Visual Form, Attention, and Binocularity[*]

March 10, 2005

And nothing *in the visual field*
allows you to infer that it is seen by an eye.
—Wittgenstein 1921/1974, 5.633

Fill my eyes with that double vision;
No disguise for that double vision.
—Foreigner, "Double Vision"

1 What is Visual Form?

Suppose that one sees a penny presented head-on. Now suppose that one tilts the penny through an angle of 60 degrees, so that it is presented to one obliquely.

One's visual experience will change. One change is that the penny looks to be presented obliquely, where it once looked to be presented head-on. This is a change in the ostensible three-dimensional spatial properties of the penny: certain parts of the penny look to have become more distant from one; other parts look to have become closer to one.

But another change involves two-dimensional spatial properties one (at least ostensibly) experiences. One way to describe the change is to use the terminology of *presentations in a visual field*, as follows: prior to the tilting, the presentation of the point at the center of the penny's circular facing surface is equidistant in the visual field from every presentation of a point along the edge of the penny. After

[*][Acknowledgements]

the tilting, this is no longer so: rather, there are two presentations of points along the edge of the penny which are closer in the visual field to the presentation of the point at the center of the penny than any other such point; and two presentations of points along the edge of the penny which are more distant in the visual field from the presentation of the point at the center of the penny than any other such point; and the distance in the visual field between the former is half the distance in the visual field between the latter. My central purpose in this essay is to describe the underlying ontology of the discourse of presentations located in a visual field: I will assume that the discourse is well enough understood at a pretheoretic level to permit engagement with this purpose.

The two questions that need to be answered are: what are presentations? and what is it for two presentations to be at a certain distance in the visual field? A visual field in which presentations are arranged is a structured entity, with a spatial structure. A structure is an arrangement of elements in a form; so the two questions can be taken as concerned with the nature of the *elements* of such a structure, and with the *form* of the structure, respectively. Accordingly, I will sometimes refer to the first question as the *question about elements* and the second as the *question about form*.

Two leading contenders in the philosophy of perception are *intentional theory*, on which the only properties one experiences in a visual experience are external properties (such as surface colors, the the shapes and locations of the regions of physical space objects do or could occupy) ostensibly instantiated in entities surrounding one; and *sense-datum theory*, on which some of the properties one experiences in a visual experience are properties *internal* to one.[1] These views are allied with certain

[1]For the sense-datum theory, see Baldwin 1992, Boghossian and Velleman 1989, Jackson 1977, Peacocke 1983, Perkins 1983, and Robinson 1994; for intentionalism see Byrne 2001, Chalmers 2005, Clark 2000, Dretske 1995, Harman 1990, Hilbert and Kalderon 2000, Horgan and Tienson 2002, Kriegel 2002, Levine 2003, Loar 2003, Lycan 2001, Rey 1998, Searle 1990, Shoemaker 1994, Siewert 1998, Strawson 1994, Thau 2002, Thompson 2003, and Tye 2000. A position allied with intentionalism for present purposes is disjunctivism (see Martin 2002); Sturgeon's "intentional trope theory" (Sturgeon, 2000) also fits the bill. Discussions of intentionalism sometimes distinguish "Russellian" and "Fregean" variants. I am not certain what this distinction amounts to; maybe

answers to the questions about form and elements.

Consider the intentionalist's position on these questions. Since only properties ostensibly external to the subject are experienced in a visual experience, a visual experience is an experiencing as of token qualities of surfaces arrayed in various directions in an egocentric space. These qualities might not in fact exist at all, if the experience is not veridical. For this reason, the token qualities should be taken to be *intentional*: potentially existing, potentially non-existent. The intentionalist's answer to the question about elements is therefore that presentations are always intentional tokens of external properties. I will call this thesis *element intentionalism*.

Concerning the question about form, it should not be assumed at the outset that there is only one system of spatial relationships, namely those concerning distances between points or regions of physical space. Spatial notions are highly abstract, and can apply to a wide range of systems of relationships among things, so long as the systems satisfy certain mathematical properties.[2] It is therefore open to debate whether spatial relationships in the visual field pertain to physical space. The following remarks on the nature of the visual field are illustrative of a likely intentionalist answer to the question about form, on which they do so pertain:

> In visual science, "the visual field" (of a normal human) is usually taken to be the region-type within which a visual stimulus can be detected, keeping the eye fixed, or the token such region for an individual subject at a particular time. Measured from central fixation, the monocular visual field is about 160 degrees in width and 135 degrees in height. Alternatively, **the visual field (for a subject at a time) may be taken to be** the scene before the subject's eyes at that time, or **the scene-as-represented by the subject's visual experience at that time**. (Byrne, 2001, fn. 36, my emphasis)

The emphasized passage suggests an answer to the question about form on which for two of these intentional tokens to be a certain distance from one another in

my argument is only against the Russellian approach.

[2]Namely, there is a distance function δ from pairs of things into positive real numbers meeting the conditions that $\delta(x, x) = 0$, $\delta(x, y) = \delta(y, x)$, and the triangle inequality. Since functions come cheap, the function should also capture some natural relationship among the things.

the visual field is for there to be a certain visual angle between the directions in egocentric space in which the intentional tokens ostensibly lie—ostensibly, because the apparent positions of the tokens might not be their genuine positions. I will call this thesis *form intentionalism*.

By contrast, a sense-datum theorist's answer to the question about elements is that at least some presentations are tokens of *internal* qualities: features of the brain, or of the mind, where these latter may be either reducible or irreducibly "dualistic". And the answer to the question about form is that for two presentations to stand at a certain visual distance from one another is for them to bear a certain neural or mental relation to one another. I will call these positions *element anti-intentionalism* and *form anti-intentionalism*.

An important proviso is required if mentalist versions of element and form anti-intentionalism are to be nontrivial competitors with element and form intentionalism: the relevant mental quality tokens must not be tokens of intentional properties (such as the property of being intentionally directed toward redness), and the relevant mental relations must not be tokens of intentional relations (such as the relation such that for two intentional tokens of redness to bear it to one another is for one's visual experience to represent them as located at a certain visual angle with respect to one another). They must rather be, for instance functional but not intentional properties or relations, or irreducibly dualistic non-intentional properties or relations.

The intentionalist's and sense-datum theorist's pairings of answers to the questions about form and elements are pure views. Combining element anti-intentionalism with form intentionalism yields a form of *projectivism*, on which experience mistakenly locates internal quality tokens in the two-dimensional space of egocentric directions. Alternatively, element intentionalism can be combined with form anti-intentionalism to yield a view with a *Kantian* flavor, on which experience presents ostensible external entities, but is structured in an ultimately nonrepresentational

fashion.[3]

A vast contemporary literature addresses the debate between element anti-intentionalism and element intentionalism (see Block 1990; Chalmers 2005; Harman 1990; Lycan 2001; Martin 2002; Shoemaker 1994; Thau 2002; Tye 2000), to which I won't add in this essay. By contrast, the debate between form anti-intentionalism and form intentionalism has received comparatively little attention. This debate is the focus of this essay. My aim is primarily to argue for form anti-intentionalism, and secondarily to propose and defend a specific version of form anti-intentionalism, according to which relations of visual distance stem from nonrepresentational features of a faculty of selective attention.

* * *

In addition to bearing on an overlooked aspect of the debate between sense-datum theory (and its ally the Kantian view) and intentionalism (and its ally, projectivism), which answer the question about form receives bears on two further issues of central concern in the philosophy of perception and consciousness.

The first is the status of a *metaphysical argument against sense-data*. A familiar concern about classical sense-datum theory runs as follows:

> If a datum seems F, it is F. So since the datum of the tilted penny seems elliptical, it is elliptical. Since sense-data are "internal", they are either in the brain of the subject or in the mind of the subject. But it is absurd to suppose that seeing the tilted penny requires the existence of something elliptical in one's brain. So they must be dualistic purely mental entities in some private space. But that's also absurd! (cf. Broad 1923, 118; Chisholm 1950, 170–1; Price 1950, 128; Crane 1992; Huemer 2001, 149–68.)

