

DEPICTIVE SEEING AND DOUBLE CONTENT

John Dilworth

[Forthcoming in (eds.) C. Abell and K. Bantinaki, *Philosophical Perspectives on Picturing*, Oxford University Press.]

A picture provides both configurational content concerning its design features, and recognitional content about its external subject. But how is this possible, since all that a viewer can actually see is the picture's own design? I argue that the most plausible explanation is that a picture's design has a dual function. It both encodes artistically relevant design content, and in turn that design content encodes the subject content of the picture--producing overall a *double content* structure. Also, it is highly desirable that a resulting double content theory for pictures should be closely integrated with a related double content account of perceptual content generally, so as to avoid suspicions of ad hoc theorizing that would apply only to pictorial content.

The resulting theory should also be able to explain the inevitable ambiguities involved in abstracting two levels of visual content from a single visible surface, as well as explaining the systematic relations between the two kinds of content. I provide an *orientational* theory--based on a recently developed spatial logic of orientational concepts--for this purpose, and show how depictive and perceptual content in general can be usefully explained in these orientational terms. This account of picturing also integrates well with a previously developed, more generic double content theory of art, and it is also plausible in cognitive science terms.

This paper provides an overview of a comprehensive double content (DC) approach to the depictive issues of seeing-in, as discussed by other contributors to this volume such as Hopkins, Lopes, Nanay and Brown. A potential advantage of the DC approach to pictorial depiction is that it has previously been systematically worked out in two much broader contexts. First, my 2005 book *The Double Content of Art* argues that all of the arts--and not just depictive arts such as painting or photography--involve similar double content structures. And second, the DC account has also been extended to perceptual content generally in several recent articles [Dilworth 2005b, 2005c]. These latter articles

securely embed the DC account in a cognitive science framework, which--among other things--enables the account to avoid the potential dangers of an over-reliance on everyday assumptions concerning experiences of pictures.

A useful point of entry for subsequent discussion is provided by Richard Wollheim's influential concept of 'seeing in' [Wollheim 1968, 1987]. In particular, Wollheim's distinction between the configurational, or design, features of a picture and its recognitional features is taken as a point of departure for such accounts. Any adequate account of depiction must explain how two distinct elements--namely, the purely physical design features of the surface of a picture, and the subject matter that is seen in the picture--are related to each other in our experience of pictures as pictures. Consequently, Lopes, Hopkins and Nanay follow Wollheim in assuming that, insofar as pictorial experience has two kinds of content, they are best distinguished as design-related, configurational content on the one hand, and subject-related, recognitional content on the other hand.

However, from a cognitive science point of view this is a confused--or at least superficial or inadequately analyzed--content-based contrast. This is because the purely physical design or configurational features of the surface of a picture must provide the physical basis for *any* kinds of content whatsoever that are associated with the picture. If, for example, the relevant two kinds of content are kinds of perceptual content associated with viewings of the picture, then *both* kinds must be derivable from viewings of the same single physical design on the surface of the picture. Thus there is a sense in which all

pictorial content must be configurational content, since all of it must be derived from the same visible design.

This is not to deny that there are two saliently different kinds of content that are thus derivable, but it is to insist that their differences cannot be explained purely in terms of their common derivation from a single design. An informal exposition of some legitimate content differences will now be provided, after which a more systematic account will be presented in subsequent sections.

A typical picture, such as Rembrandt's pen drawing of pastor Jan Cornelius Sylvius, as discussed by Hopkins, in some respects resembles an actual person seen face-to-face, while in other respects--such as in the directions of its inked lines--it instead resembles the design configuration on the surface of the drawing. Presumably some such contrast underlies Wollheim's intuitive configurational versus recognitional content distinction, as further developed by Hopkins et al. However, the comparison is superficial or preliminary only, because it is those very inked lines themselves that also must provide the respects in which the picture resembles a person seen face-to-face, since strictly all that is visible in the picture is those inked lines.

But, it could reasonably be asked, how could two different kinds of perceptual content be derived from the very same visible areas of a picture? Would they not visually compete with each other, so that, for instance, once all of the configurational, inked-lines content of the pen drawing have been accounted for, there would be no areas of the picture left

over to represent the person who is portrayed by the drawing? In reply to such concerns, clearly the visible ink lines have to do *double duty*--both providing content or information about the features of these specific drawn lines that show Rembrandt's pen-drawing style, and providing information about the person depicted in the picture as well. I have argued elsewhere [Dilworth 2005a, 2005e] that in cognitive science terms there is really only one possible solution to this potential problem of informational conflict--namely, that the two kinds of content or information must be *hierarchically* arranged, in two distinct levels, with one level encoding or representing the other.

Also, arguably the most plausible hierarchical arrangement is one in which lower-level, design-related content--providing information about the configuration of lines etc. on the picture surface--is used to encode or represent a higher level of subject-related content. This is the structure adopted by the current DC account of depiction. So on this DC account, viewers of a picture see the depicted subject in the picture by seeing the manner in which the lower-level configurational content of the picture represents that higher-level subject. The next section will give a preliminary, intuitive account of how these two levels of content could explain seeing-in, and will preview further details of the DC theory.

1. Double Content and Perceptual Ambiguity Issues

It is important to situate the DC account in a broader context of human perceptual abilities and representational processing structures. It will be shown that closely related

double content structures must be involved in normal perception of objects that are not pictures, so that the postulated DC structure of depictive perception is plausible, in that it is no more than a natural extension of more generic perceptual abilities.

