

Luke Roelofs

What are the Dimensions of the Conscious Field?

Abstract: *I analyse the meaning of a popular idiom among consciousness researchers, in which an individual's consciousness is described as a 'field'. I consider some of the contexts where this idea appears, in particular discussions of attention and the unity of consciousness. In neither case, I argue, do authors provide the resources to cash out all the implications of field-talk: in particular, they do not give sense to the idea of conscious elements being arrayed along multiple dimensions. I suggest ways to extend and generalize the attentional construal of 'field-talk' to provide a genuine multiplicity of dimensions, through the notions of attentional proximity and causal proximity: the degree to which two experiential elements are disposed to bring one another into attention when attended, or to interact in other distinctively mental ways. I conclude that if consciousness is a field, it is one organized by attentional and/or causal proximity.*

Consciousness displays phenomenological structure. But what kind of structure? In this paper I want to explore an idea which has often been suggested but rarely developed in detail, namely that consciousness is structured as a *field*. A number of prominent authors writing about consciousness have used the language of fields, often in prominent or emphatic positions. Moreover, there is a certain intuitive attractiveness to the idea: consciousness does seem to comprise elements somehow arrayed, some closer to each other, some further away, some at the centre, others further out. But, I will argue, the analyses usually presented with such talk do not actually yield the kind of structure which the idea of a conscious field suggests.

Correspondence:
Email: luke.roelofs@utoronto.ca

I take it that, to be a field, consciousness must have its diverse elements¹ laid out along two or more ‘dimensions’, i.e. standing in relations of degree, in at least two more-or-less orthogonal ways. By this definition, clear cases of fields include the visual field, a physical field of force, any graph, and space itself. Time is not a field, since it has only one dimension, though relativistic space-time is. Moreover, I take it we can appropriately call consciousness a field only if it ‘presents itself’ to us as structured by these relations, both in that we have ready introspective access to the relations themselves, and in that they exercise a decisive influence on the phenomenological judgments we are inclined to make about consciousness. Obviously the elements of consciousness bear relations of degree to each other, but it is a further claim that they are experienced as organized by these relations.

An analysis of the structure of consciousness may, moreover, open up ways to analyse the subject of experience, since the two are clearly connected. In particular, an analogy of the form subject:consciousness::eye:visual field is discussed by Wittgenstein (1921/1961, 5.633), and a similar analogy replacing ‘eye’ with ‘vanishing point’ is discussed by Sorensen (2007). But in both cases the exact meaning of the analogy remains obscure until we have some idea of how the structure of consciousness in general compares to that of the visual field.

It may be that consciousness is not a field, or can be called a field only in a misleading or confused way (I consider a proposal of this sort in Section 3). But if talk of a field does capture something about the structure of consciousness, what? What are the dimensions of this field? Where are its centre and/or boundaries? These are the questions I will pursue.

In Section 1 I review some uses of ‘field-talk’ to convey the unity of consciousness, noting that this generally provides no dimensions for elements to be arrayed in. In Section 2 I consider and reject the deflationary proposal that our tendency to use ‘field’ talk about consciousness arises simply from the fact that consciousness *represents* a spatial field. In Section 3 I review discussions of the attentional structure of the field, noting that they provide only one dimension. In Sections 4 and 5 I propose more multi-dimensional analyses, and finally in

[1] One of the points of contention in this area concerns whether these elements should be called ‘experiences’, ‘parts of experiences’, or something else (see Tye, 2003, pp. 21–41; Raymont and Brook, 2009; Dainton, 2010; Bayne, 2010, pp. 20–46; and Chudnoff, 2013, for discussion of this question). I am neutral here, but will sometimes refer to them as ‘experiences’ for convenience — if this is considered misleading, the term ‘experience’ can be replaced with the more colourless term ‘element’.

Section 6 I argue that consciousness is a *deforming* field, unlike the visual field.

1. The Conscious Field as Unified Consciousness

Many authors say that consciousness is a field, meaning that it is unified in some way. In Bayne and Chalmers (2003) the term ‘field’ is introduced as a way of referring to a total conscious state, i.e. a conscious state subsuming all of a subject’s conscious states: ‘a subject’s total conscious state might be thought of as the subject’s conscious field’ (*ibid.*, p. 5). This is followed by the use of the terms ‘subsumptive unity’ and ‘field unity’ as interchangeable (*ibid.*, e.g. pp. 8, 10, 12). Talking of a phenomenal field or conscious field is taken as equivalent to talking of unified consciousness.