Clearly, whether this argument goes through depends on the nature of visual field distance. Maybe once the nature of visual field distance relations is uncovered, it

[3]For projectivism (not necessarily of the sort I describe here), see Perkins (1983) and Boghossian and Velleman (1989, 1991). I don't know of any contemporary philosopher endorsing the Kantian view, though see [***setiya].

won't seem absurd to suppose that visual field distance relations hold among entities in one's brain or mind.

The second is the status of *representationalist* views in the philosophy of consciousness, according to which to be a phenomenal character (what an experience is like to its subject) is to be a certain sort of representational property. If form anti-intentionalism is correct, then, although the notion of a representational property is perhaps somewhat obscure, it is likely that it does not apply to the property of an experience of involving a visual field in which p and p^* are at distance d, since something's having it involves two things standing in a nonrepresentational relation to one another. And this property is, to my mind, an extremely good candidate for a phenomenal character.

* * *

Road map: in section 2, I will characterize a non-mental, non-intentional relation of *attentional spatial distance*, which I will later go on to argue is identical to distance in the visual field. In section 3, I will address the central case for intentionalism, arguing that while it may support element intentionalism, it does not support form intentionalism. In section 4, I will discuss a case against intentionalism on the basis of double vision, arguing that discussion of this issue in the literature leaves the issue unresolved. Finally, in section 5, I will provide a sort of "spectral inversion"-like case involving double vision which cuts against form intentionalism; I will also assess how this case shapes up against more familiar spectral inversion arguments against element intentionalism.

2 Attentional Space

The notion of *attention* is used in a broader and a narrower sense. Consciousness plausibly comes in degrees: if one sees x and y, but thinks about x and not y, one is more conscious of x than of y. Both senses of the notion of attention concern bringing something to consciousness in the highest degree. The broader sense has

no further connotations. Since one way to bring something to consciousness to the highest degree is to think about it, among the relations one may bear to x and thereby count as attending to x in the broad sense is *thinking about* x.

The narrower sense has connotations concerning the manner in which one makes x conscious to the highest degree. While bringing something to the highest degree of consciousness doesn't require its being perceived. Still, when one perceives something, one is thereby in a position to bring it to the highest degree of consciousness, by training on the object a selective faculty distinctive to perception: a sort of "spotlight" or "zoom lens". This spotlight need not (and in any normally rich experience, cannot) be trained on everything one perceives: rather, it selects from among all that is perceived a small number of entities to promote to the maximal level of consciousness. The spotlight is moveable: in a visual experience, it can highlight presentations at the top of the visual field or at the bottom; at the left, the right, or the center; or anywhere in between. Moreover, it is resizable: it can be focused broadly, to encompass presentations widely spread in the visual field at a low level of detail, or narrowly, to encompass presentations bunched closely in the visual field at a high level of detail. To attend to something in the narrower sense, then, is to promote it to the maximal level of consciousness by training this attentional spotlight on it.

The attentional spotlight can be *moved* and *grow* or *shrink*: movement is change of position, growth and shrinkage are changes of size. Clearly such changes of position and size require positions and sizes/distances, and hence a space within which things are located and bear distance relations to one another. I will call this space *attentional space*.

Suppose that two presentations are in the same region of attentional space. Then, if the attentional spotlight is focused in exactly that region, and the total visual field is not too complex, the subject can focus attention on those two presentations simultaneously. It follows that if two presentations are co-located in a point-sized region, so that the distance between them is zero, then if the attentional spotlight

is shrunk to a point, is focused in that point-sized region, and the total visual field is not too complex, the subject can focus attention on those two presentations simultaneously. Conversely, if the subject can focus attention on two presentations simultaneously with the attentional spotlight shrunk to a point, it would seem that the distance in attentional space between the two presentations must be zero.

I will defend *attentional space theory*: the view that for two presentations to stand at a certain visual distance from one another is for them to stand at a certain distance in attentional space.

Relations between presentations such as *standing in distance d in attentional space* are clearly mental relations. Moreover, they are not, *prima facie*, intentional relations. The attentional spotlight is related to intentionality as a sort of conduit or sluice for information; if so, the fact that presentations p and p^* are at a certain distance in attentional space seems to be primarily a matter of what one must do in order for information concerning both p and p^* to enter this sluice. Facts about how information is processed by a system are not themselves facts about the informational properties of the system: moving the linguistics books and the philosophy books to the same floor of the library might make life easier for the philosophers of language, but has no influence at all on the total content stored in the library.

Is attentional space theory committed to bizarre metaphysical views about private spaces? Not obviously. If attention is reducible at all, attentional space will be reducible, e.g., by regarding attention as a *mere* informational sluice and by regarding attentional spatial positions as ways this sluice may interact with various data structures. Nonreductivists about consciousness may regard the attentional spotlight as irreducible; if they're correct, it's already curtains for reduction, so no harm in adding in irreducible spatial features.

Is attentional space theory a version of anti-intentionalism? *Ultima facie*, it might turn out that there are strong metaphysical ties between attentional spatial relations and certain representational relations, so that the former and the latter

turn out identical (for instance, if relations are individuated modally). Strengthen the ties enough, and a position with which intentionalists would be happy might emerge. Of course, it is also compatible with attentional space theory that there are not such ties. This is the view I will eventually defend.

3 Attention and the Transparency Case for Intentionalism

The first stage in my defense of this view will be to undermine a case that many take to support intentionalism *tout court*. I will argue that it at best supports element intentionalism.[4]

According to Harman (1990, 251), "[t]he sense-datum theorist's view about our immediate experience of color is definitely not the naive view; it does not represent the viewpoint of ordinary perception. The sense-datum theory is not the result of phenomenological study". This is because

> When you see a tree, you do not experience any features as intrinsic features of your experience. Look at a tree and try to turn your attention to intrinsic features of your experience. I predict you will find that the only features there to turn your attention to will be features of the presented tree.

Let *transparency* be the claim that one can turn one's attention in visual experience only to (ostensible) external qualities: this is the first premiss of the argument. The implicit second premiss seems to be that if one experiences a quality in a visual experience, one can turn one's attention to it. These premisses imply that one only experiences external qualities in a visual experience; if so, intentionalism is true.[5]

[4] I don't think it supports this position either: see [AUTOCITE].

[5] Philosophers following Harman in advancing this style of argument include Byrne (2002), Dretske (1995), Hilbert and Kalderon (2000), Levine (2003), Loar (2003, in preparation), Lycan (2001), Martin (2002), Shoemaker (1991, 1994, 2001), Thau (2002), and Tye (1992, 1995, 2000).

Since there are two notions of attention, there is a concern about equivocation.[6] Suppose the argument runs on the broad notion of attention. Then the second premiss is plausible: if one experiences a quality in a visual experience, one can think about it. But the first premiss is not plausible: one can certainly *think about* an internal mental feature during a visual experience, and in so doing one would thereby become maximally conscious of it.

Now suppose that the argument runs on the narrow notion. Maybe the transparency claim—that one cannot turn one's attentional spotlight on any internal quality—is plausible, maybe it isn't.[7] Then why grant the second premiss in the argument, that for any quality one experiences in a visual experience, one can attend to it—can train one's attentional spotlight on it? It's not in general the case that the only qualities one experiences *outside* of visual experience are susceptible to having the visual spotlight turned on them. After all, one cannot turn a visual spotlight on the felt quality of emotions, or the imagined sound of one's own voice in linguistic thought, or on the experienced phenomenal quality of trying to do this or that (on "tryings", see Peacocke MS).

I think in addition that even in visual experiences, one experiences properties on which one cannot turn one's visual spotlight. Consider spatial features of the attentional spotlight itself, such as the property of being in motion or at rest in the visual field; the property of increasing or decreasing in size; the property of being shrunk to a point, or being focused more or less broadly. One seems to experience these properties. At least, one seems to know that one's attentional spotlight has these properties, and this knowledge does not seem to be of a purely theoretical character: rather, it seems to result from experiencing of these qualities.

