Philosophical Psychology, 2013
http://dx.doi.org/10.1080/09515089.2013.776476

Routledge
Taylor & Francis Group

The modal breadth of consciousness

Michael Trestman

Downloaded by [Indiana University Libraries] at 14:47 30 September 2013

In this paper, I argue that the intentional structure of typical human conscious experience has "modal breadth"—that the contents of experience typically include alternate possibilities. I support this claim with analyses of conscious mental processes such as the perception of temporally extended events, persistent objects, and causality, and the experience of bodily agency. While modal breadth may not be strictly necessary for consciousness per se, it is essential to many cognitive processes that are pervasive and functionally important to normal human consciousness.

Keywords: Agency; Causality; Cognition; Consciousness; Modality; Object Perception; Perception; Temporal Consciousness; Time

1. Introduction

Although our conscious experience of the world is incredibly rich, detailed, vivid, and varied, it can also seem deceptively simple or straightforward in some ways. In the normal run of things, we take our perception of a stable world of persistent objects located in a well-behaved spatiotemporal order for granted. Objects move and change, of course, but they do so in lawful, well-behaved ways that are to some extent explainable and predictable according to our intuitive understanding of causation and some general commonsense laws of nature. Coffee mugs do not blink into or out of existence, nor do they teleport across the room or fly through the air of their own volition, nor do they inexplicably grow or shrink, nor do they suddenly transmogrify into trout or petunias.

This is all very well and good, and no doubt much of the lawfulness and stability of our perceptual world is due to the lawfulness and stability of the world that we perceive, as it exists apart from our perceptions. However, there is also a lot of cognitive work, a lot of information processing, that goes on "behind the scenes" of our conscious awareness, as it were, that is absolutely essential to the stability, intelligibility, and meaningfulness of our perceptual world. Some aspects of these secret machinations are understood, at least in part. For example, Merker's (2005)

Michael Trestman is a National Science Foundation Postdoctoral Research Fellow, Indiana University
Correspondence to: Michael Trestman, Cognitive Science Program, 819 Eigenmann, Indiana University, 1900 E. 10th St., Bloomington, IN 47406-7512, USA. Email: michael.a.trestman@gmail.com

© 2013 Taylor & Francis

account of the evolutionary origins of consciousness is grounded in a rich and insightful analysis of the neurocomputational processes that integrate the information from multiple sense modalities, and compensate for the sensory transformations caused by an animal's own bodily movements in order to create a stable awareness of self-in-environment suitable for guiding action. However, much of what goes on "beneath the surface" of consciousness is poorly understood. In this paper, I will argue that the nature of normal conscious experience implies that cognition includes representations that are "modally broad"—they include among their contents non-actual (counterfactual) possibilities. This is not just the case when we explicitly reflect on how things could have been, but rather is a pervasive and fundamental structural feature of normal consciousness, and may well be an essential feature of cognitive systems that are conscious in anything like the way we are.

Downloaded by [Indiana University Libraries] at 14:47 30 September 2013

2. Preamble: Phenomenal Content and Intentional Content

The relationship between intentional content (what, if anything, a mental state represents) and phenomenal content (what, if anything, it is like to be in a mental state) is rife with controversy. Here, I will assume only a relatively weak thesis, that *some* features of the intentional content of a given mental state determine some features of its phenomenal content—some differences in what an experience represents make a difference for what it's like to have it. I do not assume that the intentional contents of a mental state wholly determine its phenomenal contents—there may be aspects of phenomenal content that have no intentional content. Nor do I assume that all aspects of the intentional content of a mental state make a difference for phenomenal content—some of the representational content of a mental state might be inaccessible to consciousness (this could be either "internal" or "narrow" content that for one reason or another doesn't become conscious, or "external" or "wide" content that in principle couldn't become conscious because it is the wrong sort of thing). Nor do I assume here any particular theory of content (i.e., what makes a mental state "about" something in the world), or of the metaphysical or explanatory relationship between intentionality and phenomenality.

Before moving on, I will say that I think that we probably need to be pluralists about intentionality and about content: a notion of intentionality is a powerful, if not absolutely essential, conceptual tool for making sense of intelligent behavior, including language and communication; but in this context, the relevant notion of "aboutness" is inseparably bound to causal and historical relations between agents and their external world. On the other hand, careful attention to conscious experience reveals that it is loaded with intentionality in its own right, apart from considerations of behavioral disposition or history (Horgan & Tienson, 2002; Husserl, 1966/1991; Merleau-Ponty, 1945/1962; Zahavi, 1999). These different sorts of considerations yield views of intentionality and content that may be genuinely distinct (i.e., externalism versus internalism, wide content versus narrow content). If this is right, then, while the different views might be irreconcilable in the sense that neither can subsume or reduce

to the other, they will turn out to be complementary, making non-redundant contributions to a full understanding of the mind. Indeed, there are probably deep substantive questions concerning the relation between them. For example, evolutionary and developmental explanations of how creatures like us have come to experience the world in a way that is laden with intrinsic intentionality no doubt have much to do with the demands on those creatures for intelligent, object-oriented, goal-directed interaction with the environment.

In any case, I do think that consciousness is characteristically (if not necessarily) intentional, in that experience presents the world as being a certain way to the subject—to deny this is just to operate with a distorted picture of the nature of conscious experience. Experience is characteristically *about* the way the world is; it presents the subject with a world. We don't experience bare, meaningless sense data like red and green patches; we experience the presence and nature of objects and events in our environment, and their relevance for our own felt urges and preferences. This implies that at least some of the intentional contents of conscious mental states are experienced, and hence that they make a difference for the phenomenal content of the experience. I will argue here that modal breadth (the inclusion of a range of alternative non-actual possibilities) is not just an intentional property of (some) experiences, but that it is the sort of intentional property that often makes a difference for phenomenal content. In some cases, this is because an awareness of the alternate possibilities is experienced, and perhaps in other cases because though it is not experienced itself, it is required as support or substructure for intentional contents that *are* experienced. What I am after is an understanding of the *intentional structure of conscious experience*—what intentional properties play what roles in constituting the sorts of experience we have.