This is reinforced in Bayne’s (2010) book, *The Unity of Consciousness*, the first chapter of which is entitled ‘The Phenomenal Field’. Here he repeats that ‘what it is for a pair of experiences to occur within a single phenomenal field just is for them to enjoy a conjoint phenomenality’ (noting that ‘this claim is stipulative... not... a substantive thesis’) (*ibid.*, p. 11). Later he says that ‘the field metaphor... accurately captures the structure of consciousness at a time’ (*ibid.*).

Searle (2000) also sometimes uses the term ‘field’ to mean simply ‘unified consciousness’, but at other times seems to associate the term ‘field’ with the stronger claim that consciousness is *holistic*. For instance, to convey the idea that consciousness is distinctively unified, he says ‘I do not just [have a list of particular experiences]. Rather I experience all of these as part of a single unified conscious field’ (*ibid.*, p. 561–2). But he also distinguishes what he sees as two opposed models of consciousness:

[T]he building block model, according to which any conscious *field* is made of its various parts, and the unified *field* model, according to which we should try to explain the unified character of subjective states of consciousness. (*Ibid.*, p. 557, emphasis added)²

In a similar vein, Uriah Kriegel unpacks the claim that ‘my overall phenomenology is unified at the time’ as meaning that ‘it does not feel

[2] This distinction is confusing, for two reasons. Firstly, if ‘conscious field’ simply means ‘unified consciousness’, then ‘unified field’ is redundant, meaning ‘unified unified consciousness’. Secondly, it seems that adherents of the ‘building block model’ can perfectly well agree that consciousness is unified, and hence that it is a field (indeed, a unified field). This is reinforced by Searle’s (2000) own use of the term ‘field’ in characterizing this position, e.g. p. 572: ‘The building block theory: The conscious field is made up of small components that combine to form the field.’

like just so many unconnected items, but rather like a single cohesive “field of experience” (Kriegel, 2009, p. 172).

So there is a pattern of people using this language to convey something about the unified character of consciousness. But in so far as unified consciousness is analysed in terms of every bit of consciousness being related in some intimate way to every other bit, or to the whole, it does not provide any dimensions. If two experiences are unified, that in itself gives no meaning to the idea that they are close to or far from each other, in any particular direction, within any particular region of the field. To this extent it is more like a *set* of elements defined by bearing a certain relation to one another, rather than a *field*.

I think that we should retain the idea of a conscious field being equivalent to unified consciousness; it is intuitively attractive to think of consciousness being phenomenologically structured as an ‘array’ or ‘space’, and this helps to capture the idea that all elements of consciousness are connected. But the ‘field’ idea inherently implies a further structure, which talk of unity does not capture.

2. The Conscious Field as External Space: A Deflationary Proposal

There is a way to deflate the idea of a conscious field, namely by interpreting it simply as the claim that consciousness *represents* the three-dimensional spatial field in which we live. There would then be no distinctive question about the structure of the conscious field, beyond the different ways that consciousness represents space.

This idea appears in the literature on the unity of consciousness as the claim that ‘consciousness is unified’ means simply that consciousness *represents* a single unified space. Dainton calls this deflationary analysis the S-thesis: ‘that simultaneous experiences are co-conscious by virtue of occurring within a single unified phenomenal space’ (Dainton, 2004, p. 9), where a ‘phenomenal space’ is a set of experiences isomorphically representing objects in objective space. The S-thesis, as I will follow Dainton in calling it, has been considered by a number of philosophers, and consistently rejected, for broadly similar reasons. In particular, I believe we should reject the S-thesis about phenomenal unity because not all experiences have spatial content, and because spatial disunity is compatible with phenomenal unity. The failure of the S-thesis, along with the fact that attentional relations can cross-cut represented spatial ones, should lead us to reject the deflationary proposal about the field-structure of consciousness.

The first objection is that not all of our experiences have spatial content, and yet all of our experiences seem to be unified. This objection is given in very similar terms by Dainton (*ibid.*, p. 9), and by Bayne and Chalmers (2003, p. 4). Various examples are given of experiences without spatial content (the first two from Dainton, the last two from Bayne and Chalmers):

- conscious thoughts which do not involve spatial imagery of any kind;
- certain types of sounds;
- an emotional experience such as melancholy;
- a conscious thought about philosophy.