I conclude that Harman's argument fails due to equivocation: if 'attend' is read broadly, the first premiss is not plausible; if narrowly, the second premiss is not plausible.

[6] Stoljar (in preparation) endorses a similar equivocation line against a use of transparency to oppose "qualia", intrinsic properties of experiences of which one is "directly aware".

[7] I don't find it plausible: see [AUTOCITE].

4 Double Vision

My ultimate complaint against form intentionalism is that it cannot adequately explain *double vision*. In this section, I will discuss a debate over whether intentionalism is compatible with double vision between Boghossian and Velleman (1989) (no) and Tye (2000) (yes): while Tye's reply to Boghossian and Velleman fails, a souped-up intentionalist view affords a reply. In the following section, I will argue against the souped-up view. I begin with some preliminaries.

4.1 Definitions and an Example

Let a *centered oriented space* be a space with certain origin and axes, perhaps one which is not stationary with respect to "objective" space. For instance, let *car-space* be the centered oriented space with its center as the point at the center of my car and its forward, up, and left directions as the mutually orthogonal rays extending from this center point toward the front of the car, toward the roof of the car, and toward the left side of the car. So long as I am sitting in my car's driver seat, I am motionless (or moving only trivially) with respect to car-space, but if I am driving, I am moving nontrivially with respect to physical space (and hence, so is car-space).

Next, let *C-space* (for a particular person) be the centered oriented space with its center as the point directly between that person's eyes and its forward, up, and left directions as the mutually orthogonal rays extending from this center point forward from the person's face, etc. Let a *C-direction* be a ray extending from the center of C-space, and the *C-distance* of an object the distance along a C-direction to the object from the center. The view described by Byrne, discussed at the top of this paper, can then be interpreted as the view that points of a person's visual field are directions in that person's C-space, and distances in the visual field are visual angles between directions in that C-space.

I will be referring to the following *Double Vision Example* later in the paper. Suppose that before one, located on the C-direction *forward*, receding from one are

a smallish pink heart, yellow moon, and orange star; that a grey background lies behind the three objects, and that both eyes are pointed at the yellow moon. If one is a normal observer, one's visual field will be symmetrical: a moon-shaped region in its middle will be occupied with presentations of yellow; flanking this region to the left and the right will be regions identical to one another in character. Perhaps the best way to describe these regions is as having a shape like that of a heart overlaying a star, with every subregion of the region of overlap between the heart and the star containing presentations of both pinkness and orangeness, and every subregion of the region of non-overlap containing either presentations of both pinkness and greyness or presentations of orangeness and greyness. The rest of the visual field contains presentations of greyness.

I don't claim that this "color mosaic" description *exhaustively* characterizes one's visual field: it may well leave out representation of C-distances, as well as of such "gestalt" properties as shape, as well as of different presentations seeming to inhere in objects which are integrated wholes. But it is at least a true *partial* characterization of one's visual field.

Having an experience involving this visual field, it will look to one as if located on the C-direction *forward*, receding from one are a smallish pink heart, yellow moon, and orange star, against a grey background. Things will not, it seems obvious, look to one as if there are *two* hearts or stars in front of one, or as if any direction significantly off of C-*forward* is occupied with anything with a chromatic color. One's experience will therefore be veridical.

4.2 A Preliminary Skirmish Over Double Vision

Contention over double vision is a long-term mainstay of the philosophy of perception literature. For instance, Hume (1739/1978, 210–1) argues that double vision refutes naive realist theories of perception on the grounds that

> [w]hen we press one eye with a finger, we immediately perceive all the objects to become double, and one half of them to be remov'd from their

> common and natural position. But as we do not attribute a continu'd existence to both these perceptions, and as they are both of the same nature, we clearly perceive, that all our perceptions are dependent on our organs, and the disposition of our nerves and animal spirits.

Broad (1923, 96; 1925, 181) and Price (1950, 28, 57, 63, 224) argue that double vision supports sense-datum theory. Barnes (1944/5, 162–3) replies that a "multiple relation theory of appearing" can treat double vision without sense-data; Jackson (1977, sec. 4.5) rejoins that it cannot.

More recently, Boghossian and Velleman (1989, 92–3) have argued that double vision is incompatible with intentionalism:

> If you press the side of one eyeball, you can see this line of type twice without seeing the page as bearing two identical lines of type. [...] Similarly, you can see nearby objects double by focusing on distant objects behind them, and yet you cannot get yourself to see the number of nearby objects as doubling. [...] None of these experiences can be adequately described solely in terms of their intentional content. Their description requires reference to areas of color in a visual field, areas that split in two [...] without anything's being represented to you as doing so. (92–3)

Consider the Double Vision Example. It doesn't look to one as if there are two pink hearts, or two orange stars, located before one. And it doesn't look to one as if the pink heart is either to the left of C-*forward* or to the right of C-*forward* (although it looks to have *parts* on both sides while also having a part on C-*forward*); nor does it look to one as if the orange star is either to the left of C-*forward* or to the right of C-*forward*. Rather, it looks to one as if there is one pink heart and one orange star, both of which are located on C-*forward*. So, Boghossian and Velleman seem to be suggesting, nothing about how things look to one is available as plausibly identifiable with such properties of one's visual field as its containing two heart-shaped regions in which pink is presented or two star-shaped regions in which orange is presented.

This last step presupposes that whenever the ostensible fact that p about C-location is part of the content of one's visual experience, it looks to one as if p. Tye (2000, 88–9) responds, so far as I can tell, by rejecting this presupposition:

> In the epistemic or conceptual sense of the term 'appears', it does not appear that the number of lines has doubled. It does not look as if the number of lines has doubled. But phenomenologically, there is, of course, a conspicuous change. This [...] is because at the level of the grouped array, there is a change in representational content. [...] Where this case differs from that in which one sees two identical lines of type *as* two such lines [...] is in a much higher conceptual layer of content. (89)

The "grouped array" is

> a representation similar in character to Marr's $2\frac{1}{2}$-D sketch ["a vital foundation for further visual processing"] [...] [with] a matrix-like structure, the cells of which are dedicated to particular lines of sight. Within each cell are symbols for various local features of any surface at that position in the field of view (for example [color]). (71)

Tye seems to be suggesting that while it does not look to one as if two two orange stars and two pink hearts (overlapping one another) flank a yellow moon, one still visually represents this, since the representation is only at the level of the grouped array, and the content that p can be represented at this level without that content's looking to one to obtain.[8] This seems to involve two levels of consciousness of representational contents: an *intermediate* level of consciousness at which the content influences the character of one's visual field, and a *higher* level at which the content in addition captures how things seem to one. The grouped array's contents are only at the intermediate level; the conceptual contents are at the higher level.

The success of Tye's solution hinges on an empirical claim, that his "grouped array" is actually part of human visual processing. Tye seems to suggest as support for this empirical claim that Marr's theory, as adumbrated in Marr 1982, is committed to the existence of the grouped array. I'll grant that if Marr's theory is committed to something like the grouped array, Tye's solution rests on an adequate empirical foundation. Otherwise, though, Tye's solution is more appropriately viewed as at best a speculative proposal which might save intentionalism if science goes the right way.

[8] Tye's approach is endorsed by Byrne (2001) and Chalmers (2005).