But in addition, it will also be important to be able to theoretically account for a significant ambiguity problem that is unavoidably associated with the suggested DC approach--both with respect to pictures, and for perception generally. The problem is that any object having a two-level DC structure is inevitably *potentially ambiguous* between various distinct DC interpretations, because by hypothesis at most one of the levels is directly or basically perceived--the upper-tier represented level must be inferred or interpreted from the data provided by the lower level. But in fact the situation is even worse than this preliminary description suggests, because there are potential conflicts between the two levels of content as well, so that strictly speaking, neither level of content is fully determined by the initial data as derived from the surface of a picture or from worldly objects generally. A legitimate interpretation would assign appropriate determinate content to each level, but the basic problem is that there would always be more than one possible way of doing this for each level.

Hence *representational ambiguity* ensues: potentially there are a range of distinct, at least minimally legitimate compatible double content interpretations of the retinal configuration, so that there is no one unique double content structure associated with a concrete perceptual state S. For example, retinal information indicating that blue light was visually received might show that object X was blue, and illuminated by ambient

neutral white light, or it might instead show that object X was white, but illuminated by ambient blue light. (There is also an indefinite range of equally compatible, though less extreme, aspect and object data combinations for colour factors).

A similar point holds for some concrete shape pattern on the retina, such as an elliptical shape. That shape might have been caused by light emanating from an actually elliptically-shaped object X directly facing the perceiver, or from a circular object X that was viewed from an oblique angle, and so on. Hence any concrete retinal shape properties are just as ambiguous with respect to double content interpretations of them as are concrete colour properties of the retinal image. But, since arguably all low level information provided by retinal images is either shape or colour information, it follows that all of the basic information derivable from concrete retinal states is potentially multiply ambiguous with respect to compatible double content interpretations of those states [Dilworth 2005c].

2. Basic Similarities in the Perceptual Content of Pictures and Non-Pictures

It might initially be assumed that perception of an external representation, such as a picture, is fundamentally unlike perception of an ordinary real object, such as a cow, which does not represent anything. However, we must be careful not to confuse the legitimate differences in representational status of the relevant objects themselves--pictures versus non-pictures--with differences in the perceptual content of a person who

is viewing each kind of item. There are at least three fundamental kinds of similarities in perceptual content that apply to both classes of object, which will now be briefly summarized.

As a preliminary, note that some sort of distinction of two kinds of perceptual content is inevitable in the case of pictures, because of the differences between their surface designs and their depicted subject. One way to express this point is that, for any given subject *S* such as a scene or a person, there could be many different pictures of it, each having the same subject *S* but with a different design content *D*₁, *D*₂, ...*D*_{*n*}. However, there is an analogous distinction between perception of an object *O*, and perception of various *aspects* *A*₁, *A*₂, ...*A*_{*n*} of that object *O*, as seen from various distinct perceptual viewpoints. To invoke a G.E. Moore-style argument, just as I can know that I see a hand when I look at one of my hands, so also can I know that in so doing, I also see some particular aspect of that hand--such as the palm of my hand, or the back of the hand. But since my perceptual content always includes both the hand, and some aspect of the hand, and since the two must be distinct--since there is only one hand, but many aspects of it--there must be two distinct kinds of perceptual content involved in one's perception of the hand, just as there must be in one's perception of a picture and its subject.

The relevant concept of an *aspect* of an object may be generalized to include any broadly aspectual or contextual conditions under which objects are perceived, including lighting, distance, atmospheric conditions such as rain or mist, and partial occlusion by other objects. This generalization is legitimate because any such factors can affect our retinal

images of objects, which provide the proximal source of all our visual information about them. So we may distinguish the *aspect* or aspectual content of retinal visual images, from their specifically object-related *intrinsic* content. Arguably aspect and intrinsic content together provide a generic two-level DC structure of perceptual content, which is closely analogous to the two-level DC structure of design and subject content in the case of pictures.

The second fundamental similarity in perceptual content for pictures and non-pictures is provided by the mediating role of retinal images themselves. Indeed, it is arguable that such retinal images are themselves closely analogous to external pictures such as projected screen images, in that, physically speaking, a retinal image literally is a projected image focused on the retina by the eyeball. In addition, the double perceptual content of external pictures must always be perceptually mediated via the retinal images of viewers of the pictures, so that inevitably there must be an intimate relation between the contents of perception of external pictures and of non-pictures. Or in more detail, when a picture is perceived, the image of the picture that is formed on the retina of the perceiver will approximately duplicate the main qualitative features of the surface design of the picture itself. Consequently, a perceptual interpretation of the design of the picture would involve--as applied to the closely similar retinal image--the same standard double content extraction procedures as would be used to extract content from any retinal image of worldly objects. In other words, since all visual perceptual interpretation is mediated by retinal images, it should not be surprising that those retinal images that represent

external pictures are internally processed in similar structural ways to other more miscellaneous retinal images.

As for the third fundamental similarity in perceptual content for pictures and non-pictures, it is that generic perceptual content is just as fundamentally ambiguous as is pictorial content--as will be further demonstrated in the next section.

3. More on Double Content and Representational Ambiguity in Perception and Depiction

This section provides a more detailed theoretical overview of issues concerning generic perceptual double content and its inevitable ambiguities. These generic perceptual issues will also be closely related to pictorial double content issues [Dilworth 2005a-c].

Normal perception of an actual object X involves a particular viewpoint of the perceiver with respect to X, so that different aspects Y of object X will be seen from differing viewpoints. Perception of such an aspect Y provides information both about object X itself, and about various concrete aspectual factors Y1, Y2, ..., such as intervening objects, lighting and atmospheric conditions, and so on, that together make up the totality of what can be perceived from a given perceptual viewpoint during observation of object X. A primary task of normal perception is to extract from such viewpoint-dependent aspects correct information about object X itself, free of the potentially distracting

aspectual factors. An early stage of this perceptual extraction involves the processing of sensory information derived from a concrete sensory input state S , such as a concrete pattern of retinal excitation caused by light rays emanating from a given viewpoint-dependent aspect of X .