Toward this end, we must distinguish between an intentional content of an experience playing a relatively *explicit* or *implicit* role in determining the phenomenal content of the experience. An experience may have intentional contents that are themselves experienced; on the other hand, it may have contents that are not directly experienced, but play some other role in determining the phenomenal content of the experience. For example, suppose you hear a song that you heard once on a very happy occasion, and hearing the song makes you happy. The role of the happy memory might play an *explicit* role in the experience—you consciously experience the memory in some way, perhaps as some distinctive images of the place, people, and events involved, together with a realization that the song was playing, and that this association is why you currently feel happy. At the other extreme, perhaps you don't consciously remember any of the specifics during the course of this experience, or make the connection—you just feel happy. In this latter case, your memory of the happy occasion plays an *implicit* role in determining the phenomenal content of the experience of hearing the song (it feels happy), but its intentional content is not itself experienced (you don't experience any sort of recollection of the occasion).

The implict/explicit distinction regarding the role intentional contents play in determining the phenomenal content of an experience should be seen not as binary but graded. There is probably a large spectrum between the above extremes,

Downloaded by [Indiana University Libraries] at 14:47 30 September 2013

concerning how a memory of a happy occasion can contribute to the experience of hearing a song. For example, you might experience some fleeting images of the place or people involved, without putting the whole thing together, in which case the memory's contribution might be best described as partially explicit, somewhere between the extremes above. A full discussion of this issue is beyond the scope of this paper, but the basic distinction will be relevant for some of the discussion below.

3. Temporally Extended Events, Surprise, and Uncertainty

I think that the modal structure of conscious cognition is intimately related to its temporal structure, and so I will begin with a discussion of the conscious awareness of temporality. There is an old argument in the literature, dating at least to Husserl and Brentano (Dainton, 2010), that our experiences often (or always) have *temporal depth*—i.e., that they include an awareness of a brief temporal interval, centered on the experienced present but extending back into the experienced past and forward into the experienced future.[1]

Premises:

1. To experience an event or process as unfolding over a temporal duration is to have an experience with temporal depth.
2. We often experience events and processes as unfolding in time.

Conclusion: therefore, we often have experiences with temporal depth.

The first premise is supported by analysis of the classic example of the experience of listening to a melody. In one sense, when one listens to a melody, one only *hears* a single note at a time—only one note is actually impinging on the senses and is experienced as currently occurring. But in another sense, one can certainly hear a *melody* as such, that is, precisely an ordered sequence of notes extended over time in a distinctive way. The melody of which it is a part influences the experience of hearing the note that is currently playing at any given time—a single note played by, e.g., a violin can *sound* (in the sense of *how it sounds to a listener*, i.e., what the listener's experience of hearing it is like) one way or another, depending on what note(s) preceded it. For example, it can satisfyingly resolve a tension established in a chord progression, or sound discordant and off-putting. The previous notes, although gone, still play a role in the experience of the current note. The past notes are experienced *together* (co-experienced) with the current note—they are *part* of the content of the experience of hearing the current note *as* a note in a melody, and hence hearing the melody itself.

However, the notes in the melody are not experienced *as happening* together with the current note—that would be a different experience, an experience of several notes being heard together. The past notes are experienced *as past*—as having preceded the current note in sequence, one after another. If at one time we hear, say, a C, and then a moment later we hear a D, we will experience the D as having been preceded by the C. But if, a moment after that, we hear an E♭, we may hear the E♭ as being preceded by

Downloaded by [Indiana University Libraries] at 14:47 30 September 2013

the D, which itself was preceded by the C. Both the preceding C and D are "retained" in the experience of hearing the E♭, but they are not retained in exactly the same way—the C is retained as preceding the D, as having already been past when the D was heard. Such a sequence of iterated retention constitutes an immediate subjective past.[2] In this way, a temporal sequence of events is experienced as such—that is, all experienced together but *as* stretched out over a temporal duration, i.e., *as a temporally extended sequence.* When we hear a melody, we hear it as extended in time, and any description of the contents of such an experience must capture that fact.

Not only the lingering awareness of previous notes but also the expectation of what notes are to come next may make a difference in our experience. Convention or association from experience might make us expect that such and such a note is to come next, for example to relieve a harmonic tension. Often such an expectation of the next note is *part* of the experience of hearing a note as part of a melody. This is a clear example of temporal depth, since the experience of hearing the note as part of the melody *includes* the lingering awareness of the previous notes and the expectation of what is to come.

In support of premise 2, the lesson offered by this example is supposed to generalize. Just as, in order to hear a note as part of a melody, one must retain previous notes as just having been heard and expect upcoming notes to be heard, in order to experience many kinds of temporally extended processes, we must hold retentions and expectations together with the impression of what is happening at the moment. To take another simple example, suppose you are watching a ball roll across a table to your left, and at a given moment, it is in the center of the table. This is a different experience from seeing the ball sit still in the center of the table, or from seeing the ball in the center of the table as it rolls to your right—despite the fact that in each case the ball is currently in the center of the table. If these experiences are different, it must be because an awareness of the past and future of the ball are *part* of the content of the experience. An experience of the ball *as moving* implies an awareness not only of its current location, but of its previous and expected locations as well.

So far I have discussed the idea that the content of an experience of perceiving a temporally extended process is at least partly constituted by retentions of previous sensory impressions as well as expectations for the way sensory impressions will change over time. We experience our current impressions as situated in a temporal stream of changing impressions, and this is essential to the intentional structure of experience, which presents us with a world of temporally extended processes and persistent objects. But as we shall see, similar considerations support the conclusion that consciousness is modally broad as well as temporally deep.

To start with, I would argue that not only are previously retained impressions further retained, as discussed above, to form a sequence of iterated retentions stretching back and constituting a subjective past, but as well it seems that previous *expectations* or "protentions" are retained as well, to yield a more complex temporal structure. This seems evident in the experience of surprise, which I think is best understood as an experience of the violation of prior expectations. An awareness that one's prior expectations were violated implies a tacit awareness of those expectations,

Downloaded by [Indiana University Libraries] at 14:47 30 September 2013

and of a mismatch between those prior expectations and what one now experiences or has just experienced.[3]

However, while surprise seems to imply some awareness of prior expectations, and the mismatch between them and how things turned out, it is a further question whether those prior expectations are *part* of the current experience of surprise, or whether they can contribute to the conscious experience without being themselves conscious. The latter would be the case if a mismatch between prior expectations, which were not themselves conscious, caused one to experience a sensation of surprise that did not represent anything about *what* was surprising—this sensation of surprise could have the content that *some* prior expectation was violated without including the prior expectation among its contents. Do prior expectations play an implicit or an explicit role in the experience of surprise? That is, are prior expectations themselves unconscious components of the intentional content of the experience, or are they actually experienced, i.e., are they components of the phenomenal content of the experience as well?