One response to this objection would be to say that these other, non-spatial, experiences are not strictly phenomenally conscious, or that in so far as they are conscious it is by having associated sensory components. Tye takes something like this line, in that he regards ‘the phenomenology of conscious thoughts [as] derive[d] fundamentally... from the phenomenology of their associated linguistic, auditory images’ (Tye, 2003, p. 79), but I follow Bayne and Chalmers in finding this implausible (Bayne and Chalmers, 2003, p. 7). There is something it is like to think, just as there is ‘something that it is like to get a joke, to be puzzled about a problem, and to see that an argument is fallacious’ (*ibid.*, p. 6).³

A second objection to the S-thesis is that we can imagine hypothetical cases in which a subject has spatially disunified, but phenomenally unified, experiences (Bayne argues that there may also be actual examples of this in ‘heautosopic hallucinations’, 2010, pp. 261–2). Dainton, Bayne, and Tye all establish this through thought experiments involving a single brain receiving two disconnected streams of perceptual and/or bodily sensations.

Bayne’s example is ‘Borgy’, three bodies whose brains are connected by radio transmitters so as to function as smoothly together as the hemispheres of a normal brain (*ibid.*, pp. 262–6). Dainton’s example involves two bodies, each without a brain but with radio receivers that pick up signals sent by a single brain, held in a secure location elsewhere (2004, pp. 9–10). Tye’s example, designed to show the separability of bodily and perceptual spaces, involves a person buried in

[3] An alternative response might be that thoughts, etc. in fact have spatial content of a degenerate sort. One might say, for instance, that thoughts are experienced as occurring ‘in the head’, behind the eyes and between the ears, and are thus spatially unified with other experiences. I have not seen this approach developed in sufficient detail, and am unsure of how compelling it could be.

sand but with their eyes and ears attached to long nervous stalks projecting above the sand and exposed to a 3D film of a rollercoaster ride (2003, pp.76–8).

The S-thesis implies that in these cases, because the subject lacks spatial unity, they will also lack phenomenal unity — under the pressure of these disorganized inputs, their consciousness will ‘split’ to cope with the different streams. All three authors consider and ultimately reject this idea; there is little reason to think this would happen. We often manage to process apparently conflicting sensory inputs, such as when subject to illusions, or wearing vision-inverting goggles (this point is made in Tye, 2003, p. 78, and Bayne, 2010, p. 264).

If the S-thesis is false, then phenomenal unity is both irreducible to spatial unity and central to the idea of a conscious field: this suggests that the idea of a conscious field goes beyond the representation of spatiality. Here is further support for that conclusion: attentional relations are important to conscious structure, but need not coincide with represented spatial relations. For instance, two objects might be right beside each other in the visual field, but one be the focus of attention and the other completely peripheral and unattended. Conversely, two objects perceived as being far apart might be ‘close together’ in attentional terms, if we were consciously focused on comparing them, or discerning differences between them.

I conclude that we should not accept the deflationary proposal. If consciousness is a field, we will not understand the structure of this field by investigating the *content* of conscious representations, but by investigating the structure in which those representations *themselves* are arrayed.

3. The Conscious Field as Attentional-Focus-and-Periphery

We have looked at how the idea of a conscious field relates to the unity of consciousness, and to the representational content of consciousness. A third approach to the conscious field aims to characterize it in relation to attention. Here consciousness is described as a field with a centre and a margin: what is attended is at the centre, while what is unattended is at the margin. Kriegel, Searle, and Bayne all sometimes speak in this way,⁴ and William James provides an earlier example:

In most of our fields of consciousness there is a core of sensation that is very pronounced... [While] thinking and feeling, [you] are getting

[4] For instance, Bayne describes the phenomenal field as ‘a small band of focal experience surrounded by an experiential penumbra’ (2010, p. 79), and Kriegel observes that ‘Usually, the phenomenological field has a center/periphery structure’ (2009, p. 173).

through your eyes sensations of my face and figure... The sensations are the centre or focus, the thought and feelings the margin, of your actually present conscious field. (James, 1899/1983, p. 18)

It is unclear, however, whether this language of centres and peripheries is more than a convenient image for expressing the observation that we attend to some things and not to others. Sebastian Watzl claims that it is more than an image, and that ‘The structure of attentional space can be treated with some of the same formal precision, with which we can, say, treat the structure of space-time’⁵ (Watzl, 2011, pp. 159–60). His ‘structuralist’ account of attention holds that:

Consciously attending to something consists in the conscious mental process of structuring one’s stream of consciousness so that some parts of it are more central than others. (*Ibid.*, p. 145)

Watzl describes this ‘structuring’ as the production of an ‘attentional space’, constituted by ‘attentional relations’ between experiences (*ibid.*, pp. 158–60). The simplest of these relations, and the only one Watzl discusses in any detail, is ‘peripherality’, where something is at the centre of consciousness (i.e. attended to) iff all other experiences are ‘peripheral to’ it. Experiences outside the centre can also be peripheral to each other, when one is ‘more peripheral than’ the other, i.e. further from the centre of attention (*ibid.*, p. 160). The peripherality relation is taken as primitive, irreflexive, antisymmetric, and transitive (*ibid.*, p. 160).