Hence, in symbolic form, the initial or low-level informational content Z' derivable from such a concrete sensory state S thoroughly intermixes data Y' --about specifically aspectual worldly factors Y --with data X' about X itself. Hence, in order to extract useful information about object X itself, the perceptual system must somehow transform or decode the low-level data $Z' = Y'(X')$ into upper-level data X' that is specifically about object X itself. (In typical perceptual cases, the extracted aspect content factors Y' , which provide information about worldly aspectual factors Y , would be discarded as being irrelevant to the main perceptual task). Consequently, perceptual data has a two level or double content structure in which low-level, intermixed information $Y'(X')$ of a raw, broadly aspectual kind, has to be decoded to extract the upper level objectual information about object X .

However, there is a fundamental epistemic problem associated with this task, namely that a concrete sensory input state S , such as a concrete pattern of retinal excitation as discussed above, provides no guaranteed or automatic procedure for separating out the genuinely X -related information X' from more miscellaneous information Y' that is about the aspect Y and its aspectual factors Y_1, Y_2, \dots , rather than specifically about X . The perceptual state S potentially contains information both about object X and about the

ambient, aspectual conditions under which X is being perceived. But since two distinct content factors Y' and X' have to be extracted from the single retinal state, arguably there is a range of distinct possible two-factor interpretations, each of which is compatible with the same concrete retinal state.

The general point that perceptual data is fundamentally ambiguous or indeterminate in various respects is widely accepted in vision science--see, e.g., Purves & Lotto [2003] for further examples. However, the specific double content structures summarized here, and defended in more detail elsewhere as noted, provide a novel, DC-based theoretical approach for articulating what are, arguably, the two most salient kinds of content involved in such perceptual cases of informational ambiguity.

Turning now to external representation cases such as pictures, arguably a closely analogous structure, along with similar possibilities of representational ambiguity, holds in their case as well. Just as a concrete perceptual state S has a low-level, broadly aspectual content of form $Z' = Y'(X')$, so also does a concrete external representation R have a low-level, broadly design or configurational content of the same form $Z' = Y'(X')$ --which content Z' is the raw or unprocessed design information derivable from the visible surface of the relevant picture. In addition, in both perceptual and external representation cases, an upper level content X' must be extracted by some means--where X' is content that is about an object X in generic perceptual cases, and about a subject matter X of the picture in a specific pictorial case.

As to the status of the decoded lower-level aspectual or design content factors Y' in both generic object and picture perception cases, this information may simply be discarded in utilitarian cases, such as everyday perception of objects or vacation snapshots, in which all that is desired is accurate information about the actual properties of the objects that were seen or photographed. But arguably artistic pictures gain much of their significance--including their typical inflective differences from face to face seeing of an actual object X--from perception of such lower level design factors Y'.

As for the issue of how potential double content ambiguities could be resolved in generic perceptual cases, various views are possible. Purves & Lotto [2003] argue that broadly historical statistical factors, such as probabilities of previous correlations between a given kind of low-level retinal data pattern and actual facts about the objects being perceived, are sufficient to resolve the ambiguities. I would add that in many cases of normal perception of non-representational objects, the potential ambiguities could be resolved by more comprehensive current perceptual observation. Potential ambiguities in allocation of colour information to either aspectual or object factors could be reduced by examining the ambient illumination directly. Or potential shape ambiguities could be reduced or eliminated by viewing an object from a range of viewpoints.

By contrast, in typical pictorial cases, artistic intentions help to resolve ambiguities. Here is an example that involves artistic intentions and the nature of artistic media [Dilworth 2005a]. Presumably Rembrandt intended viewers of his pen drawing (as discussed previously) to see it as depicting a person. On such a legitimate interpretation, arguably

only some of the pen lines could be concurrently noticed as such by a viewer, because noticing *all* of them *as pen lines*--i.e., as lower-level content only--would be perceptually inconsistent with concurrently noticing Rembrandt's intended depicted subject at all. Or, to put the matter another way, though all of the pen-line-derived data must be initially processed as lower-level configurational content, in a legitimate interpretation some of it should ultimately be interpreted as transparently representing the subject--i.e., as upper-level content only--while other elements may be semi-transparent, i.e. providing content elements at both levels, while still others might provide low-level content only that provides no subject information (such as the frame or signature).

Semi-transparent cases, involving both levels of content, could be illustrated as follows. For these cases too there could be both appropriate and inappropriate interpretations. For example, Rembrandt's drawing could be incorrectly interpreted as indeed depicting a person, but nevertheless one who has strange pen-like markings all over his face. This possible interpretation is semi-transparent (involving both levels of content) because the pen markings are seen as being *on the depicted person's face*--rather than just as being meaningless lower-level pen-strokes--but it is nevertheless wrong because it would betray a fundamental misunderstanding of the artistic medium of pen drawing. Or, as Gombrich remarks in a similar case, a drawing of Tivoli should not be interpreted as depicting Tivoli as being bounded by wiry lines [Gombrich 1962: 78]. A DC-based theory capable of explaining the basic cognitive factors involved in such ambiguity cases will be provided in the following sections.

To sum up the points raised in this section, the two-level perceptual content associated with viewing a picture is potentially ambiguous between various distinct two-factor content interpretations of its design. The potential ambiguities would need to be resolved via constraints provided by artistically relevant factors, such as the artist's intentions, or relevant media conventions such as those applying to drawings rather than paintings. In addition, more generic kinds of perception of worldly objects involve structurally similar kinds of perceptual double content--except that in their case, the standards for appropriateness of double-content interpretation are provided by normal standards for veridical perception of objects and their properties.