I think that typically, but not always (and therefore, not necessarily), the prior expectations are an explicit part of the experience of surprise. When one experiences surprise, one is typically consciously aware—indeed keenly focused—on what is surprising, i.e., on just what violates one's prior expectations. This will be discussed further below, as it bears on our discussion of modal breadth, but I will point out here that one reason to think this is that the alternative—to experience surprise without being aware of *what* is surprising—is very disconcerting. I have in mind cases like first seeing a familiar person who has changed something drastic about her appearance, and knowing that something is different but being unable to identify it. In this case, there is surprise at the discrepancy between expectation and perception, coupled with an added dimension of strangeness at being unaware of those expectations, of which one is usually aware.[4]

The experience of surprise seems to involve an awareness of the violation of prior expectations, and therefore at least a minimal element of modal breadth—one previously expected the present to be one way, and it turned out to be different. But we soon discover that the modal structure is richer and more complex than this.

Our experience of temporal processes is grounded in a structure of retentions and protentions. But are these retentions—and more to the point, our protentions—singular and linear? Are our expectations structured as a single, ordered sequence of impressions stretching forward into the future? I think not. Often, our experience of the expected future is *open* in such a way that its intentional content cannot be structured as a single linear sequence. This openness can take at least two different forms, characteristic of different sorts of experience, and involving different sorts of intentional structure: uncertainty and controllability (involved in the experience of agency—to be discussed later).

We probably experience several different types of uncertainty, some of which may not involve modal breadth, but I think there is a distinctive type of uncertainty involved in perceiving the unfolding of quasi-predictable processes. By 'quasi-predictable' I mean not totally unpredictable, but not totally predictable either. Return to the example of the ball rolling across the table, but now suppose that it is a billiards

Downloaded by [Indiana University Libraries] at 14:47 30 September 2013

table. The ball is rolling towards a corner pocket, but it is going slowly and the angle is a little funny; will the ball drop, or will it stop short, or will it bounce off the edge and back toward the center of the table? When the ball is still all the way on the other side of the table, you can't tell—these are all open possibilities as far as you can see. Initially, your expectations are divided among the three possibilities, but as you watch the ball roll toward the pocket, your expectations for the future will resolve or collapse onto a single expectation. For example, perhaps as it gets close, you can tell it has enough momentum to make it to the pocket, and you can tell the angle is okay. Now, how can we make sense of the intentional structure of this experience? To experience the ball *as rolling* implies an awareness of the future of the ball—farther along toward the pocket than where it is now. And indeed, an awareness of the ball as *rolling toward the pocket* implies an awareness of a sequence of future moments, at each of which the ball is seen to be closer to the pocket. If one saw the ball *as bound to drop into the pocket*, that is, if in watching the ball, one had a clear, certain expectation that the ball will drop, this might be understood as a single linear sequence of expected momentary impressions of the ball approaching the pocket and then falling in. This would be an experience with temporal depth but no modal breadth. And if the uncertainty involved just were a kind of doubt or lack of confidence that the ball will go in, the intentional structure of the experience might be understood as just a particular kind of attitude or feeling of doubt towards that linear sequence.

Perhaps in some cases this is how things are. But in watching the ball roll toward the pocket, aware that either it will stop, drop, or bounce back, but unaware of which, it seems that at least three distinct sequences of expected impressions are held in mind as possible, but non-compossible (and therefore not possibly all *actual*), alternate futures. The (at least) three futures in this range of possibilities are among the contents of such an experience, and they seem to play a direct role in the nature of the experience. It might be possible to experience a kind of totally open uncertainty—not uncertainty about which of a small number of possible sequences will ensue, but a total, blank openness toward the future—but I think this is atypical. Indeed, the contrast of that blank openness with experiences like the one described reveal that an experienced uncertainty between multiple possible future sequences of events includes among its explicit (i.e., experienced) contents awareness of those sequences of events. Such an experience is, therefore, not just modally broad but explicitly so.

4. The Experience of Agency

Our experience of the future has a different kind of openness in the context of the experience of agency: *controllability*. The basic idea is that to experience agency is, at least in part, to be aware of multiple alternate possible futures as possible, in a way that is under one's control. This idea is illustrated vividly in the following anecdote by Temple Grandin:

> On a bright, sunny day, I was driving to the airport when an elk ran into the highway just ahead of my car. I had only three or four seconds to react. During those few

Downloaded by [Indiana University Libraries] at 14:47 30 September 2013

> seconds, I saw images of my choices. The first image was of a car rear-ending me. This is what would have happened if I had made the instinctive panic response and slammed on the brakes. The second image was of an elk smashing through my windshield. This is what would have happened if I had swerved. The last image showed the elk passing by in front of my car. The last choice was the one I could make if I inhibited the panic response and braked just a little to slow the car. I mentally "clicked" on slowing down and avoided an accident. It was like clicking a computer mouse on the desired picture. (2000, p. 17)

As Grandin describes her experience, it is clearly modally broad, and in as explicit a way as could be. But, even assuming that it is an accurate description of her own experience in this case, is this typically the way we experience agency, as a set of possibilities, out of which we can select one to "click on" and enact? It is worth noting that Grandin thinks that her experience is *atypical*—she is autistic, and thinks that her highly visual, detail-oriented thought processes are substantially different than those of typical adult humans, the latter supposedly being more language-oriented and abstract. Grandin suggests that her imagistic experience of possible futures is not something that most of us share, and hence that *typical* human experience is not explicitly modally broad, at least in quite the way she describes. If she is right, do we experience the future as open (modally broad) in some other way, or not at all?