I think Watzl’s account is a step in the right direction, but it stops short of giving us an adequate account of how consciousness might be a field. This is because it is ‘centrocentric’, overly focused on the focus; the only component of attentional structure that is explained is the peripherality relation, and this only provides a single dimension: closeness to focal attention. If this were the only dimension to speak of, there would be no reason to plot attention as a *centre*, since it could just be taken as one extremity of a line, with ‘fringe’ contents at the other.

Intuitively, though, we can make sense of two equally peripheral experiences being ‘close to’ each other, and another two being ‘far apart’. Consciousness is experienced as an organized whole, in which some elements are more closely connected than others. For instance, I might be focusing my attention on my computer screen, with a number of contents somewhat peripheral to this: plans for what to write in the

[5] Watzl does not speak of consciousness as a ‘field’, but does speak of it as involving a sort of ‘space’. I take it that anything which could be called a ‘space’ could just as well be called a ‘field’.

next section, memories of the previous page, awareness of my legs, awareness of my feet, etc. Some of these might be peripheral to others, but they need not be — they might all be perfectly equally-situated in respect of peripherality. But my awareness of my legs and of my feet are clearly closely connected, as are my intentions and memories related to this paper. Both form ‘clusters’ whose members are more related to each other than to members of other clusters.

Attentional structure does seem to be the right thing to investigate, but we should seek an account that can capture ‘peripheral connections’ between unattended experiences, and thereby capture the idea that consciousness is a field, i.e. something with more than one dimension.

4. The Conscious Field as a Terrain of Attentional Proximity

In light of Section 3’s discussion we should look for an attentional relation which can hold between two experiences independently of their relation to the present focus of attention, and which is a matter of degree. My proposed candidate is ‘attentional proximity’, the propensity of experiences to transfer attention to each other. For instance, recall the example in which I have four peripherally conscious experiences:

- (a) bodily awareness of my legs;
- (b) bodily awareness of my feet;
- (c) intentions about the next section;
- (d) memories of the last page.

We wanted a way to capture the idea that (a) and (b) are ‘close together’ in the conscious field, as are (c) and (d), though the two pairs may be very ‘far apart’. I think we can capture this by observing that if I attend to (a), (b) will be more likely to enter attention, and vice versa. If I have been ignoring a mild pain in my foot, I am more likely to notice it if my attention is drawn to my leg. The same can be said of the different experiences related to this paper.⁶

On this analysis, the conscious field is a sort of attentional terrain, through which attention moves. It moves quickly and easily over short ‘distances’, and with more difficulty over long ones. We are intuitively inclined to speak of consciousness as a ‘field’ because we

[6] We might talk either of ‘me easily moving my attention’ or of ‘attention easily/frequently moving’, depending on whether we wished to think of attention as something the subject actively does, or as an event in a deterministic causal chain.

recognize it as something that ‘we’ can move (our attention) around in.

It is likely that attentional proximity strongly correlates with many representational relations. For instance, things represented as spatially close to each other tend, other things being equal, to be attentionally close to each other; the same goes for things which have a common ‘theme’ (as the above example illustrates). But neither sort of representational unity can completely explain attentional proximity: we may for instance have a strong association between two objectively unrelated ideas, while certain objective connections never occur to us. This is one reason why we should reject attempts to reduce attentional structure to the structure of a represented field.

This proposal lets us make sense of consciousness being a field with distances between its elements. But what about the directions in which there are these distances? These can be constructed from the ratios between distances. For instance, given that distance $A - B = X$, and distance $B - C = Y$, we can determine the dimensions of the space containing A, B, and C by looking at the distance $A - C$. If it is $X + Y$, we can arrange them on a single line, in the order $A - B - C$; if it is $X - Y$, or $Y - X$, we can again arrange them on a single line, with A and C on the same side of B. If it is some other value, we can arrange them in a triangle, thereby establishing two dimensions. Further points D, E, F... will either be fitted into this plane based on their relative distances from A, B, and C, or will have a set of distances requiring further dimensions. For instance, four equidistant points would need to be represented as the corners of a tetrahedron, while five equidistant points would imply more than three dimensions. Dimensions, that is, can be defined implicitly by the structure of distances.