4. An Orientational Double Content Theory: Basic Concepts

Because of the importance of double content ambiguities for the understanding of both perception in general and depiction in particular, it is appropriate to develop a theoretical structure that is specifically geared to the analysis and articulation of the full range of relevant factors in such double content ambiguity cases. The following theoretical structure was initially developed [Dilworth 2005a, chs. 8-10 and 12] to explain pictorial ambiguities that become salient in cases of changes in *spatial orientation* of a picture, and hence it is appropriately described as an *orientational* theory of double content ambiguities.

As a preliminary, since a given picture--such as Rembrandt's pen drawing--could be interpreted in more than one distinct double content way, we shall describe each of those ways as characterizing a distinct *pictorial content* (PC) PC1, PC2, .. of the concrete design provided by the picture. The orientational theory to be provided shows how these various possible distinct double content interpretations PC1, PC2, ... of the picture relate both to each other, and to the basic raw or uninterpreted information that can be perceptually extracted from the picture surface via the concrete retinal configurations of a perceiver of the picture.

To begin, here is a basic example that shows a need to introduce orientational concepts. A picture is capable of being interpreted in at least two different pictorial ways, involving distinct pictorial contents PC1 and PC2, as follows. If a picture of an upright house is spatially inverted, an ambiguity becomes apparent: the inverted picture can either be interpreted as an inverted picture of an upright house (PC1), or as an upright picture of an inverted house (PC2).

The relevant pictorial contents PC1 and PC2 must be distinct, because each involves a distinct subject matter--namely, an upright house in the first case, and an inverted house in the second case. But the two distinct pictorial contents PC1 and PC2 are associated with only a single concrete picture. Hence it becomes necessary to distinguish at least three different kinds of items, and three corresponding relevant orientational factors--respectively, those applying to the concrete picture, to each distinct pictorial content PC1 and PC2, and to the distinct subjects S1 and S2 of each pictorial content.

An adequate orientational theory must articulate all of the relevant orientational factors involved in such cases. It is immediately clear from the example that at least two distinctive orientational factors are involved, namely the spatial orientation of the concrete picture P on the one hand, and the differing spatial orientations of the subject matters S1 and S2 of each pictorial content PC1 and PC2 on the other hand. But arguably a third kind of orientational factor, which concerns the spatial orientation of the distinct contents PC1 and PC2 themselves, must also be involved.

This is so for the following reason. The orientation of the subject matter of a pictorial content (PC) cannot be determined independently of knowing the orientation of the PC itself. For example, if a picture P depicts an inverted house, then it must be possible to see that the depicted house is inverted when the relevant PC is viewed in its normal upright orientation. But this is only possible if the PC itself has an orientation that is independent of the orientation of its subject matter S. And finally, arguably a fourth, broadly environmental orientational factor E is required as well, so that the orientation of other items relative to environment E can be described.

The most basic orientational concept in the current DC theory is that of a *field*. A field--to be discussed further in the following section--is an abstract structure of different specific qualities of some generic, homogeneous kind, such as some interrelated spatial positions, or some interrelated colours, or shapes. In the case of a spatial field, a field is a three dimensional structure of abstract spatial points, the coordinates of which could be

used to describe or index corresponding positions in some region of three dimensional space. Each of the four orientational factors described above would involve, among other things, its own distinct spatial field, so that four implementations or tokens of the same abstract spatial field type would be required to fully explain the orientational structure of the relevant example.

Second, a concept of *field orientation* needs to be defined. This concept concerns the way in which two distinct copies of the same field type relate to each other. For example, if a concrete pictorial token T has had a spatial field F(T) assigned to it, then the field orientation of token T with respect to the environmental field E is defined in terms of the correspondence or alignment between the differing coordinates of field E and field F(T).

Third, a concept of *intrinsic orientation* is also required. Whereas the concept of field orientation (or alignment) is a relational concept specifying the relations of two distinct though type-identical fields, the concept of intrinsic orientation instead presupposes that there is some functionally defined standard that specifies *a unique internal positional structure* for elements of the field (in some particular application of the field structure). In particular, a *top* distinguished element or layer must be specifiable, which thereby counts as an *intrinsic top* of the relevant intrinsically oriented field. In contrast, a *field top* of an item X is whichever side or region of X is aligned with the intrinsic top of the relevant environmental field E, whether or not X has an intrinsic top--and, if it does,

whether or not its current field orientation with respect to E results in its intrinsic top being aligned with its field top.

For example, a spatial environmental field E is normally regarded as having an intrinsic orientation in which elements of the field that are more distant from the Earth's surface are taken to be *higher than* elements that are closer to the Earth, so that the intrinsic top of the field would be the layer of elements that are at the highest altitude (most distant from the Earth). Or, as another example, a car has an intrinsic orientation, with its intrinsic top being the top of the roof, because that is the functionally based orientation in which cars are normally driven. However, by contrast, miscellaneous items of hardware such as small hand tools do not have an intrinsic orientation because they lack any unique, functionally specified orientation in which they are normally used.

The concept of intrinsic orientation is particularly important in analyzing pictures. Both PCs PC1, PC2, ..and their corresponding subject matters S1, S2, .. possess an intrinsic spatial orientation, and in the analysis to be given concrete pictures P do so as well, as part of the analysis, even though such pictures need not independently have an intrinsic orientation in their own right, simply as physical objects. Also, it will turn out that when the orientational theoretical framework is generalized--to include colour, shape, etc. fields as well as spatial fields--pictorial contents,, subject matters and pictures will continue to possess an intrinsic orientation and an intrinsic top element with respect to their relevant fields.