The possibility of substantial interpersonal differences is a general problem for consciousness studies—if you and I carefully attend to and describe our experience of a similar-seeming psychological event, and the descriptions differ, how do we know if this is because our experiences really differ, or if one of us is misdescribing the experience (Schwitzgebel, 2011)? Grandin's description of the experience of agency as a choice among explicitly represented possible futures strikes a chord of familiarity with me. However, if this sounds right for both her and me, but not for most people, is this because I am an unusually "visual" thinker, as she is? Or is it just because we have spent more time thinking about the experience of agency than most people, and have adopted similarly theory-laden ways of conceptualizing it? It is well known that there are significant individual differences in cognitive style, and various abilities, and these may well correspond to substantial differences in aspects of our experience. I think that a person who can listen to an orchestra play a symphony and then produce an accurate score for the entire performance has something very different going on cognitively than what happens for me—very plausibly, such a person experiences the symphony quite differently. But how can we tell?

To say the least, this is a thorny issue. I will try to focus on what I take to be the simplest, most cognitively fundamental form of the experience of agency, and also the form which seems to me to most clearly involve modal breadth—*bodily* agency. As I will argue below, bodily agency pervades perception and cognition, at least implicitly and probably, for most people, explicitly as well. Because bodily agency is modally broad, modal breadth is therefore a pervasive feature of conscious experience in humans—and likely anyone else who has it.

Downloaded by [Indiana University Libraries] at 14:47 30 September 2013

Downloaded by [Indiana University Libraries] at 14:47 30 September 2013

5. Bodily Self-Awareness

As conscious embodied subjects, we are aware of our bodies in (at least) the following two distinguishable ways: first, we are aware of our bodies as physical objects in a world of physical objects. We can see our arms and legs, for example, and see how they interact with other objects in our surroundings. But, secondly, we are also aware of our bodies as the *medium* of our perception of and action upon the world. Our bodily self-awareness is uniquely "subjective" or "subjectual" in the sense that our bodies are "of the subject"—literally part of oneself *as* subject: what happens to my hand happens *to me*; what my hand does, *I* do. Of course, this is patently not true in a case like alien hand syndrome. But the fact that this pathology consists in a loss of the subjective experience of the hand—the hand is experienced as *other,* as beyond control—supports my point about the "normal" experience of the body.

One's hand is a physical object, and one can be conscious of it as such. I can look down and see my hand resting on a table next to a cup. But unlike the table and the cup, my hand is also an organ of perception by which I experience my environment, and an organ of action, by which I can move or grasp another object. And much the same is true of the entirety of my body (Zahavi, 1999, pp. 92–103).

The awareness of one's own body's potentiality for action and perception seems clearly to be inherently loaded with temporal depth and modal breadth, at least implicitly. An awareness of one's body as medium of action implies a temporally forward-looking component. If I am consciously aware, for example, that I can move my hand *thus* (in some particular way, such as to grasp a cup that is before me), my awareness is future-oriented, since it is an awareness of something (the movement) that occurs in the future, if it occurs at all. Therefore, to the extent that the experience of bodily agency pervades human consciousness—plausibly, quite a large extent, especially given the role of bodily agency in normal perception (discussed below)—human consciousness is pervaded by temporal depth. Further, it seems compelling to me that modal breadth is essential to bodily agency, and hence that modal breadth as well is a pervasive feature of consciousness.

An experience of bodily agency is an experience of one's body as a medium of action, as something that is under one's direct control. It is, for example, to feel that one can move one's hand in a particular way, that one can reach for something. But to experience one's body as under one's control is precisely to experience a plurality of alternate futures as possible—I can move my hand *this* way or not. If there is only one "choice" of how to move, then it is of course not really a choice at all—if there are not multiple possible movements, then one is not in control. The experience that one's hand *will* move thusly is not an experience of control, but a prediction—something one can experience regarding a bird or a trashcan as well as one's own body. An action experienced as under one's control as an agent is experienced as chosen from among a range of possible actions. This implies that the experience of bodily agency is modally broad. Of course, modal breadth is not *sufficient* for agency—as I argued in the last section, watching a bird fly or a trash can roll down a driveway probably involves modal breadth as well—but it seems to be *necessary* in a deep way. What more is

involved in agency, what sort of attitude or functional role distinguishes possibilities experienced as *choosable* as opposed to possible but not under one's control, is a complex issue that merits discussion beyond the scope of this paper.

One might well think it plausible, even without further argument, that this awareness of one's body as medium of action—what we might call the experience of *bodily agency*—is ubiquitous in normal consciousness (or perhaps even essential to it), or at least consciousness that is anything remotely like our own. Moreover, the temporal depth and modal breadth inherent in the experience of bodily agency can be seen to permeate consciousness further when we consider the intimate relation between this kind of experience and the perceptual experience of objects.

Downloaded by [Indiana University Libraries] at 14:47 30 September 2013

6. Object Perception

This intimate relation between action and perception has been discussed at length by many authors in philosophy and psychology, such as Clark (1997); Gibson (1979); Grush (2004); Hurley (2002); Husserl (1966/1991; see also Zahavi, 1999, pp. 92–103); Merleau-Ponty (1945/1962); O'Regan and Noë (2001); Thelen and Smith (1996); and Varela, Thompson, and Rosch (1992). Indeed, some authors have argued for this reason that bodily agency is strictly *necessary* for perception (Noë, 2004), although we need not concern ourselves with establishing this strong claim.

Much of this literature has focused on the ways in which normal perception requires active exploration of the environment by the perceiving animal. Much of the information needed for vision, for example, is gathered in temporally extended processes of moving eyes, head, and body. In many cases, static snapshot-like images do not contain enough information to uniquely determine the spatial layout of a scene, and this problem is made even worse by the fact that our eyes can only focus sharply on a tiny little bit of the visual scene at once. Even if one is seated and one's head is held still, one's eyes move about to extract the relevant information. In fact, it seems that *action* per se (in the sense of motion controlled by the animal), rather than just motion, is required for the development of normal object perception (Held & Hein, 1963)—which suggests an important role for modally broad cognition as a building block in the ontogeny, and perhaps the constitution, of functional perception.

Many authors have also argued—and this is more to the point for our current purposes—that much of the content of object-perception concerns *affordances*—potentials for interaction offered to the perceiver by the object, together with a sense of the motivational relevance of those potentials (Gibson, 1979). There has been a great deal of controversy over the proper interpretation of Gibson's notion of affordance,[5] as well as over Gibson's (1979) claim that the contents of perception are wholly constituted by affordances. Interpretive issues aside, I think that the notion as characterized above is the one used in contemporary cognitive science. Further, I am interested here only in the weaker claim that affordances are common or pervasive among the contents of object-oriented perception and cognition, perhaps in

combination with "classical" properties, because if affordances are pervasive among these contents, then so are modal breadth and temporal depth.