If we analysed the structure of consciousness in terms of a field of attentional proximity, we would obviously need an independent account of attention. Somewhat awkwardly, the account of attention most congenial to this analysis — Watzl’s structuralist account — defines it in terms of the structure of consciousness. But explaining structure by attention, and attention by structure, is circular. This is not a huge problem for this analysis of the conscious field, since there are many other accounts of attention available, but if we can generalize this analysis to remove its dependence specifically on the notion of attention, that would be worth exploring; the next section suggests one approach to doing so.

5. The Conscious Field as a Terrain of Causal Proximity

We can broaden the notion of ‘attentional proximity’ to include other forms of interaction as well as transfers of attention. More precisely, we might speak of ‘direct’ causal relations among conscious experiences, as opposed to relations with conscious experiences which depend on the holding of other causal relations.⁷ We could also speak of ‘sensitive’ causal relations, those which occur in virtue of the fine-grained experiential and representational properties of their relata. We could then say that consciousness is a field in which different experiences are laid out according to their ‘causal proximity’, their tendency to affect each other in direct, sensitive ways, which includes but is not limited to transferring attention.

The form that these causal relations take will somewhat depend on the sorts of experiences involved. Two conscious thoughts might interact by informing and revising each other’s content. Two conscious intentions might interact by mutual adjustment. In general, when two conscious experiences interact, it will commonly involve the production of a new experience whose content and properties reflect both. A perception of red and a perception of blue might interact by producing a perception of contrast, for example. This third experience might then have distinctive properties — e.g. two flavours might blend together into one more pleasant than either.⁸ A different sort of interaction, which might be seen as the failed version of the above sort, is the generation of dissonance or tension. Two firmly-held but contradictory beliefs, for instance, will often interact not by producing a consistent joint belief but by producing a (vague or precise) awareness that there is a problem.

However, it is not necessary to specify exactly what sorts of interaction are in question, since all the above forms of interaction tend to go together — the conditions under which two experiences will produce a sense of dissonance are roughly the same as those under which they would produce a conjunctive belief, or produce a third belief which is

-
- [7] For instance, if I hear someone’s voice, which leads me to reflect on my life so far, which leads me to become sad, we would not count that auditory perception as direct cause of the sadness. This distinction is meant to be rough and intuitive.
- [8] ‘Gestalt unity’, as briefly discussed by Tye (2003, p. 13) and Bayne and Chalmers (2003, p. 6), is a special case of this, distinguished by the third experience having very prominent novel features due to very specific, but less prominent, features of the first two. While the mentioned authors all say that Gestalt unity is quite rare, I agree, but also think that the production of simpler sorts of novelty (e.g. two colours producing an experience of contrast) is ubiquitous.

a consequence of their conjunction, or transfer attention from one to the other. We can fudge this question by speaking of causal proximity in general.⁹

This analysis of the dimensions of consciousness connects neatly with at least one account of the *unity* of consciousness, which we observed ‘field-talk’ was often employed to capture. Shoemaker argues that the mechanisms which make a state conscious are also the mechanisms which make two states co-conscious, namely ‘be[ing], at least to some degree, integrated into a larger set of mental states... [with] aptness to generate perceptual and introspective beliefs... availability for use as premises in reasoning... [and] availability for the rational control of behaviour’ (Shoemaker, 2003, p. 63). I take it that the sort of interactions indicated here are the same sort of interactions I appealed to in defining causal proximity, and so the causal proximity account suggests that to be part of a unified consciousness (to be in the field at all) an experience must be at some distance from other items in the field, meaning that it must be ‘at least to some degree’ disposed to interact in the relevant ways. The causal proximity account does not entail a functionalistic account of unity, but it dovetails nicely with such an account if the latter is found independently attractive.

It is worth noting that some of the forms of interaction that constitute causal proximity might obtain between conscious experiences and unconscious states.¹⁰ For example, the standing unconscious belief that tigers are dangerous might interact directly and sensitively with the conscious perception of a tiger, making it salient and producing a sense of fear (at least, that is one way to describe the complex way that knowledge and perception interact).