Unfortunately, Marr's theory does not seem to be committed to a representation like Tye's grouped array. To show this, I will sketch Marr's theory of binocular vision. To begin with, note that binocularity manifests in two ways:

> The experiential difference between monocular and binocular vision is independent of the double images of unfocussed objects produced by binocular vision. The extra way depth is indicated in binocular vision is present when you look in a child's stereoscope, and there need not be any double images when you do. (Peacocke, 1983, 13)

Binocularity begins with monocular representations from the left and right eye. In Marr's view, the phenomenon Peacocke describes of extra depth without double images is due to the *fusion* of these representations, namely the formation on the basis of these monocular representations of a third, binocular, representation. This binocular representation is labelled with information about the distances of points on represented surfaces, which is calculated on the basis of the disparity between the distance relative to its surroundings of the left representation of the point and the distance relative to its surroundings of the right representations of the point. Marr's story about binocularity (Marr, 1982, 111–58) is exclusively concerned with how the binocular representation is calculated on the basis of the left and right representation.

When fusion occurs, the left and right representations are entirely suppressed from consciousness. Evidently, fusion does not always occur, in which case *diplopia*, or double vision, is the result: no binocular representation is calculated, and the left and right representations are not entirely suppressed from consciousness.

The left and right representations of a point are fused only when the disparity disparity is small; otherwise, diplopia results. A curved plane known as the "horopter", on which the point of intersection between L-*forward* and R-*forward* lies, contains the set of points such that there is no disparity between their left and right representations. The upshot of the requirement that disparity be small is that fusion will only result for points within a small distance of the horopter, an area

known as "Panum's fusion area". Outside of Panum's fusion area, no binocular representation is formed, and diplopia results.

This is *all* Marr says about diplopia. His theory is *silent* as to why, e.g., in the Double Vision Example, presentations of pink and orange overlap, rather than presentations of pink and yellow overlapping, or why they fit together into a unified visual field, or why diplopic representations are promoted to consciousness at all (see esp. p. 145).

(Or, for that matter, why a point seen with diplopia looks to have a C-direction and a C-distance. My speculative views on this matter are that positions in the visual field are somehow associated with C-directions; when x is seen with diplopia, x looks to be in the C-direction associated with the position of the visual field between the two positions the presentations of x occupy. The mathematics determining which C-distance x looks to be at are slightly intricate and need not detain us here (see Marr 1982, 155–6), but the rough idea is that the C-distance x looks to be at is further from the horopter, the farther apart the presentations of x are.)

In particular, it is no part of Marr's view that a single binocular representation is formed with the false content in the Double Vision Example that two pink heart-shaped regions are before one, each of which overlaps an orange star-shaped region. There is thus no basis for Tye to claim empirical support for this aspect of his story, which is essential to his treatment of double vision. Tye's proposal is therefore highly speculative.

Worse still, Tye's speculative proposal is a particularly bad fit with the overall Marrian approach. His false binocular representation needs to be formed by integrating true monocular representations. This step would seem to violate Marr's "principle of least commitment", a principle that "guide[s] the design of algorithms, and [...] probably ought to be satisfied by any serious candidate for an early visual process in the human visual system", namely "not doing something that may later have to be undone" (Marr, 1982, 106). Since forming Tye's false representation provides no new information by which the fused representation can be calculated

beyond that which is already in the monocular representations, there seems to be no purpose in generating it; and since it does not by itself influence how things look, the representation needs to be suppressed from consciousness, thereby undoing the formation of the representation. By Marr's lights, then, the formation of Tye's grouped array is not a "serious candidate for an early visual process in the visual system".

(Treating fusion would considerably complicate the following discussion. Henceforth, I will ignore it, assuming that all binocular vision involves diplopia.)

4.3 Dual-Monocular Form Intentionalism

Tye goes wrong, I believe, because he assumes that the only space which binocular vision represents things as located in is C-space. This assumption is understandable: the only space anything ever *looks* to be in, in binocular vision, is C-space. However, Tye if anyone should recognize that things can be visually represented to be a way they do not look to be: the representation might be at Tye's intermediate level of consciousness.

In *monocular* vision, things do not look to be in C-space, but rather in either of two monocular spaces, which I will call *L-space* and *R-space*. These are centered oriented spaces, with the center of L-space (R-space) at the center of the left (right) lens or retina, and axes defined so that L-*forward* (R-*forward*) points straight out of the left (right) eye.[9] These spaces are represented by left (right) monocular representations. They move with respect to one another when one's eyes swivel in their sockets: look at something near, and L-*forward* and R-*forward* cant inward to intersect at a point closer to one; look at something far, and they cant outward to intersect at a more distant point.

Close the left eye, and one experiences a shift of perspective such that things look to be in R-space; open it, and perspective shifts back, so that things no longer

[9]That is, is on a line passing through the center of the retina, the center of the lens, and the center of the pupil.

look to be in R-space but now look to be in C-space. Evidently, binocular vision not only provides for representation of a space not represented in monocular vision in which things look to be located, but also interrupts the representation of monocular spaces, so that things do not look to be located in them. It seems that representation of monocular space is at Tye's higher level of consciousness during monocular vision, but is pushed down to Tye's intermediate level of consciousness during binocular vision.

Recognizing representation of L- and R-space affords a intentionalist analysis of the Double Vision Example. Assume that there is representation that pinkness is presented in a certain L-location (10 degrees, I will suppose, to the right of L-*forward*, and at an L-distance *near*) and a certain R-location (10 degrees to the left of R-*forward*, and *near*); there is representation that orangeness is presented in a certain L-location (10 degrees to the left of L-*forward*, and *far*) and a certain R-location (10 degrees to the right of R-*forward*, and *far*); and there is representation that yellowness is presented in a certain L-location (at L-*forward*, and *middle*) and a certain R-location (R-*forward*, and *middle*). These assumptions do not posit representations beyond those posited in Marr's theory, nor do they violate Marr's principle of least commitment.

During binocular vision, representation of L- and R-location is only at Tye's intermediate level of consciousness—nothing looks to one to be in an L-direction or an R-direction—while representation of C-location is at Tye's higher level of consciousness (though the influence it has on the composition of the visual field does not manifest in in the visual locations in which qualities are presented, but in much more subtle ways).

Say that two directions d and d^* *visually coincide* iff either they are both L-directions or both R-directions and are identical, or one is an L-direction and the other is an R-direction and, in binocular vision, presentations ostensibly located in d and d^* overlap. In the Double Vision Example, the L-direction in which orangeness is presented and the R-direction in which pinkness is presented visually

coincide; as do the L-direction in which pinkness is presented and the R-direction in which orangeness is presented; as do the L- and R-directions in which yellowness is presented.

Let *dual monocular form intentionalism* be the view that locations in the visual field are identical to pairs of visually coinciding monocular directions: for two presentations to be in the same visual field location is for the monocular directions they are on to visually coincide; if two presentations are both ostensibly located on L-directions (R-directions), for them to be at a certain visual distance is for the angle between these L-directions (R-directions) to have a certain value; if one presentation is ostensibly located on an L-direction d and another presentation is ostensibly located on an R-direction d^*, for them to be at a certain visual distance is for the L-direction d^{**} visually coincident with d^* to be such that the visual angle between d and d^{**} has a certain value.

5 Attentional Shear

The dual monocular form intentionalist's notion of visual coincidence between an L-direction d and an R-direction d^* deserves further probing. When d and d^* are visually coincident, what makes them so? Of any two directions, it is not easy to believe that their visual coincidence should be a primitive fact.

In this section will argue that what makes d and d^* visually coincident (when they are) is the fact that there is a point in attentional space to which both bear a certain relation; but, unfortunately for the intentionalist, whether two directions stand in this relation to the same point is contingent.

Let S be the visual spatial property (instantiated in the Double Vision Example) of having a presentation of pink and a presentation of orange be co-located in the visual field. Let A be the attentional property of having a presentation of pink and a presentation of orange be co-located in attentional space. Let R be the intentional property of having a presentation of pink and a presentation of orange be ostensibly

located in visually coincident directions. In subsection 5.1, I will argue for the claim of attentional space theory S and A that are necessarily coinstantiated. In subsection 5.2, I will argue against the claim of dual monocular form intentionalism that S and R are equivalent. The latter argument involves what I will call *attentional shear*. Attentional shear resembles "spectral inversion", widely discussed in debates over element intentionalism, so in subsection 5.3, I will compare attentional shear and spectral inversion.