When one views, e.g., an apple, the content of the perception doesn't just concern "objective" or "action-neutral" properties like shape and color. Rather, one views it as graspable, edible, etc. Properties like shape and color are perceived in relation to the ways the perceiver might interact with it. The color isn't just seen as "red," but rather in connection to properties like ripeness and what the apple might taste like. There are any number of such affordances that one might see in an apple, and which ones actually make up the content of your perception of the apple depends on your current goals and what else is going on in your awareness. For example, if you are hungry, your perception of the apple will be different than if you are angry and looking for something to throw. In the former case, you will probably attend to more to things that indicate its edibility, and your perception will be largely constituted by imagining what it would be like to bite into the apple. In the latter, you might attend more to its shape and imagine how it would feel in your hand as you throw it, how heavy it is and what kind of impact it might have on the intended target. What pattern of light the apple casts on your retina will obviously inform these quick, largely automatic imaginings that make up your perception; but the content of your perception is a function not just of the information contained in the light that stimulates your eyes, but also of your own goals, interests, and awareness of your bodily capacities for action—the apple is perceived in relation to these aspects of the perceiver, and these relational aspects pervade the content of the experience.

To perceive the apple as graspable is to be aware of a *possible future* grasping (moreover, a possible future grasping that is *choosable*); to perceive it as tasty is to be aware of a *possible* (and choosable) *future* tasting. These generalizations apply to perception and cognition in many ways and at many levels. Many of them may concern unconscious cognition—that is, it may well be that modal breadth is a largely *implicit* feature of cognition, playing a role "beneath the surface" of consciousness. But I think it probably plays an explicit role in conscious experience as well. The experience of seeing an apple when one is hungry is no doubt different than seeing an apple when one is angry and looking for something to throw, and part of that difference lies in the content of one's perception of the apple: what one notices about it, what features of the apple one is aware of.

It might be objected that, surely, not all perceptual experience is quite so goal-directed as these examples would suggest, but the perceived properties that make up the content of the perception of an object are embodied and action-oriented—hence *oriented toward a range of possible, but non-compossible, futures*—in more subtle ways as well.[6] This is because the perception of an object is at least partly constituted by expectations of how *else* the object might appear (Grush, 2004; Noe, 2004; Zahavi, 1999, pp. 92–103). When one has a normal perceptual experience such as, for example, seeing a mug, the content of that experience is not just the *image* of the mug, i.e., how the mug now appears, but also includes an implicit understanding that there is something there (the mug) which now appears just so, but could also appear differently. That is, it is an experience of the appearance *of* a mug, which *could* appear

Downloaded by [Indiana University Libraries] at 14:47 30 September 2013

differently (either in the same or in a different sense modality), and so *transcends* the way it actually appears at a given moment. For example, perceiving the mug as having a particular shape might be constituted by anticipating that if one moves one's head a certain way, then the image of the mug will deform in a particular way, or that if one were to grasp it, it would feel a certain way. Seeing the mug *as hot* (in virtue of seeing steam rising from it) might be constituted by an expectation that it will feel a certain way if it is touched, or even that another person who is seen reaching for it will appear to handle it carefully.

This is what it is to *perceive an object*—to interpret a sensory experience (such as a seen image) as the appearance of an object that transcends this appearance. The experience of perceiving an object is an experience of some *thing* appearing, a thing that could appear otherwise.[7]

To return to a previous example, if I look at an apple, my experience of seeing it *as an apple* is at least partly a matter of (implicitly) interpreting the image or visual appearance with which I am confronted as an appearance *of something* that could appear differently. My expectations for how the apple could appear differently in part constitute the intentional content of my perception of the apple. For example, if I see it *as round*, this is in part a matter of certain expectations of how the image of the apple will deform (e.g., smoothly, without clear, discontinuous edges) if I move my head or rotate it manually, and perhaps in part a matter of my expectation of how it will fill my hand if I grasp it. Seeing it as having a certain weight is in part a matter of my expectation of how much it will resist my effort to lift or accelerate it; seeing it as ripe and tasty is a matter of my expectations of how it will feel and taste in my mouth if I bite into it.

I submit that often, and probably almost always, these expectations are explicit; they make a difference for the conscious experience of the perception. A different set of sensory expectations of this sort would make for a *different experience* of seeing the apple (e.g., seeing it as a convincing-looking, but hollow and unconvincing-tasting, plastic apple). Moreover, these expectations are (at least usually) not definite and linear, but have some modal breadth to them, reflecting at least some degree of uncertainty in the content of the perception. For example, say I want to eat the apple and reach for it to feel how ripe it is. If I just bought the apple yesterday at the store, I may expect it to be pretty firm and hard, whereas if it has been sitting on the counter for a month, I will probably expect it to be pretty squishy. But if it has been there for a few days or a week, neither a crisp feel nor a squishy feel would be particularly surprising; my expectations seem to be distributed over a range of ways the apple could feel when I squeeze it. Similarly, when I look at one side of the apple, I expect that if I turn it, I will see smooth, continuous deformation of the visual contour of the apple, rather than a sharp edge, or an eyeball, or a tiny steering wheel. But a blemish or brown spot would be far *less* surprising. My expectations seem to be distributed (unequally!) over these possibilities. And again, the distribution of my expectations over the salient possibilities seem to partly constitute the content of my experience of seeing the apple. The uncertainty in the content of my perception of the apple—whether it is ripe or rotten, real or plastic,

Downloaded by [Indiana University Libraries] at 14:47 30 September 2013

etc.—is partly constituted by my divided or distributed expectations of how my sensory experience of the apple will unfold in the immediate future. To the extent that this uncertainty is *experienced* as a division between multiple possibilities, the multiple possibilities are an explicit part of the conscious experience. So at least to the extent that this sort of uncertainty is characteristic of conscious experience, consciousness is modally broad.