Fourth in this list of basic orientational concepts is a concept of *uprightness*, along with a correlative concept of *inversion*. An item *X* is *upright* just in case it has an intrinsic top that is currently aligned with the top of an intrinsically oriented environmental field *E*--i.e., in the spatial case, if *X*'s intrinsic top is currently spatially above any other parts of *X*, relative to the environmental field. Correspondingly, an item *X* is *inverted* in spatial cases if its intrinsic top is currently spatially below any other parts of *X*. It will also be convenient to use a more inclusive concept of inversion for any case in which item *X* is not upright.

Finally, a fifth basic orientational concept comes in the form of a principle: the *OSMI* principle, or *oriented subject matter invariance* principle [Dilworth 2005a, chs. 9, 12]. As a simple example, we normally assume that if a concrete picture *P* of an upright house is spatially inverted, so that the field orientation of picture *P* relative to environmental field *E* changes by one hundred and eighty degrees, nevertheless the field orientation of the subject matter *S*--which is defined relative to the intrinsic top of the relevant pictorial content *PC1*--does *not* change during the rotation, in that the picture remains one that is interpreted as being a picture of an upright house, rather than its changing into a picture of an inverted house.

5. Applications of Orientational Concepts

The orientational concepts just defined will now be applied both to concrete retinal tokens Tr , and to aspect $Y'(X')$ and object X' contents. A parallel application to concrete pictures P , double pictorial contents $PC = Y'(X')$ and subject matters S will also be described. The orientational concepts provide ways in which to keep track of changing relationships between these three kinds of items, plus of their relations to a fourth item--an environmental, intrinsically oriented field E . However, these structured orientational relations are not just a perspicuous way of describing the relevant relations. In addition, they are intended to be a useful tool in articulating a broadly truth-functional or correctness-based view of depiction--which closely associates depictive content with a correctness-based view of perceptual content as well. Thus, on the present view, pictures can correctly or incorrectly represent their subjects, just as perception can correctly or incorrectly represent worldly objects.

Consequently, the current orientational DC framework must fully support a clear distinction between correctly and incorrectly depicted subjects. It does so by identifying *orientational uprightness of subject* in some respect--such as colour, shape, or spatial orientation--with *correctness of subject* in that respect--and generic inversion in the relevant respect with incorrectness of subject. Similarly, in the case of object-related perceptual content X' , uprightness of its relevant field is identified with correctness of object-related content.

Next, here are some further details on the structure of a field. Because of the emphasis on correctness of subject matter, the elements of the relevant field are defined in the following way. First, there is an intrinsic top element, consisting of the relevant kind of correct subject matter--such as a particular correct spatial orientation, or a correct colour, or a correct shape, and so on. As a result, all four copies of the relevant field have the same intrinsic top and intrinsic orientation. Environmental field E provides a fixed reference standard, relative to which the other three intrinsically oriented fields may have differing field orientations.

Second, the rest of the elements in the field are defined in a structured relational way that corresponds to (what could be called) *relevant psychophysical alternatives* for the kind of property being considered. For example, with spatial orientation, the relevant alternatives are other spatial orientations to which a picture could be rotated. In the case of colours, because of our psychophysical makeup, the other colour elements in the field would roughly correspond to some standard 'colour wheel' structure for the colours. In the case of perceptual shapes, the actual shape of the subject matter would be the intrinsic top of the field, with other retinal shapes produced by actual viewing of such an object from various angles making up the rest of the field elements.

In depictive cases, stylistic and medium-related considerations also serve to structure other characteristic kinds of field, in ways to be discussed further in section 6. For example, a medium such as pencil drawing would use fields in which top elements were fully realistic renderings of a subject, with increasingly graphic elements such as linear

shading and cross-hatching becoming more prominent in the less realistic, more inverse elements of a field.

As an initial application of the orientational framework, consider a colour change case. During stormy weather, daylight that is normally a neutral white colour tends to take on a bluish tinge--as evidenced by any number of undesirably bluish photographs taken under such conditions. As perceivers with sophisticated perceptual constancy mechanisms, we are usually able to compensate for such cases of abnormal ambient lighting. Hence normally, a white sheet of paper continues to look white to perceivers of it, whatever the weather conditions, or time of day, may be. But colour information extracted directly from a perceiver's concrete retinal image of the paper in such cases would still contain an unusually high amount of bluish light. How are we able to correctly perceive the inherent colour of the paper as being white under such conditions?

The orientational framework in this case would involve a colour field with a white intrinsic top, since white is the correct colour of the paper. As a preliminary, recall that low level aspectual double perceptual content is of form $Y'(X')$, involving aspectual factors Y' and objectual factors X' , and that this double content information has to be extracted from a concrete retinal image token Tr . Also, since the paper X is indeed white, and it is correctly perceived as white by normal perceivers, the paper-related perceptual content X' of such a perceiver must be upright relative to the relevant environmental field E .

Now in the orientational model, the colour that the paper is perceived to have is the colour that is the current *field top* (relative to field E) in the object-related content X'. Since this is a case of correct perception of colour, the field top of X' is also identical with the white intrinsic top of the colour field, so that the paper is perceived to be white. But nevertheless, the colour-related properties of the concrete retinal image token Tr include a predominance of bluish-light related properties, so that consequently token Tr's copy of the relevant colour field is *not* upright, but instead it is generically *inverted* with respect to reference field E. Hence there is a potential compatibility problem: how can the actual white colour of the paper be made consistent with the physical bluish properties of retinal token Tr, which consequently makes token Tr's own copy of the colour field non-upright relative to field E?

Clearly the answer is that perceivers can correctly compensate for the bluish ambient light conditions by adopting a full double content perceptual interpretation Y'(X') of the retinal image token Tr in which the *aspectual* content Y'(X'), and its included aspectual colour factors Y', are *non-upright*, or inverted, in the orientational model, *to the same degree as token Tr itself*. Such a double content interpretation is fully compatible with the non-upright orientation of token Tr, in that the predominance of bluish light properties in token Tr is fully explained as being entirely due to aspectual ambient light factors, rather than its being due to the inherent colour of the paper.