5.1 *S* and *A* are Equivalent

First I will argue that necessarily, if S is instantiated in an experience, so is A. Suppose for reductio that S is instantiated in Bill's experience and A is not. Since S is, presentations of pink and of orange overlap somewhere in the visual field; since A isn't, Bill can't simultaneously attend to pink and to orange with the spotlight shrunk to a point (even though the experience is not particularly complicated). Are these compatible? Suppose Bill is attending to the presentation of pink in this overlapping region. Either Bill can shift visual attention from pink to orange without moving his attentional point-spotlight, or he cannot. Suppose the former. Because he can't attend to both simultaneously, it would not seem to Bill as if pinkness and orangeness overlap in the visual field. Since the experience isn't complex, he wouldn't have the phenomenology of *hunting* for orangeness. Rather, he would more likely seem to have an experience of "disambiguating" a surface property ambiguous as between pink and orange, in much the way that when a Necker cube flips from one orientation to another, one has an experience of disambiguating the figure first in the one orientation, then in the other. Now suppose the latter: that in order for Bill to shift his visual attention from pinkness to orangeness, he must move his point-spotlight. Once again, it would not seem to Bill as if pinkness and orangeness overlap: rather, the phenomenology would be either as if presentations of pinkness and orangeness are at some distance from one another in his visual field; or as if he saw something ambiguous between pink and orange, but inducing the shift between

disambiguations requires flicking the location of the spotlight to a point elsewhere in the visual field then back to the original point.

If A isn't instantiated in Bill's experience, it certainly would not seem to Bill as if S is instantiated. One could insist that despite this, S is instantiated. But this insistence would come at the cost of admitting that spatial properties of the visual field are sometimes bizarrely inaccessible to reflection.

Second, I will argue that if A is instantiated in an experience, so is S. Suppose for reductio that A is instantiated in Bill's experience but S is not. Since A is, Bill can simultaneously attend to pink and to orange with the attentional spotlight shrunk to a point; since S isn't, pink and orange don't overlap in the visual field. But how could this be? The attentional spotlight is shrunk to a *point*. Points have no extension. There is no room in Bill's attentional spotlight for the simultaneously attended-to presentations of pinkness and orangeness to do anything but overlap. It would certainly seem to Bill as if S were instantiated. One could despite this insist that S is not instantiated; once again at the cost of a bizarre lack of transparency.

The conclusion that S and A are necessarily coinstantiated falls short of full attentional space theory: I haven't argued that visual field distances and attentional space distances are equivalent when those distances exceed zero. Unfortunately, providing such an argument would require a sharper characterization of the notion of attentional distance than I am at present in a position to provide. Still, having established the special case helps to make attentional space theory plausible, and will in addition provide crucial support in the upcoming argument against intentionalism.

5.2 S and R are Inequivalent

I will describe two subjects, Reg and Irreg, who have differing visual fields with the same representational properties. Both are in the surroundings of the Double Vision Example. Reg is a normal subject, so that his visual field fits the description I gave in discussion of that example. But Irreg is abnormal: while the presentations in his

visual field due to the stimulation of his left (right) eye (his *L- (R-)presentations*) are the same as those for Reg, the L-presentations and the R-presentations are sheared with respect to one another in attentional space: he suffers *attentional shear*.

For a loose analogy, suppose that two slide projectors are casting distinct images onto the same area of movie screen. Now suppose that the beam from one of the projectors is shifted slightly to the left while the other is shifted slightly to the right. The overall configuration will change; Reg is comparable to the "before" condition, Irreg to the "after".

I argue as follows: I will describe two ways for a visual system to work at the computational level—one for Reg, one for Irreg. I will then argue that if a person with the first visual system is normal, a person with the second suffers attentional shear, so that they differ phenomenologically: S is instantiated in Reg's experience but not in Irreg's. Then I will argue that despite this, their visual experiences could share representational contents: R is instantiated in both. Hence S and R are inequivalent. Throughout, I will grant the element intentionalist thesis that presentations are intentional color-instances.

Computation

The descriptions I will give greatly simplify the actual computational realization of experiences, but I doubt that that they distort it much. I don't intend to rule out by these descriptions that experiences involve an irreducibly mental element.

A visual experience (for Reg or Irreg) involves the tokening of a pair of matrix-like monocular representations with a distance function defined over pairs of cells in the matrix, the *L-* and the *R-representation*.[10] Each cell in the L- (R-)representation bears the relation of *responding* to a certain L- (R-)direction, such that standardly, the color of the surface in that direction causes a certain symbol to be tokened in that cell: e.g., a yellow surface in a certain direction causes a certain symbol to be

[10]I'm assuming for simplicity of exposition that there is no third binocular representation of the sort Marr describes.

tokened in the cell responding to that direction. I will assume a correlational theory of visual intentionality, so that the tokening of that symbol in the cell therefore represents that something yellow is in the associated direction, so that the symbol is a symbol for yellow. I will characterize the symbols in the cells of these matrices in terms of what properties they are symbols for, or of their intentional properties; no other classification will be relevant in this context. (This is because I am granting element intentionalism for the sake of argument.) Distances between cells in the L- (R-)representation are measured in degrees: if the L- (R-)cell c responds to the L- (R-)direction d and the L- (R-)cell c^* responds to the L- (R-)direction d^*, then the distance between c and c^* is n degrees iff the visual angle between d and d^* is. Let the *central cell* of the L- (R-)representation be the cell responding to L- (R-)*forward*.

The L- and R-representations are linked together by a device for copying symbols from the L- and R-representations to a "central executive" processor or module or "box". This device moves over the matrices and copies regions of variable size. Its state at a time can be characterized in terms of an L- and/or an R-representation cell over which it is centered, and a radius, such that it can copy the information stored in any cell within the radius distance of the central cell to the central executive. For some set of L-cells and some set of R-cells, there is a bijective function between these sets such that if the copying device is centered over a cell in one of these sets, it is also centered over the image of this cell under the bijection.

The principal computational difference between Reg and Irreg concerns which of their L- and R-cells are paired up by this bijection. The bijection for Reg is much like that for a normal person: it pairs the L-cell responding to L-*forward* with the R-cell responding to R-*forward*; *mutatis mutandis* for the rest (e.g., the L-cell responding to 20 degrees to the left of and 10 degrees above L-*forward* is paired with the R-cell responding to 20 degrees to the left of and 10 degrees above R-*forward*). However, for Irreg, the L-cell responding to the L-direction 10 degrees to the right of L-*forward* is paired by the bijection with the R-cell responding to the R-direction 10 degrees to the *left* of R-*forward*; *mutatis mutandis* for the rest (e.g., the L-cell

responding to 20 degrees to the left of and 10 degrees above L-*forward* is paired with the R-cell responding to *40* degrees to the left of and 10 degrees above R-*forward*).

Suppose that Reg and Irreg are both confronted with the scene in the Double Vision Example. In this case, their L- and R-representations will have the same character: each L-representation will consist of a central moon-shaped region filled with symbols for yellow, a star-shaped region filled with symbols for orange centered on a point 10 degrees to the left of center, and a heart-shaped region filled with symbols for pink 10 degrees to the right of center; and each R-representation will consist of a central moon-shaped region filled with symbols for yellow, a heart-shaped region filled with symbols for pink centered on a point 10 degrees to the left of center, and a star-shaped region filled with symbols for orange 10 degrees to the right of center. (I will assume that these regions are small relative to the angle 10 degrees.)

I see no reason to deny that Reg and Irreg, so described, are possible. With the aid of some assumptions about how these computational realizations are related to phenomenology and to representation, it can be argued that Reg and Irreg's visual fields are different—the spatial property S is instantiated in Reg's experience but not in Irreg's—though their experiences have the same representational content.