7. Perception of Causality

If one thinks that causation is a matter of some sort of conditional dependence among events, an assumption which is shared among a very broad range of views of causation (e.g., Lewis, 1973, 2000; Menzies, 2008; Woodward, 2003), then the fact that our perceptions of objects and processes are modally broad helps to demystify how we can experience causal relationships. The fact that we seem to perceive causal relationships has puzzled philosophers at least since Hume, who noted that we never observe causation directly, nor infer it logically; all we ever observe is the "constant conjunction" of objects or events (Hume, 1999; Morris, 2009). But if part of what it is to perceive an event is to be aware of a range of ways it might be unfolding/might have unfolded, and part of what it is to perceive an object is to be aware of a range of its potentialities for disclosing hidden aspects to us—for changing and interacting with other objects, etc.—and if causation is some sort of relation of conditional dependence among possible events, then it becomes much clearer how, in simply observing a sequence of events involving objects, we can perceive causality. On this view, Hume's mistake—which set up the puzzle—was in mischaracterizing the nature of perceptual experience, by denying its modal breadth, due to his commitment to thorough empiricism about human psychology.

To run this backwards as an argument for modal breadth:

> Premises:
>
> 1. We do have experiences *as of* the perception of causality, i.e., we *seem* to perceive causality (this is meant to be less metaphysically committal than saying that we *perceive* causality in a factive sense of 'perceive'—it is a claim about the contents of our experience, not whether that experience accurately reflects reality).
> 2. Causality has something to do with relations of conditional dependence among events; a fact about whether x causes y is a *modally broad fact*, in that it has to do with a set of events, some of which are non-compossible with others.
>
> Conclusion: Our experience of causality (i.e., as of the perception of causality) is modally broad.

The contents of our experience of causality include relations among a set of events, some of which are non-compossible with others. Our experience of causality is modally broad. So at least to the extent that experience of causality is characteristic of consciousness, consciousness is modally broad (contra Hume).[8]

Downloaded by [Indiana University Libraries] at 14:47 30 September 2013

8. Is Modal Breadth Essential to Consciousness?

So far, I have argued that modal breadth is a pervasive feature of conscious experience, but I have not touched on the issue of whether it is *essential* or *necessary* to conscious experience per se. Is there any reason to think this? Further, even if it is not a strictly necessary condition for conscious experience, is there some reason to think that it is at least *important* in some non-contingent way? Or is just a quirk that it happens to pervade our conscious experience to such a large extent?

To start off with, even if there is no good reason to think that modal breadth is strictly necessary for consciousness per se, I think it is necessary for consciousness that resembles anything like our consciousness. This follows from what I have argued throughout the paper—that modal breadth is a pervasive feature of human conscious experience, and is essential to many of the core features of conscious cognition, including perception of temporally extended processes, object perception generally, perception of causality, and the experience of agency. A consciousness that lacked modal breadth would lack all of these features, and hence be very alien to us. I will discuss this further below.

It may be that, while modal breadth is not strictly necessary for consciousness, temporal depth *is*. Philosophers of mind going back to James (1890) have held that the temporally extended, *streaming* nature of consciousness is essential to consciousness, and Husserl argued for this point at length. While I will not concern myself with the fine points of Husserl interpretation (see Zahavi, 2005 for discussion), the argument, which I find compelling, seems to be as follows:

What is essential to consciousness is its *subjective* nature: it is not just representation, but representation *for* a subject; not just information gathering or processing, but information that is presented *to* a subject. But how can a state have this dimension? How can a single state have both representational content, and a subject to whom the content is presented? The articulated temporal structure described in the above sections is a solution to this structural requirement. In holding previous, current, and expected impressions (together with the retentions of prior retentions and protentions) together, the articulated sequence itself emerges as something that transcends any of the momentary impressions that can be present for it at a time. The articulated, sequential nature, on this view, constitutes the *streaming* of consciousness, and it is the stream itself that *is* the subject, most fundamentally. The momentary impressions come and go, entering the stream and passing away after being retained for a few iterations. But the stream *transcends* any of these impressions, and is in a sense aware of that transcendence precisely because of the multiplicity of impressions, and their transience, which is represented in the sequential asymmetry (iterated retention/protention) of the structure.

If consciousness were possible without this sort of streaming, iterated sequential structure, what would it be like? It could only be temporally punctate—a disconnected set of unitary experiences—but not organized into any kind of persistent subjectivity. Experience at one moment would not be unified in any way with experience at another moment; they would not be parts of *the same* experience in any way. It is the streaming nature of consciousness that integrates punctate momentary impressions into a single

Downloaded by [Indiana University Libraries] at 14:47 30 September 2013

process of experiencing the world as it changes over time. Without this temporalized, streaming structure, there could be no identity of subjectivity over time—no persistent *subject* for whom various experiences *are* experiences. But that is impossible, if what is distinctive about consciousness is precisely its subjectivity. A conscious experience *must*, by its nature, be *for* a subject. And it is in the temporal extension of an articulated retention-impression-protention sequence that that *transcendent* subject (i.e., a subject that is more than just the content of its representations) is constituted.

If this is right, then temporal depth might be a necessary property of cognitive systems that are conscious. At the very least, it seems like a prerequisite for identity over time for conscious systems; the relations of iterated retention are what make an experience I had a few moments ago *my* experience at this moment—it is retained as part of my current experience.[9]

I would argue that temporal depth might be essential to another core feature of consciousness that has not been discussed so far: affect. Affect is the positive (attractive, pleasurable, desirable, soothing, etc.) or negative (aversive, painful, fearful, etc.) valence or evaluative component of experience. While it may not be strictly necessary for consciousness per se, a consciousness that lacked affect would be starkly minimal. Moreover, on a broad range of views about the foundations of ethics, affective consciousness of one sort or another is thought to be essential to *moral worth*, such that the only candidates for moral worth would be beings capable of affective consciousness of some sort, for example of suffering, of doing well or badly, or of having preferences. This makes affective consciousness of special interest, even if it is non-identical with consciousness in its most basic sense.

I would argue that experiences that have an affective, evaluative, or motivational component are intrinsically temporally deep, particularly in the future direction. The basic idea is that a constitutive part of experiencing something as painful (or nasty tasting, gross smelling, socially uncomfortable, or otherwise unpleasant) is *wanting* it to stop or to not happen again, whereas a constitutive part of experiencing something as pleasurable (or delicious, nice-smelling or feeling, socially rewarding, or otherwise pleasant) is *wanting* it to continue or happen again. If one didn't want something to stop, it wouldn't really be painful, and if one didn't want it to continue or recur, it wouldn't really be pleasurable.[10] Wanting something to stop or to continue seems to refer to a (possible) future event—the cessation or continuation of the sensation or event.[11] So if it is right that something like this wanting is inherent to being affective, then any experience that is affective has temporal depth.

