings, or Doing Things with Pictures". In: *The Monist* 58.2, pp. 216–235.
- (1978). "Pictorial speech acts". In: *Erkenntnis* 12.1, pp. 55–71.
- Koenderink, J.J. and A.J. van Doorn (2003). "Pictorial space". In: ed. by H. Hecht, R. Schwartz, and M. Atherton. MIT Press.
- Korsmeyer, C. (1985). "Pictorial assertion". In: *The Journal of aesthetics and art criticism* 43.3, pp. 257–265.
- Kratzer, A. (2012). *Modals and conditionals: New and revised perspectives*. Vol. 36. Oxford University Press.
- Kripke, S. (1972). "Naming and Necessity". In: *Semantics of Natural Language*. Ed. by Donald Davidson and Gilbert Harman. Revised edition published in 1980 as *Naming and Necessity* (Cambridge, MA: Harvard University Press). Dordrecht: Reidel, pp. 253–355, 763–769.
- (1977). "Speaker's Reference and Semantic Reference". In: *Midwest studies in philosophy* 2.1, pp. 255–276.
- Kulvicki, J. (2006). *On Images: Their Structure and Content*. Oxford: Clarendon.
- Lepore, Ernie and Matthew Stone (2015). *Imagination and Convention: Distinguishing Grammar & Inference in Language*. Oxford University Press.
- Lewis, D. (1973). *Counterfactuals*. Blackwell Publishers.
- (1978). "Truth in Fiction". In: *American Philosophical Quarterly* 15, pp. 37–46.
- (1980). "Index, Context, and Content". In: *Philosophy and Grammar*. Ed. by Stig Kanger and Sven Ohman. Dordrecht: D. Reidel, pp. 79–100.
- Lopes, D. (1996). *Understanding Pictures*. Oxford: Oxford University Press.
- MacFarlane, J. (2009). "Nonindexical contextualism". In: *Synthese* 166.2, pp. 231–250.
- McDowell, John (1984). "De re senses". In: *The Philosophical Quarterly* (1950-) 34.136, pp. 283–294.
- Murray, Sarah E and William B Starr (2017). "The structure of communicative acts".
- Newall, Michael (2011). *What is a Picture?: Depiction, Realism, Abstraction*. Palgrave Macmillan.
- Nickel, Bernhard (2016). *Between logic and the world: An integrated theory of generics*. Oxford University Press.
- Novitz, D. (1975). "Picturing". In: *Journal of Aesthetics and Art Criticism* 34.2, pp. 145–155.
- Panofsky, Erwin (1924 (1991)). *Perspective as symbolic form*. Ed. by Christopher Wood. Zone books New York.

- Peacocke, C. (1992). *A Study of Concepts*. The MIT Press.
- Rescorla, Michael (2008). "Predication and cartographic representation". In: *Synthese*.
- Roberts, C. (1996). "Information structure in discourse: Towards an integrated formal theory of pragmatics". In: *Working Papers in Linguistics-Ohio State University Department of Linguistics*, pp. 91–136. ISSN: 0473-9604.
- Rogers, S. (1995). "Perceiving pictorial space". In: *Perception of space and motion* 5.
- Ross, J. (1997). *The Semantics of Media*. Kluwer Academic Publishers.
- Schier, F. (1986). *Deeper Into Pictures: An Essay on Pictorial Representation*. Cambridge University Press.
- Siegel, S. (2011). *The Contents of Visual Experience*. Oxford University Press.
- Stalnaker, R. (1968). "A theory of conditionals". In: *Ifs. Conditionals, Belief, Decision, Chance, and Time*. Dordrecht, D. Reidel Publishing Company, pp. 41–55.
- (1978). "Assertion". In: *Syntax and Semantics, Volume 9: Pragmatics*. Ed. by Peter Cole and Jerry Morgan. Academic Press, pp. 315–332.
- Willats, J. (1997). *Art and Representation: New Principles in the Analysis of Pictures*. Princeton: Princeton University Press.