Phenomenology

My first assumption about phenomenology is that there is no funny business monocularly as concerns the placement of presentations in attentional space due to the placement of symbols in representations, so that the presentations and relations among them in attentional space solely attributable to the stimulation of the left (right) eye are the same for Reg and for Irreg. For instance, due to the character of his L-representation, Reg's attentional space contains, from left to right, a star-shaped region filled with presentations of orange, a moon-shaped region filled with presentations of yellow, and a heart-shaped region filled with presentations of pink, standing at a certain attentional space distance (which could be assumed harmlessly

to be a distance of 10 degrees) from one another; and so does Irreg's.

My second assumption is that the selective copying device realizes the attentional spotlight in a natural way: so that, for instance, if the selective copying device is shrunk down so that it covers a single matrix cell, the attentional spotlight is shrunk to a point; and so that, for instance, the attentional spotlight selects a presentation of yellowness just in case the copying device has a symbol for yellow within the radius distance of the cell on which it is centered.

Given these assumptions, A, and therefore S, is instanced in Reg's experience but not in Irreg's. Suppose that Reg's copying device has radius zero, and is centered over the L-cell 10 degrees to the left of center (which is occupied by a symbol for orange), and therefore also over the R-cell 10 degrees to the left of center (which is occupied by a symbol for pink), then, by the first assumption, his attentional spotlight is shrunk to a point; and, by the second assumption, he can attend both to a presentation of pink and a presentation of orange. Putting both these results together, a presentation of pink and a presentation of orange are co-located in attentional space; and this is property A.

If Irreg's copying device is centered over an L- (R-)cell containing a symbol for orange, then it is centered over an R- (L-)cell containing a symbol for grey. And if it is centered over L- (R-)cell containing a symbol for pink, then it is centered over an R- (L-)cell containing a symbol for pink. There is no way for it to be centered over a cell containing a symbol for pink and a cell containing a symbol for orange. So suppose that it has radius zero, and that therefore (by the first assumption) his spotlight is shrunk to a point. Then there is no way for it to have both such symbols within its radius. But then, by the second assumption, the attentional spotlight does not select both a presentation of pinkness and a presentation of orangeness. Since this is so no matter where it is centered, a presentation of pinkness and a presentation of orangeness are not co-located in attentional space; hence A (and therefore S) is not instanced.

There is therefore a phenomenological difference between Reg and Irreg concern-

ing the composition of their visual fields: S is instanced in Reg's, but not in Irreg's, experience.

I leave as an exercise to the reader the task of showing that Irreg's visual field is centered on a heart-shaped region filled with presentations of pink; flanking the center on each side, 10 degrees off center, are moon-shaped regions filled with presentations of yellow overlapping presentations of grey; flanking the center on each side, twenty degrees off center, are star-shaped regions filled with presentations of orange overlapping presentations of grey; presentations of grey everywhere else.

Representation

Five sorts of properties are candidates for representation in Reg's and Irreg's experiences: qualities, such as colors and (perhaps) shapes; L- and R-locations, which divide into L- and R-directions and L- and R-distances, which I will treat simultaneously; and C-locations, which divide into C-directions and C-distances, which I will treat separately. I will argue that they differ in representation of none of these.

Qualities I can't see how to deny that for both Reg and Irreg, the pink heart looks pink and cardioid; the yellow moon looks yellow and lunar; the orange star looks orange and sidereal.

L- and R-locations For neither does anything look to have any L- or R-location. The representation is at Tye's intermediate level of consciousness. There is therefore no clear way to motivate anything beyond a simple correlational theory of reference as applying here, so that which L- (R-)location the heart, moon, and star are represented as being in is fixed by the L-location responded to by the cells that represent them. These patterns of responsiveness are the same for both, so that the represented L- and R-locations are the same for both.

C-directions My claim is that for both Reg and Irreg, the pink heart looks to be located in C-*forward*; as does the yellow moon; as does the orange star.

It's easy to show this for Reg. Reg is just like a regular person, and if a regular person were having an experience like Reg's, this is how things would look for him.

Which C-direction(s) do the three objects look to Irreg to be located on? Admittedly, it's obscure how to answer this question. Still, two assumptions about the C-direction an object looks to be on when it is seen with diplopia seem natural: first, that it looks to be on the direction associated with the region of attentional space directly between the two regions in which the two presentations are located; and second, that the center of attentional space is associated with the direction C-*forward*. Given the composition of Irreg's attentional space, it follows by the first assumption that each of the objects looks to be on the direction associated with the center of attentional space; and it follows by the second assumption that each looks to be on C-*forward*. A "backward-facing" correlational theory of reference delivers the same verdict: standardly, when an object causes an experience of any of the three sorts the three objects are causing in Irreg, the object is on C-*forward*. Finally, a "forward-facing" theory of reference on which visual content goes by the action or belief it is disposed to cause also delivers the same verdict: there is no reason to suppose that Irreg couldn't know to, for instance, point in C-*forward* when having one of these experiences and desiring to point in the direction in which the experienced object is located.

C-distances For Reg, the heart looks to be at a certain C-distance *near*, the moon looks to be at *middle*, and the star looks to be at *far* (once again, Reg is a fully normal observer). Could the objects could also look to have these distances to Irreg?

Say that the point in C-space at which both eyes are aimed (at which L-*forward* and R-*forward* intersect) is the *target*. The distance to the target is a function of the distance $2d$ between the eyes and the angle α at which the eyes are canted with respect to C-*forward*, namely $d \tan \alpha$. If, moreover, the disparity between

the projections of an object on C-*forward* onto the right and the left retina is 2β degrees, then the distance to the object is $d\tan(\alpha+\beta)$ (if the left retinal projection is clockwise of the center and the right retinal projection is counterclockwise, as with an object closer than the target, then the value of the former projection minus the value of the latter will have a positive value; if the other way around, as with an object farther from the target, a negative value). For entities not on C-*forward*, their C-distances are close enough to these values for present purposes.

Suppose that the distance in attentional space between the L- and R-presentations of an object located on the target is γ. For a normal observer, $\gamma = 0$; for Irreg, $\gamma = 20$. For normal observers, the C-distance of an object such that the disparity between its L- and R-presentations is β is reckoned by the visual system to be $d\tan(\alpha+\beta)$, so that this value is the C-distance the object looks to be at. This is a special case of reckoning the distance as $d\tan(\alpha+\beta-\gamma)$, for $\gamma = 0$, so that Irreg could accurately see the C-distances of objects if his visual system reckoned them as the value $d\tan(\alpha+\beta-20)$.

I see no reason that Irreg's visual system couldn't work in this way (nor do I assume that it *must* work this way). Hence, I see no reason that C-distances couldn't look to Irreg the same way they look to Reg. Nor do I assume that the value the visual system assigns to γ has any particular influence on phenomenology, or any influence at all: either way, phenomenology differs between Reg's and Irreg's experience, though the content is the same. Moreover, both backward- and forward-facing causal theories deliver the verdict that it must work this way.

Content is *multiply realizable* by experiences with differing visual fields, contra intentionalism.

5.3 Comparison to Standard Spectral Inversion Arguments

I have argued, granting element intentionalism, that, while Irreg's case of attentional shear results in his visual field being different from Reg's, there is no reason to suppose that the representational content of Irreg's visual experience differs from

that of Reg's experience. This isolates form intentionalism as the target of my argument. Element intentionalism is more widely discussed than form intentionalism; in particular, a widely discussed case against element intentionalism, the *spectral inversion* objection, closely resembles my attentional shear objection. After briefly rehearsing the spectral inversion objection, I will highlight two of its vulnerabilities which my attentional shear objection lacks.