One might take from this the idea that unconscious beliefs and other states can somehow occupy positions in the conscious field, deter-

[9] In a fully worked out version of this account, we might commit to a definite list of relevant forms of interaction, as well as a weighting scheme for how they relate to distances in the conscious field. Alternatively, we might refrain from picking any one such scheme, as long as we were willing to accept that the structure of the resultant field is consequently *vague*. Similar remarks apply to different ways of characterizing a ‘propensity’ to interact — under different conditions, different causal interactions become more or less likely. Presumably we wish to say ‘under normal conditions’, and it is clearly a vague matter which conditions are normal.

This might be analysed as a form of ontic vagueness, in the structure of consciousness itself, or as a form of representational vagueness in our introspective impression of consciousness (leaning heavily on the idea that the field is the structure in which consciousness ‘presents itself’ to the subject, becoming thereby an object of awareness). Both of these options will be controversial, especially the postulate of ontic vagueness: if we wished to avoid both that might be a reason to prefer the attentional proximity account.

[10] I thank Ole Koksvik for pressing me on this point.

mined by their causal proximity to other states. That is, one might think that our impression of consciousness as a field is more like a conscious impression of the mind as a whole, with both conscious experiences and unconscious ‘background states’ together forming a single interconnected whole.¹¹

Alternatively, one might think that the conscious field necessarily includes nothing but conscious experiences. In particular, one might feel that while conscious unity does not give *structure* to the field, it is still a necessary condition for being in the field at all, and unconscious items plausibly cannot be phenomenally unified with anything. Similarly, one might think that while attentional proximity is only one element of causal proximity, it is a crucial element in that being at some attentional distance from something is a necessary condition for being at any causal proximity from something. Since it is plausible that an unconscious item cannot be at any attentional distance from anything because it is not a potential target of attention, this would entail that unconscious items cannot be in the conscious field.

I wish to remain neutral on whether to make attentional proximity or phenomenal consciousness necessary ingredients in causal proximity, and thereby also remain neutral on whether unconscious items might occupy positions in the conscious field. Depending on how we want to think about the relation of consciousness to the unconscious, we might be intrigued by the idea that items within the conscious field might not themselves be conscious. But we might also rule this out on either of two principled grounds: that only conscious experiences can be phenomenally unified, or that only conscious experiences are potential targets of attention.

Causal proximity not only lets us analyse consciousness as a field without presupposing any view of attention, it also suggests (without demanding) a (partial) analysis of attention in terms of this structure. What is attended is at the centre of the conscious field, and the centre of a field is, overall, the point closest to many other points at once, whereas the periphery is on average more distant from other points. That is, the causal proximity analysis suggests that an experience is the focus of attention when it is made, on average, more causally

[11] An even more radical move would be to construe the unconscious/conscious distinction as itself a matter of degree, possibly correlating with ‘degree of attention’ as a development of the idea that attention is necessary and sufficient for consciousness (cf. Prinz, 2011), so that the unconscious is simply that which is furthest from the attentional focus, out in the far periphery of the conscious field. James seems to suggest something of this sort when he writes that ‘my present field of consciousness is a centre surrounded by a fringe that shades insensibly into a *subconscious* more’ (1909/1977, p. 103, emphasis added).

proximate to the rest of the field's contents — more able to influence them and be influenced by them.

Attending to, for instance, a belief will greatly increase the likelihood that any conflicting beliefs or perceptions we have will come to mind and be used to revise it; attending to a percept brings to mind all of the things we could compare with it, conclude from it, or do with it. As Nagel says of the unity of consciousness, 'for elements of experience or other mental events occurring simultaneously or in close temporal proximity, the mind which is their subject can also experience the simpler relations between them *if it attends to the matter*' (Nagel, 1971, p. 407, emphasis added).

This analysis of attention is only partial; for one thing, it leaves out all reference to the inputs and outputs to the mind (this is how Watzl, 2011, pp. 18–21, differentiates a 'structuralist' analysis of attention from a 'selection' analysis). Those might be essential to attention, or they might not. In practice, it seems likely that the 'consuming systems' of the mind, such as those involved in memory-formation, verbal report, and motor control, will draw primarily from what is at the centre of attention: what is attended is privileged both by the internal structure of consciousness and by the channels leading out from consciousness.