But is affect inherently modally broad as well as temporally deep? It might be thought that affect is inherently *contrastive*. The idea here is that for some sensation to be laden with affect is precisely for the experience of it to include a conscious preference that it cease (if it is negative affect) or continue (if it is positive), and that this sort of preference is a matter of comparing two possible alternative futures, one in which the sensation ceases, and one in which it continues. Something more elaborate, but fundamentally similar in structure, would have to be true of affective experience regarding not simple sensations but events. So fearing something would be preferring

Downloaded by [Indiana University Libraries] at 14:47 30 September 2013

that it not happen, etc. If something like that is right, then modal breadth as well as temporal depth might be essential to affective consciousness.

These considerations may give us reason to think that temporal depth is strictly essential to consciousness, but what of modal breadth? The features discussed above, such as object perception, perception of causation, affect, and the experience of bodily agency, are pervasive features of human consciousness, as I have argued. But moreover, I think that these features are also very close to the core function of *cognition* broadly, as it occurs in nature, so it is likely that any consciousness that evolved naturally would be modally broad. Cognition serves to integrate sensory information for control of behavior. That is its *function*, in a sense that explains why animals with the capacity for cognition have survived, proliferated, and diversified—they did so because cognition allows for adaptive, sensory-information-driven control of behavior. That is not to say that all possible cognitive systems must be modally broad. Indeed, some cognitive systems, or at least behavioral control systems, may be at least possible in principle that do not include integrated representations at all, such as those organized as a "subsumption architecture" (Brooks, 1991).

A cognitive system that is capable of constructing detailed representations of objects, of tracking causal relationships, and of comparing different alternative courses of action (for example, in terms of their expected hedonic consequences), is a powerful, and *general,* solution to the problem of bodily control. Although to argue for this point is beyond the scope of this paper, I argue elsewhere (Trestman, forthcoming) that all animals with large, active, complexly articulated bodies and distal perception (e.g., vision) do, as a matter of fact, employ integrated, spatially explicit, object-oriented cognition.[12] This is evident not just from the behavioral sophistication of animals in all major taxa that include animals with large, active, complexly articulated bodies and distal perception, but also from the adaptive flexibility of their behavior, over both ontogenetic and phylogenetic timescales.

That is not to say that all of these animals must be conscious; that is a separate question. Perhaps consciousness requires extra cognitive properties, such as temporal depth—as I have just argued. However, these considerations of functional requirements on cognitive control systems are relevant when we consider *how* consciousness might arise in nature. For example how might a cognitive system with temporal depth (perhaps the only sort of cognitive system that could be conscious) arise?

The perception of objects, perception of temporally extended processes, affect, and awareness of the body's capacities for action (i.e., bodily agency) all have clear adaptive value in organizing sophisticated, flexible, intelligent behavior for animals with complex bodies and distal senses. But if something along these lines is correct, than the adaptive value of these features may help to explain the evolution of temporal depth, and hence the emergence of consciousness.

If modal breadth *is* implicit in the perception of objects and temporally extended processes and in bodily agency, and if these were among the features which drove the evolution of consciousness (because these features are essentially temporally deep and temporal depth is essential for consciousness), then modal breadth might be

Downloaded by [Indiana University Libraries] at 14:47 30 September 2013

"evolutionarily essential" to consciousness in the sense that consciousness would not have evolved without being modally broad. In other words, even if it is temporal depth that is truly essential to consciousness, it might be that temporally deep cognition is only adaptive if it is also modally broad—and so modal breadth can be expected to be ubiquitous as a feature of consciousness. This would of course fall short of modal breadth being strictly necessary for consciousness, but it would still be a significant fact about consciousness as it exists in nature.

9. Conclusion

I have presented arguments here that consciousness is temporally deep and modally broad—that its intentional structure is that of a branching, iterated retention-impression-protention series. We may therefore call the model developed here the BItRIP model. However, this is only the beginning. It is to be hoped that further research may lead to an understanding of the temporal and modal structures of consciousness at a much finer and more detailed resolution. There has been some discussion of the fine temporal structure of consciousness (see Dainton, 2010 for a review), but not so for modal breadth. As well, much work remains to be done in relating these structural features as apprehended from a first-person, phenomenological perspective, to the third-person empirical reality of consciousness as it can be studied from the perspectives of neuroscience, evolution, and behavior.[13]

Acknowledgements

Many thanks to Adam Sennet, Bernard Molyneux, Bert Baumgaertner, and two anonymous referees for their helpful feedback on earlier drafts of this paper.

Notes

[1] For extended discussions of this argument and ways of interpreting it, see Thompson (2007, pp. 317–328) and Zahavi (1999, chapter 5).

[2] I say 'immediate' here to distinguish the phenomena of retention from the more intentionally complex phenomena of recollection or episodic memory.

[3] This argument from surprise is hinted at by Zahavi (2005, p. 57).

[4] I suspect that this happens most frequently with the perception of other people because of the specialized way we perceive faces and bodies—our attention is automatically drawn to facial expressions and body language that express emotions and intentions, rather the physical details.

[5] For example, Nanay (2010) claims that Gibson's notion of affordance was inherently normative—including a sense of what one *should* do with an object, rather than just what one *could* do. While I think Gibson thought that motivational relevance was a component of at least some perceived affordances, construing this as inherent normativity seems to me an overly strong and uncharitable reading. In any case, even if Nanay is correct about the interpretation of Gibson, I think the weaker notion is the familiar one that is a standard concept in cognitive science.

Downloaded by [Indiana University Libraries] at 14:47 30 September 2013

[6] I use 'perceptual experience' somewhat loosely; unlike some authors, I don't have a strong intuition about what it means for an aspect of an experience that involves perception to be *part* of the perception, as opposed to be a part of the experience that merely attends the perception rather than partly constituting it. I also don't think much hangs on drawing sharp boundaries around the part of an experience that is perceptual versus non-perceptual. Conscious experiences are heterogeneous, complex, and wildly varied.