The conclusion of a spectral inversion argument is that presentations are not ostensible external quality-instances. In particular, letting "red*" be that feature, instances of which are standardly presented to normal subjects when seeing red objects, it is concluded that red* is not red, and is moreover rather some internal quality R.[11]

In such an argument, it is supposed that Invert is "spectrally inverted", or abnormal in that seeing red (green) things causes his visual field to contain presentations of green* (red*) and so forth. If Invert sees a cucumber and Nonvert (who is normal) sees a cherry, both Invert's and Nonvert's visual fields will contain a presentation of red*; still, it is argued, Nonvert's experience will thereby represent green while Invert's will represent red, so that the ostensible colors of which Nonvert and Invert are aware differ; hence red* is not an ostensible color. Moreover, if it is assumed that the only external property a visual field presenting red* would thereby represent is a color (for dissent, see Shoemaker 1994; Thau 2002) it follows that red* can't be an ostensible external property at all, and must rather be internal.[12]

[11]This is the conclusion relevant to intentionalism as I have defined it, as a thesis about the visual field. More commonly, intentionalism is understood as the thesis that visual*phenomenal characters*, what visual experiences are like to their subjects, are representational properties. This position—*pc-intentionalism*—is equivalent to intentionalism as I have characterized it on the assumption that the phenomenal character of a visual experience determines and is determined by the nature of the visual field involved in the experience. Upcoming footnotes contain a running discussion of pc-intentionalism.

[12]A converse case involving both Invert and Nonvert seeing a cherry cuts against supervenience of visual field properties on representation.

The conclusions of this argument are readily converted into an objection to pc-intentionalism given that the representational properties of a visual experience determine and are determined by the representational properties of the visual field involved in the experience. By this assumption, Invert's and Nonvert's experiences have the same representational properties; by the as-

The first of two complaints about this argument appeals to transparency (attribution issues are not straightforward, but one could plausibly regard Harman (1990), Tye (1995), Dretske (1995), and, in a sense, Byrne (2001) as pushing this line)). As I argued in section 3, above, transparency is plausible interpreted as the claim that the attentional spotlight can only be turned on external features. On this interpretation, presentations end up as ostensible instances of external features, since surely nothing is a presentation which one can't turn the attentional spotlight on. Hence, supposing that if presentations are external, they are ostensible or represented colors, and granting that the colors Invert's and Nonvert's visual fields represent differ, Invert's and Nonvert's color experiences of the cherry don't involve the same presentations after all.[13]

A parallel complaint isn't available against the attentional shear argument, since, as I also argued, transparency isn't plausible understood as the thesis that we only *experience* external properties in visual experience. We experience attentional space distances between presentations, so Irreg can experience his two cardioid regions of attentional space containing presentations of yellowness as being at a non-zero distance in attentional space from one another.[14]

The second of two complaints is based on the thought that which color properties are represented by a visual field are partly determined by the internal character of the visual field (Segal, 1991; Hardin, 1997; Hilbert and Kalderon, 2000). For instance, the color represented in a region of visual field in which red* (green*)

sumption of the previous footnote, their experiences have different phenomenal characters, contra pc-intentionalism.

[13]Concerning pc-intentionalism, some have argued that "gazing again outward" is an exhaustive infallible guide to which phenomenal character one's experience has, and that any property one's experience is revealed to have by gazing again outward determines and is determined by which external properties one's experience represents (cf. Evans 1982; Tye 2000; Byrne 2001). The connection to transparency should be clear, and the failure of transparency should undermine the claim that if gazing again outward has the second property, it also is an exhaustive guide to phenomenal character.

[14]A related complaint is that the argument assumes, perhaps illicitly, that color phenomenology is internally determined (Lycan, 2001). The assumption occupying the analogous position in my argument is that the selective copying device realizes the attentional spotlight in a natural way. I have trouble imagining how alternatives to this claim might be true.

is presented must be a "hot" ("cool") color; or the colors represented in regions of visual field in which red* and orange* are presented must be more similar to one another than either is to the color represented in a region of visual field in which green* is presented. Holding fixed the assumptions about regular influence on Invert's visual field of various colors, it follows from these conditions on extension that when Invert sees the cucumber, he represents it as red, contra the causal theory of content (assuming, that is, that only red could satisfy the hotness and similarity conditions associated with red*).

Of the five sorts of properties Irreg's experiences represent, one way to press this complaint concerns the representation of C-distance. Perhaps because of the arrangement of color presentations in his visual field alone Irreg's experience upon seeing the scene in the Double Vision Example is nonveridical, since it represents the heart at *middle*, the moon at *far*, and the star at a still farther distance. After all, these are the distances Reg (or any other normal subject) would represent if the arrangement of color presentations in his visual field were the same as Irreg's. So perhaps this arrangement alone suffices for representation of these distances, in the same way that—as on the objection to spectral inversion—containing a presentation of red* suffices for representing red.

A significant disanalogy between color perception and C-distance perception threatens this line. In a color experience, there is perhaps already enough in the nature of the presentation alone to determine which color it represents: red* stands in a certain position in the network of relations of similarity and exclusion the color*s bear to one another; so perhaps this suffices to select red as represented, on the grounds that it is the color which stands in the analogous position in the network of relations of similarity and exclusion the colors bear to one another.

However, there is not enough in an experience involving L- and R-presentations of x a certain distance apart in attentional space to determine the C-distance of x: an additional inferential step must be taken, involving knowledge of the separation of the eyes, the angle at which they are canted in the head, and a trigonometric

calculation. And why can't this calculation be parameterized to account for a subject's normal separation in attentional space (the value γ) for seen objects on the target?

The remaining way in which this approach might be plausibly pressed—which doesn't face the same difficulty—is against the representation of L- and R-direction. When the experience of a normal subject involves a presentation of x in the position α degrees from the center of attentional space, it is represented (at Tye's intermediate level of consciousness) either at L-(α degrees from *forward*) or at R-(α degrees from *forward*). Perhaps this is so for any possible subject of visual experience. If so, Irreg's experience in the Double Vision Example represents the heart at L- and R-*forward*, the moon at L-(10 degrees to the left) and R-(10 degrees to the right), and the star at L-(20 degrees to the left) and R-(20 degrees to the right), and is therefore nonveridical.

I am not sanguine about the prospects of this line, either. In the case of color experience, the higher-order properties allegedly determining the property represented when red* is presented are knowable by reflection on experience to anyone who has had an experience in which red* is presented. But since L- and R-directions in binocular vision are represented merely at Tye's intermediate level of consciousness, nothing about them is knowable by reflection on experience. Monocularity bears only a mere causal connection to spatial representation in binocular vision: it has no immediate impact on rational reaction to experience. If, despite this, there is representation of monocular space in binocular vision, it seems that only a correlational theory of reference would be relevant: it is unclear what would ground the assignment of any thicker condition which would yield the result the opponent of the argument requires.[15]

[15]Two further approaches to spectral inversion involve interpreting presentations as ostensible external secondary qualities (Shoemaker, 1994) or—concerning pc-intentionalism—interpreting phenomenal characters as functions from context-like entities into colors (Chalmers, 2005): characters in the Kaplanian sense. Assessing the prospects of either analogous approach to attentional shear is beyond the scope of this discussion.

6 Conclusion, and a Remaining Issue

I have argued in this paper that spatial aspects of the visual field do not supervene on its representational properties: that two visual fields can share their representational properties but differ in their spatial properties. I have additionally argued that the spatial aspects of the visual field are very closely tied to the possible states of an "attentional spotlight", a device which selects and highlights visual information. If consciousness is reducible at all, so are the spatial properties of this device, so that acceptance of a visual field with spatial properties poses no special problems for a naturalistic theory of mind, contra to what many opponents of sense-datum theory have argued.

A structured system involves both structure and "stuffing". For all I have argued here, the stuffing in the visual field is purely representational: I did not adjudicate between the sense-datum theory and the Kantian theory here. Debate over intentionalism has focused exclusively on stuffing at the expense of structure. But this is a mistake, since intentionalism also entails claims about the status of structure.

The positive theory leaves an open question. Suppose that one sees a tilted penny, so that an elliptical region of attentional space is filled with presentations of a homogeneous color. It seems as if in virtue of this fact alone, one represents an object with an elliptical frontal-plane projection. Why is this? What is the relation between attentional space and representation of frontal-plane projection in virtue of which this intrinsic connection obtains? The answer to this question will have to wait for another occasion.