An alternative interpretation of the retinal image token Tr will now be considered. To make the alternative interpretation plausible, and to link it to a case of depiction, suppose

that a colour photograph T_p was taken of the original piece of white paper X from the same position, and at roughly the same time, as the original perceptual observation of X that produced the bluish retinal image Tr . Then a subsequent viewing of that photographic token T_p could produce a retinal image Tr_2 that was qualitatively very similar to the initial retinal image Tr of the actual sheet of paper. However, suppose that the lighting conditions have now changed to neutral white light, so that, in viewing the photograph T_p , there is no automatic perceptual compensation by the viewer for the bluish tinge of retinal image Tr_2 . Instead, the photograph is perceived by its viewer as depicting a bluish sheet of paper, rather than a white sheet of paper. How is this difference from the original perception of the paper as being white to be explained?

To begin, since the sheet of paper X depicted by the photograph--i.e., the subject matter of the photograph--is in fact white as before, the appropriate colour field is also the same as before, with the same white intrinsic top. However, the depictive token depicts the subject matter--the paper--as being bluish in colour rather than white. Consequently, in the orientational model the colour field for the subject matter is no longer upright relative to the environmental reference field E . Instead, its field top is the bluish colour that the photograph is perceived as depicting with respect to its subject matter, rather than the actual white colour of the intrinsic top of the subject matter field. Hence, the subject matter field must be *non-upright*, or broadly *inverted*, relative to the environmental field E , which indicates in the model that the colour of the subject matter is incorrectly depicted by the photograph.

It still remains to be explained how this broadly inverted subject matter interpretation can be rendered compatible with the full current double content interpretation $Y'(X')$ of each of the two tokens--namely, bluish retinal token Tr_2 , and the bluish concrete depictive token Tp itself. But perhaps it is clear enough that the broadly inverted subject matter, with its incorrect bluish field top, fully explains by itself how the bluish properties of retinal token Tr_2 and depictive token Tp have been processed. Clearly it has been perceptually assumed that the full informational content of the bluish qualities of concrete tokens Tr_2 and Tp is exhausted by the bluish qualities of the relevant subject matter, namely the paper.

Consequently, in the orientational model the pictorial content PC itself--i.e. the full, low level aspectual content $Y'(X')$ --would normally be being interpreted as being *upright* in its field, i.e., with a white field top. Hence the full aspectual content $Y'(X')$ that constitutes the pictorial content PC is assumed to be a neutral white colour, and consequently assumed to have no impact on the relation between the bluish qualities of tokens Tr_2 and Tp and the perceived bluish subject matter of picture P . Or, in terms of the specifically perceptual content of token Tr_2 itself, the aspectual content Y' is assumed to involve nothing more than neutral ambient white light that could have no effect on the colour content of the perceived object--which ambient light conditions, as it happens, were in force during viewing of the photograph itself. Consequently, these viewing conditions help to explain why the photographic picture is misperceived as depicting a bluish sheet of paper.

As for notation, the following simplified version is appropriate for pictures: upright picture content (UPC), inverted picture content (IPC); and upright subject matter (US), inverted subject matter (IS). Or, for perceptual content in general, including the special case of perception of a picture, we may distinguish an upright aspect (UA), inverted aspect (IA); and upright object (UO), inverted object (IO). Then the case of double content ambiguity given above for retinal tokens Tr is one in which a given kind of bluish retinal image can be given an orientational double content interpretation either as an inverted aspect-upright object (IA-UO) case, or as an upright aspect-inverted object (UA-IO) case. As for a more detailed analysis involving the relevant concrete tokens as well, upright token (UT) and inverted token (IT) factors are also required. In the example given, the correct colour of the paper was white, but both the retinal token Tr and the picture P were bluish. Consequently these were both inverted token (IT) cases, or more fully, IT-IA-UO and IT-UA-IO cases.

6. The Artistic Functions of Design Content

In this section the context will be broadened to briefly consider characteristic ways in which artists can exploit orientational factors in their pictures to produce results that normally could not be seen in face to face seeing. These factors will serve to sharply distinguish artistic cases of depiction from more generic perceptual cases. Such phenomena, as previously discussed by such writers as Podro [1998] and Lopes [2005], can readily be explained in terms of Wollheim's two factor concept of 'seeing in', so it is

important that the current rival DC view should also be able to explain such phenomena, in order to effectively compete with Wollheim's influential account.

To begin, I have argued elsewhere that perceptually based double content phenomena in the arts are not confined to depictive cases of visual artworks, but instead they are endemic in all forms of art [Dilworth 2005a]. But such a synoptic approach to the arts in general thereby takes on the burden of explaining all facets of artistic meaning, not just those that are particularly salient in a given art form. In particular, it is not enough for a double content based approach to the arts to explain just standard, purely representational kinds of content or subject matter.

Instead, the approach must also explain factors in artistic meaning specifically associated with each particular art form and with artists themselves, including factors such as the idiosyncratic expressive means used by a particular artist within a given medium, more general stylistic considerations, artistic intentions and attitudes, and so on. My own attempt to explain both kinds of meaning--of both artistic and representational kinds--closely associates each kind with a particular kind of content in the two level structure. Standard representational content is explained in terms of upper level subject matter content, while broadly artistic kinds of content involving expressive artistic styles, attitudes and intentions are explained in terms of lower level aspectual content. This section will concentrate on showing how characteristic aspectual factors are recruited by visual artists for their own expressive purposes.