By contrast to the attentional centre, experiences which are unattended, way off in the 'fringe' or 'periphery', will on average be further apart from each other and from the rest of the field. In terms of causal proximity, this means they are relatively isolated and inert — for instance, two contradictory beliefs might sit unattended on the fringe for a long time without generating any tension. Peripheral experiences might form relatively tight clusters with each other, but these clusters will themselves be remote from the rest of the mind.¹²

6. The Conscious Field as Deforming

We can get a better sense of what an account of the conscious field must account for by considering the distinction between *deforming* and *non-deforming* fields. In a deforming field, the relations among elements change over time; constituents of the field can become closer together or further apart. In a non-deforming field, the constituent points and regions always retain their relations to each other, though

[12] If a particular form of interaction was especially 'easy' between two experiences, perhaps because it had occurred so many times before, or because of some innate feature of the mind's architecture, it might occur without needing any attention, and even without entering consciousness; this would mean two elements being extremely close to each other but extremely peripheral to the attentional focus.

things may move in the field from one point to another. Two very prominent examples of fields, the visual field and the Euclidean spatial field, are non-deforming. But consider the field of bodily sensations (proprioception, pains, itches, etc.); plausibly, these are arrayed spatially, but as the body changes shape, over seconds or years, points in this field change their relations to each other. And relativistic physics tells us that objective space(-time) can also deform, though this remains challenging to imagine.

If we think of consciousness as a field of attentional or causal proximity, should we think of it as deforming or non-deforming? One reason to think of it as deforming is that it seems that habituation, learning, etc. can change the relevant relations between elements. But a more subtle reason comes from the need to make sense of the shifting centre of attention.

It is an obvious fact about attention that it can ‘move’: we can shift our attention from one experience to another, just as we can shift our visual focus from one image to another. The attentional focus is not at all like the geometrical centre of a circle, which is fixed in place relative to the rest of the circle.

In fact there are two ways to express this point: that attention can move within consciousness, and that experiences can move in relation to attention. An example of the first phrasing comes from Searle: ‘within my conscious field at any given time I can shift my attention at will from one aspect to another’ (Searle, 2000, p. 564); an example of the second from James: ‘some feeling connected with your own body may have passed from a marginal to a focal place’ (James, 1899/1983, p. 18). I will assume in what follows that this difference is merely cosmetic: the point is simply that there is relative motion.

But while the analogy of attentional focus to visual focus is attractive, it is hard to see them as shifting in the very same way, for the visual focus shifts *only* relative to external space, not relative to the non-deforming visual field itself. Yet there is no analogue to external space for the attentional focus to shift relative to.

To see this more clearly, let us distinguish objective fields, represented fields, and fields of representations. Consider first a paragraph describing the structure of a gravitational field (e.g. saying where it is strongest): this is a set of representations (sentences, words) which represent a field, but are not themselves a field (they are relevantly arrayed in only one dimension). If the paragraph is accurate, its represented field will correspond to an objective (gravitational) field. Compare this with a wordsearch, crossword, or Scrabble game: here a set of representations (words) are arrayed in a two-dimensional field, but

do not represent any sort of field. There is a field of representations, with no represented field.

Finally, consider a map. Different symbols on the map represent different geographic features, and the spatial relations among the symbols represent the spatial relations among the features. Here we have both a field of representations and a represented field, and the two are isomorphic — we might say that the field of representations isomorphically represents a field.¹³

The visual field (or any other sensory field) is a field of representations isomorphically representing an objective field: different images in the visual field stand in certain phenomenal relations (‘visually-above’, ‘visually-to-the-left-of’), which represent their objects as standing in certain objective spatial relations.¹⁴ But now we can see that the visual focus, as a point in the field of representations, does not in fact shift at all. If a point in the field is, say, ten degrees above and to the left of the focus, it will always be ten degrees above and to the left of the focus. This reflects the fact that the fovea cannot move relative to the rest of the retina.

Rather, what shifts is the entire field of representations, relative to the represented field (or vice versa). When the eyes move relative to the head, or the head moves, a given ‘point in the visual field’ (such as the focus) represents a succession of different points in perceived space. When we shift our visual focus, what we are doing is shifting one field relative to another field.

If the attentional focus were like the visual focus in this way, then we should explain attentional shifts by positing *two* distinct fields, shifting relative to each other. But it seems a tall order, and phenomenologically unmotivated, to identify two distinct fields covering everything we are conscious of.

However, if we understand the conscious field as deforming, we have an alternative way to make sense of a field having a mobile centre. If the relations between elements can ‘stretch’ and ‘contract’ (changing ‘distance’), or twist and bend (changing ‘direction’), then we can suppose that at any given time, one element is such that all

[13] There are cases with both sorts of field, but no isomorphism: consider a graph of the average strength of an electrical field over time. Points in the graph are representations arrayed in a (two-dimensional) field, and what they represent is a (three-dimensional) field, though the two fields do not show any point-to-point correspondence.