References

Baldwin, Thomas, 1992. 'The Projective Theory of Sensory Content'. In Tim Crane, editor, *The Contents of Experience*. Cambridge: Cambridge University Press.

Barnes, Winston H. F., 1944/5. 'The Myth of Sense-Data'. *Proceedings of the Aristotelian Society*, 45:***–***. Reprinted in Swartz 1965.

Block, Ned, 1990. 'Inverted Earth'. In James Tomberlin, editor, *Action Theory and the Philosophy of Mind*, volume 4 of *Philosophical Perspectives*, 53–79. Atascadero: Ridgeview.

Boghossian, Paul A. and J. David Velleman, 1989. 'Colour as a Secondary Quality'. *Mind*, 98:81–103. Reprinted in Byrne and Hilbert 1997.

Boghossian, Paul A. and J. David Velleman, 1991. 'Physicalist Theories of Color'. *The Philosophical Review*, 100:67–106. Reprinted in Byrne and Hilbert 1997.

Broad, Charlie Dunbar, 1923. *Scientific Thought*. London: Routledge and Kegan Paul. Selections reprinted in Swartz 1965 as 'The Theory of Sensa'.

Broad, Charlie Dunbar, 1925. *Mind and Its Place in Nature*. London: Kegan Paul, Trench, Trubner & Co.

Byrne, Alex, 2001. 'Intentionalism Defended'. *The Philosophical Review*, 110:199–240.

Byrne, Alex, 2002. 'Something About Mary'. *Grazer Philosophische Studien*, 63:123–40.

Byrne, Alex and David R. Hilbert, editors, 1997. *Readings on Color: The Philosophy of Color*, volume i. Cambridge, MA: The MIT Press.

Chalmers, David J., 2005. 'The Representational Character of Experience'. In Brian Leiter, editor, *The Future for Philosophy*. Oxford: Oxford University Press.

Chisholm, Roderick M., 1950. 'The Theory of Appearing'. ***, ***–***. Reprinted in Swartz 1965.

Clark, Andy, 2000. *A Theory of Sentience*. Oxford: Oxford University Press.

Crane, Tim, 1992. 'Introduction'. In Tim Crane, editor, *The Contents of Experience*. Cambridge: Cambridge University Press.

Dretske, Fred I., 1995. *Naturalizing the Mind*. Cambridge, MA: The MIT Press.

Evans, Gareth, 1982. *The Varieties of Reference*. Oxford: Oxford University Press.

Hardin, Clyde L. 1997. 'Reinverting the Spectrum'. In Alex Byrne and David R. Hilbert, editors, *The Philosophy of Color*, volume i of *Readings on Color*. Cambridge, MA: The MIT Press.

Harman, Gilbert, 1990. 'The Intrinsic Quality of Experience'. In James Tomberlin, editor, *Action Theory and the Philosophy of Mind*, volume 4 of *Philosophical Perspectives*, 31–52. Atascadero: Ridgeview. Reprinted in Harman 1999.

Harman, Gilbert, 1999. *Reasoning, Meaning, and Mind*. Oxford: Oxford University Press.

Hilbert, David R. and Mark K. Kalderon, 2000. 'Color and the Inverted Spectrum'. In Steven Davis, editor, *Color Perception: Philosophical, Psychological, Artistic, and Computational Perspectives*. Oxford: Oxford University Press.

Horgan, Terence and John Tienson, 2002. 'The Intentionality of Phenomenology and the Phenomenology of Intentionality'. In David J. Chalmers, editor, *Philosophy of Mind: Classical and Contemporary Readings*. Oxford: Oxford University Press.

Huemer, Michael, 2001. *Skepticism and the Veil of Perception*. Lanham, MD: Rowman and Littlefield.

Hume, David, 1739/1978. *Treatise on Human Nature*. Oxford: Oxford University Press.

Jackson, Frank, 1977. *Perception: A Representative Theory*. Cambridge: Cambridge University Press.

Kriegel, Uriah, 2002. 'Phenomenal Content'. *Erkenntnis*, 57:175–98.

Levine, Joseph, 2003. 'Experience and Representation'. In Quentin Smith and Alexander Jokic, editors, *Consciousness: New Philosophical Essays*. Oxford: Oxford University Press.

Loar, Brian, 2003. 'Phenomenal Intentionality as the Basis of Mental Content'. In M. Hahn and B. Ramberg, editors, *Reflections and Replies*. Cambridge, MA: The MIT Press.

Loar, Brian, in preparation. 'Transparent Experience and the Availability of Qualia'. MS, Rutgers University.

Lycan, William G., 2001. 'The Case for Phenomenal Externalism'. In James Tomberlin, editor, *Metaphysics*, volume 15 of *Philosophical Perspectives*, 17–36. Oxford: Blackwell.

Marr, David, 1982. *Vison*. New York: W. H. Freeman.

Martin, Michael G. F., 2002. 'The Transparency of Experience'. *Mind and Language*, 17:376–425.

Peacocke, Christopher, 1983. *Sense and Content: Experience, Thought, and Their Relations*. Oxford: Clarendon Press.

Peacocke, Christopher, MS. 'Awareness of Action???' Talk at Santa Barbara, Winter 2004.

Perkins, Moreland, 1983. *Sensing the World*. Philadelphia: Hackett.

Price, Henry Habberly, 1950. *Perception*. London: Methuen, second edition.

Rey, Georges, 1998. 'A Narrow Representationalist Account of Qualitative Experience'. In James Tomberlin, editor, *Language, Mind, and Ontology*, volume 12 of *Philosophical Perspectives*, 435–58. Oxford: Blackwell.

Robinson, Howard, 1994. *Perception*. London: Routledge.

Searle, John, 1990. 'Consciousness, Explanatory Inversion, and Cognitive Science'. *Behavioral and Brain Sciences*, 13:585–642.

Segal, Gabriel, 1991. 'In Defense of a Reasonable Individualism'. *Mind*, 100:485–494.

Shoemaker, Sydney, 1991. 'Qualia and Consciousness'. *Mind*, 100:507–24. Reprinted in Shoemaker 1996.

Shoemaker, Sydney, 1994. 'Phenomenal Character'. *Noûς*, 28:21–38. Reprinted in Byrne and Hilbert 1997.

Shoemaker, Sydney, 1996. *The First-Person Perspective and Other Essays*. Cambridge: Cambridge University Press.

Shoemaker, Sydney, 2001. 'Introspection and Phenomenal States'. *Philosophical Topics*, 29:***–***.

Siewert, Charles, 1998. *The Significance of Consciousness*. Princeton, NJ: Princeton University Press.

Stoljar, Daniel, in preparation. 'The Argument From Diaphanousness'. MS, Australian National University.

Strawson, Galen, 1994. *Mental Reality*. Cambridge, MA: The MIT Press.

Sturgeon, Scott, 2000. *Matters of Mind*. London: Routledge.

Swartz, Robert J., editor, 1965. *Perceiving, Sensing, and Knowing*. Berkeley: University of California Press.

Thau, Michael, 2002. *Consciousness and Cognition*. Oxford: Oxford University Press.

Thompson, Brad, 2003. *The Nature of Phenomenal Content*. Ph.D. thesis, University of Arizona, Tucson, AZ.

Tye, Michael, 1992. 'Visual Qualia and Visual Content'. In Tim Crane, editor, *The Contents of Experience*. Cambridge: Cambridge University Press.

Tye, Michael, 1995. *Ten Problems of Consciousness*. Cambridge, MA: The MIT Press.

Tye, Michael, 2000. *Consciousness, Color, and Content*. Cambridge, MA: The MIT Press.

Wittgenstein, Ludwig, 1921/1974. *Tractatus Logico-Philosophicus*. Atlantic Highlands, NJ: Humanities Press. Translated by David F. Pears and Brian F. McGuinness.