Fortunately, the orientational framework established in the previous sections is particularly useful in articulating specifically artistic factors in the content of depictive and other artworks. A leading idea is that an art medium, such as the medium of painting or etching, is closely associated *with its own, medium-specific orientational fields*. These fields have fully realistic intrinsic tops, such as particular shapes, textures or colours, but what makes them artistically appropriate is that their non-top elements are medium-specific. For example, an etching uses groups of purely linear elements to depict subject matter. So the subtle shading on a woman's face would be depicted using combinations of dense cross-hatched lines for darker shadings, along with looser patterns of lines to depict lighter shadings.

Nevertheless--as previously discussed--the etched linear design on the surface of a print is not intended to be a realistic picture of a woman's face as one that is covered with a network of fine lines. The lines are part of the medium, or depictive means of etching, not part of the subject matter it depicts. In orientational terms, correct perception of what an etching depicts is of form IP-IPC-US. A correct interpretation of the wiry linear elements in the concrete design understands them as expressing aspectual depictive content of form Y'(X'), with this aspectual content being inverted because its field top is etching-specific linear elements rather than the realistic facial contours that are the intrinsic top of the relevant field or fields. Hence, just as we can perceptually compensate for an oblique view of a shape by interpreting it as an IA-UO case, so also can we perceptually compensate for etching-specific aspectual content in a picture by a structurally similar IPC-US interpretation. Whereas by contrast, a naive incorrect

interpretation of an etching as depicting a woman with wiry lines on her face would be using a dual, UPC-IS orientational interpretation instead.

Artists may also use their own inimitable stylistic fields for specifically expressive purposes. For example, many of the Swiss sculptor Giacometti's wiry, vertically elongated sculptures may initially look, to the uninitiated, like three-dimensional pictures, having the outline shapes of very thin or spindly people. Arguably, however, they are instead expressive sculptural pictures of ordinary-shaped persons. Their apparent distortions in shape require an IP-IPC-US interpretation in which the prominent aspectual divergences in shape from realistic US shapes express the drained and marginalized lives of ordinary individuals in modern societies. By contrast, an incorrect IP-IPC-IS interpretation would instead interpret such works as inexpressive, pointless realistic sculptures of spindly people [Dilworth 2008].

Colour can also be used for aspectual expressive purposes, such as in the case of Andy Warhol's various serigraphs of Marilyn Monroe. One of these has a strong red colour in the area depicting her face [Dilworth 2005a, ch. 12: 258]. The uninitiated might think that this incorrectly depicts Marilyn as having a red face rather than her actual pink face-- a strange kind of deliberate error by Warhol, of orientational form UPC-IS or IPC-IS. However, much more artistically appropriate is a dual IPC-US interpretation in which the red colour has the aspectually expressive function of allowing us to see Marilyn, with her ordinary pink face, through the garish, red-light district glow of hyperactive media

publicity. Indeed, the whole series of Marilyn serigraphs, with their various colours, only makes any artistic sense if thus interpreted in IPC-US expressive colour terms.

To briefly summarize these points, the current account has a ready explanation of such cases in terms of inverted aspectual expressive content (IPC). Since this aspectual content is inverted, it is not realistic and hence does not resemble the object seen face to face. But at the same time, it is not seen as being visually inconsistent with the correct subject matter content either, since it is perceived as expressive pictorial (IPC-US) rather than as incorrect subject matter content (UPC-IS). This seeing of the red-faced Marilyn picture is not like seeing Marilyn face to face, but nor is it like incorrectly seeing Marilyn as having a red face. Instead, the potent alchemy of aspectual inversion enables us to see Marilyn as viewed via Warhol's idiosyncratic expressive vision.

7. Other Conceptions of Pictorial Content

The current double content (DC) theory of depiction distinguishes aspect representation and content from object representation and content. It postulates that this specific representational or content-based distinction is applicable both to perceptual content generally, and to depictive content in particular. However, there are various other extant distinctions of kinds of representation or content in the literature. This section surveys some of these, with an emphasis on distinctions that are particularly relevant to depiction and to the current collection. In order to avoid any invidious comparisons, after a brief discussion of Wollheim's seeing-in view the other views are presented alphabetically by

the names of their proponents. The comments are brief because their main purpose is to distinguish other content distinctions from those of the current DC view.

7.1 Wollheim

Recall that Richard Wollheim [1968, 1987] introduced the two factor seeing-in view which is taken as basic by the other participants, according to which seeing of both the physical design of a picture and its content typically occurs during seeing in. Though the DC view denies Wollheim's claim that the physical design as such of a depictive token is seen during depictive seeing, it is possible to sympathetically reinterpret his view so as to render it consistent with the DC view, using as a basis other views of his concerning artistic media that license such a reinterpretation [Dilworth 2005a, ch. 6]. Nevertheless, for present purposes it is simpler to regard Wollheim's views of content as directly entailing a theory of depiction that is incompatible with the DC view.

7.2 Hopkins

Robert Hopkins in his [1998, ch. 6] suggests that in some cases what can be seen in a picture is separate from what the picture depicts. An example he gives [p. 128] is of a photograph of a dog. If the photograph is over-exposed, what is seen in it might be a greenish dog, because the over-exposure has resulted in an overall greenish tint to the picture--even though what the photograph depicts is a dog of no particular colour. Hopkins also argues that in part the separation occurs because there are standards of correctness for what a picture depicts, but not for what can be seen in the picture.

Arguably this general distinction--of what can be seen in a picture from what it depicts--is a legitimate one for the following reason--though it is not one that is explicitly discussed by Hopkins. Wollheim [1987: 67-71] has argued that any general account of representation, including of iconic kinds such as depictive representation, must acknowledge that it is possible to represent objects that are not particular actual objects--such as a painting of a lake that depicts no particular actual lake. In such a case, one could see the lake in the painting, even though there is no actual lake depicted by the painting, and hence no standards of correctness for what it depicts.