[14] This isomorphism is only partial, in so far as the visual field is two-dimensional, or two-and-a-half-dimensional, yet represents a three-dimensional space. Thus for instance the relation ‘visually-smaller-than’ can represent both the ‘objectively-smaller-than’ relation and also the ‘objectively-further-away-than’. I thank Ole Koksvik for pressing me on this point.

relations of a certain sort to that element are ‘contracted’, to be ‘shorter’ than they otherwise would be, and are also ‘twisted’, to be more uniform in length. That element would then come to be the centre, but this centre could shift, so that another element came to ‘pull’ all other elements closer to it, while relations to the previous centre ‘relaxed’ and became ‘longer’.

So I conclude that if consciousness is a field, it is one which changes its own shape, ‘scrunching itself up’ first around one point, and then around another. To use familiar metaphors, when we shine a ‘spotlight’ on that experience, it is displayed on the stage of the ‘Cartesian theatre’ — but the actors and audience are themselves just parts of the theatre, and the performance is the theatre contorting so as to touch some parts of itself with others.

Conclusions

In this paper I have tried to develop an observation: for consciousness to be a field, it should allow for elements to be at variable distances along multiple dimensions, and existing analyses do not provide for this. I then proposed a way to allow for this, by thinking of elements of consciousness as arrayed according to their causal, or specifically attentional, proximity. Neither unity, nor the attentional focus, provides the necessary structure for consciousness to be a field; this may be unimportant, for consciousness may not be a field. Yet talking about it as a field has been attractive to many authors, and in so far as they wish to regard consciousness as phenomenologically structured in a field-like way, I hope that the suggested analyses, in terms of either attentional or causal proximity, are of interest.

References

- Bayne, T. (2010) *The Unity of Consciousness*, Oxford: Oxford University Press.
- Bayne, T. & Chalmers, D. (2003) What is the unity of consciousness?, in Cleeremans, A. (ed.) *The Unity of Consciousness: Binding, Integration, Dissociation*, New York: Oxford University Press.
- Chudnoff, E. (2013) Gurwitsch’s phenomenal holism, *Phenomenology and the Cognitive Sciences*, **12** (3), pp. 559–578.
- Dainton, B. (2004) Précis of ‘Stream of Consciousness’, *Psyche*, **10** (1).
- Dainton, B. (2010) Phenomenal holism, *Royal Institute of Philosophy Supplement*, **67**, pp. 113–139.
- James, W. (1899/1983) *Talks to Teachers on Psychology and to Students on Some of Life’s Ideals*, Cambridge, MA: Harvard University Press.
- James, W. (1909/1977) *A Pluralistic Universe*, Cambridge, MA: Harvard University Press.

- Kriegel, U. (2009) *Subjective Consciousness: A Self-Representational Theory*, New York: Oxford University Press.
- Nagel, T. (1971) Brain bisection and the unity of consciousness, *Synthese*, **22** (3/4), pp. 396–413.
- Prinz, J. (2011) Is attention necessary and sufficient for consciousness?, in Mole, C., Smithies, D. and Wu, W. (eds.) *Attention: Philosophical and Psychological Essays*, Oxford: Oxford University Press.
- Raymont, P. & Brook, A. (2009) Unity of consciousness, in Beckermann, A., McLaughlin, B. & Walter, S. (eds.) *Oxford Handbook of Philosophy of Mind*, Oxford: Oxford University Press.
- Searle, J. (2000) Consciousness, *Annual Review of Neuroscience*, **23**, pp. 557–578.
- Sorensen, R. (2007) The vanishing point: A model of the self as an absence, *The Monist*, **90** (3), pp. 432–456.
- Shoemaker, S. (2003) Consciousness and co-consciousness, in Cleeremans, A. (ed.) *The Unity of Consciousness*, New York: Oxford University Press.
- Tye, M. (2003) *Consciousness and Persons: Identity and Unity*, Cambridge, MA: MIT Press.
- Watzl, S. (2011) Attention as structuring of the stream of consciousness, in Mole, C., Smithies, D. & Wu, W. (eds.) *Attention: Philosophical and Psychological Essays*, pp. 145–173, Oxford: Oxford University Press.
- Wittgenstein, L. (1921/1961) *Tractatus Logico-Philosophicus*, Pears, D. & McGuinness, B. (trans.), London: Routledge.