I have suggested elsewhere that in order to explain such cases, a general distinction is needed between cases of *external* representation--in which some actual object is represented--versus cases of *internal* representation, in which some object is represented, whether or not the representation also externally represents an actual object [Dilworth 2005a, ch. 11]. As applied to Hopkins' greenish dog, a greenish dog can be seen in the picture because the concrete depictive token does indeed internally represent a greenish dog, even though the depictive token does not externally represent a greenish dog.

Perhaps there are other ways to defend Hopkins' distinction, including his own explanatory framework, but the main point is that his distinction is indeed a legitimate one that is distinct from, and hence does not conflict with, the DC distinction.

As for Hopkins' current contribution in this volume, he makes a good case that unitary, rather than divisive, accounts of seeing in have much to recommend them. To that

extent, his analyses support both his own unitary theory, and the current unitary DC theory as well.

7.3 Kulvicki

In his book [2006], John Kulvicki distinguishes, roughly following Haugeland [1991], the *bare bones* from the *fleshed out* content of an image, and a similar distinction plays a more minor role in his current contribution (this volume). However, there are two significant problems in attempting to compare his two kinds of content with my aspectual and subject matter kinds of content. Firstly, on my account it will typically be possible to find more than one depictive double content interpretation of a given depictive token. But those alternative interpretations would not typically differ in level of specificity or richness as do Kulvicki's two kinds of content.

Secondly, Kulvicki's interesting claim in the current volume--that there is limited diversity and competition in viewers' interpretive schemes, which fails to match the diversity in picture production schemes--can be directly questioned in various ways such as the following. First, since depictive tokens always permit dual or ambiguous interpretations, those interpretations do semantically compete with each other. And second, unusual production schemes, such as those of fisheye photography, do produce legitimate pictures and depictive interpretations--pace Kulvicki--in that they involve specialized shape fields that define the relevant art medium. Arguably a bulgy fisheye picture of a lemon is just as much a picture as is a wiry etching of the same lemon, once their double content structure is properly understood.

7.4 Brown

John Brown's pioneering investigations, in addition to being of significant value in their own right, could be regarded as providing further evidence for the current DC theory.

Brown is able to successfully show, via further development of Hopkins' concept of separation, that standard conceptions of seeing-in are overly narrow and simplistic, in that they neglect various kinds of possibilities of separation between what can be seen in a picture and its official subject. From the perspective of the DC theory, arguably such additional cognitive possibilities are only to be expected, for reasons such as the following.

If the DC orientational theory is correct, specific artistic media involve socially constructed, artistically relevant fields, which are arrived at in the following way. Groups of artists experiment with specific media, such as the medium of pen and ink drawing, and--in collaboration with informed audiences--discover how pen and ink marks may be used, both to depict people and places, and to express an artist's own attitudes and ideas about those subjects. However, what this means is that the relevant cognitive fields are *social* constructions--the products of whole cultures over many years--whose depictive functioning inevitably outruns the specific intentions of any one artist as applied to any one specific drawing. Consequently, precisely because the social functioning of a field thus outruns individual intentions, that artist is unable to fully communicate, to the

viewer, exactly how each penned mark is intended by him to characterize his official intended subject in the current particular picture.

Nevertheless, according to the DC theory, the lower-level pen marks are inevitably capable of representing *some* higher-level subject matter in *some* way. But since not all of this higher-level content can be attributed to the official subject--because of the intentional gap between social meanings of fields and individual artistic intentions--the DC theory predicts that there would, in general, be surplus higher-level represented content that would piggyback on content associated with the officially intended subject. But this widespread, separated surplus higher-level content or meaning is exactly what Brown's insightful investigations have uncovered. Also, presumably artists have learned to control--via prior observations of their own earlier pictures--these additional kinds of higher-level content for further artistic purposes of their own. And, to conclude, my thanks to editors Catharine Abell and Katerina Bantinaki for their very helpful comments on earlier versions of this essay.

Western Michigan University

References

- Dilworth, John 2005a. *The Double Content of Art*, New York: Prometheus Books.
- Dilworth, John 2005b. The Double Content of Perception, *Synthese* 146: 224-243.
- Dilworth, John 2005c. The Twofold Orientational Structure of Perception, *Philosophical Psychology* 18:187-203.
- Dilworth, John 2005d. The Perception of Representational Content, *British Journal of Aesthetics* 45: 388-411.
- Dilworth, John 2005e. A Double Content Theory of Artistic Representation, *The Journal of Aesthetics and Art Criticism* 63: 249-260.
- Dilworth, John 2007. In Support of Content Theories of Art, *Australasian Journal of Philosophy* 85: 19-39.
- Dilworth, John 2008. The Abstractness of Artworks and its Implications for Aesthetics, *The Journal of Aesthetics and Art Criticism* 66: 341-353.
- Gombrich, Ernst 1962. *Art and Illusion*, 2nd ed. London: Phaidon Press.

Haugeland, John 1991. Representational Genera, in *Philosophy and Connectionist Theory*, eds. W. Ramsey, S. Stich, and D. Rumelhart, Hillsdale, NJ: Lawrence Erlbaum.

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

Kulvicki, John 2006. *On Images*, Oxford: Clarendon Press.

Lopes, Dominic McIver 2005. *Sight and Sensibility*, Oxford: Clarendon Press.

Podro, Michael 1998. *Depiction*, Cambridge, Mass.: Harvard University Press.

Purves, Dale, & Lotto, R. Beau 2003. *Why We See What We Do*, Sunderland, MA: Sinauer Associates.

Wollheim, Richard 1968. *Art and its Objects*, Cambridge: Cambridge University Press.

Wollheim, Richard 1987. *Painting as an Art*, London: Thames & Hudson.