[7] This understanding, or interpretation, by the perceiver, of an image in terms of an object is *not* done deliberately, inferentially, *after* perception, or even consciously—it is automatic, subconscious, and *constitutive* of normal perception and perceptual experience. Indeed, it is a somewhat artificial mode of perception to see *the image of the mug*, rather than the mug. Perhaps artists can train themselves to do this for the purpose of capturing an image in a visual medium, but it is clearly a distinct mode of seeing from normal visual perception, which has as its content the *object* putatively in the world, rather than the image of the object.

[8] I would like to point out that the disagreement with Hume is not on premise 1—he also thought that we *seemed* to perceive causal relationships, but on the conclusion, which is at odds with Hume's commitment to thorough empiricism, which implied that all of the contents of experience were composed of sense data. But if we relax that assumption, and allow the intentional contents of experience to be modally broad, we are able to account for the nature of the experience of causation in a more satisfying way.

[9] I am not suggesting that this suffices as an account of *personal* identity. Being a *person* might well consist in more than just being conscious, although plausibly consciousness is required for personhood.

[10] Although the case of the masochist, who takes pleasure in pain, is puzzling. Perhaps the right thing to say is that enjoyed pain is both painful and pleasurable, but *more* pleasurable than painful.

[11] There are various more specific theories, particularly focusing on pain, that capture this basic idea in various ways (Aydede, 2009). But many of them share a commitment to the claim that the content of an affective or evaluative experience is inherently future-oriented. Though there are alternative theories of pain that do not impute to it a motivational component as a necessary constituent (Aydede, 2009), it seems hard to deny that consciousness is pervaded by motivational aspects, such as immediate desires, anticipations (both eager and fearful), urges, and imperatives, which seem clearly to be future-oriented. One might well object that an account of pain that left out its temporal depth would fail to acknowledge just what is *painful*, i.e., aversive, motivating, affective about it.

[12] Similar arguments can be found in Merker (2005). Merker also argues that such functional benefits drove the evolution of consciousness, although he does not make the connection to temporal depth and modal breadth, which as I argue here, are crucial to bridging the gap between cognition and consciousness per se.

[13] For example, the BItRIP model can serve as a useful framework for understanding psychological phenomena such as chronostasis, subjective time-dilation, and variation in flicker fusion threshold (between human subjects, for the same subject in different contexts, or even between different species); see Trestman (unpublished manuscript).

References

Aydede, M. (2009). Pain. In E. N. Zalta (Ed.), *Stanford encyclopedia of philosophy* (Spring 2010 ed.). Retrieved from http://plato.stanford.edu/entries/pain/

Brooks, R. (1991). Intelligence without representation. *Artificial Intelligence*, *47*, 139–159.

Clark, A. (1997). *Being there: Putting brain, body, and world together again*. Cambridge, MA: MIT Press.

Downloaded by [Indiana University Libraries] at 14:47 30 September 2013

Dainton, B. (2010). Temporal consciousness. In E. N. Zalta (Ed.), *Stanford encyclopedia of philosophy* (Fall 2010 ed.). Retrieved from http://plato.stanford.edu/entries/consciousness-temporal/

Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston: Houghton Mifflin.

Grandin, T. (2000, Winter). My mind is a web browser: How people with autism think. *Cerebrum*, *2*(1), 14–22.

Grush, R. (2004). The emulation theory of representation: Motor control, imagery, and perception. *Behavioral and Brain Sciences*, *27*, 377–396.

Held, R., & Hein, A. (1963). Movement-Produced stimulation in the development of visually guided behavior. *Journal of Comparative and Physiological Psychology*, *56*(5), 872–876.

Horgan, T., & Tienson, J. (2002). The intentionality of phenomenology and the phenomenology of intentionality. In D. Chalmers (Ed.), *Philosophy of mind: Classical and contemporary readings* (pp. 520–533). Oxford: Oxford University Press.

Hume, D. (1999). *An enquiry concerning human understanding*. Oxford: Oxford University Press.

Hurley, S. L. (2002). *Consciousness in action*. Cambridge, MA: Harvard University Press.

Husserl, E. (1991). *On the phenomenology of the consciousness of internal time (1893–1917)* (J. B. Brough, Trans.). Dordrecht, NL: Kluwer. (Original work published 1966).

James, W. (1890). *The principles of psychology*. New York: Henry Holt.

Lewis, D. (1973). Causation. *Journal of Philosophy*, *70*, 556–567.

Lewis, D. (2000). Causation as influence. *Journal of Philosophy*, *97*, 182–197.

Menzies, P. (2008). Counterfactual theories of causation. In E. N. Zalta (Ed.), *Stanford encyclopedia of philosophy* (Fall 2009 ed.). Retrieved from http://plato.stanford.edu/entries/causation-counterfactual/

Merleau-Ponty, M. (1962). *Phenomenology of perception* (C. Smith, Trans.). London: Routledge Press. (Original work published 1945)

Merker, B. (2005). The liabilities of mobility: A selection pressure for the transition to consciousness in animal evolution. *Consciousness & Cognition*, *14*, 89–114.

Morris, W. (2009). David Hume. In E. N. Zalta (Ed.), *Stanford encyclopedia of philosophy* (Fall 2010 ed.). Retrieved from http://plato.stanford.edu/entries/hume/

Nanay, B. (2010). Action-Oriented perception. *European Journal of Philosophy*, *20*(3), 430–446.

Noë, A. (2004). *Action in perception*. Cambridge, MA: MIT Press.

O'Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences*, *24*(5), 939–972.

Schwitzgebel, E. (2011). *Perplexities of consciousness*. Cambridge, MA: MIT Press.

Thelen, E., & Smith, L. B. (1996). *A dynamic systems approach to the development of cognition and action*. Cambridge, MA: MIT press.

Thompson, E. (2007). *Mind in life: Biology, phenomenology, and the sciences of the mind*. Cambridge, MA: Harvard University Press.

Trestman, M. (forthcoming). The Cambrian explosion and the origins of embodied cognition. *Biological Theory*. Available at http://michaeltrestman.wordpress.com/cv-and-publications/

Trestman, M. (unpublished manuscript). *How swiftly does the stream of consciousness flow?*

Varela, F. J., Thompson, E. T., & Rosch, E. (1992). *The embodied mind: Cognitive science and human experience*. Cambridge, MA: MIT press.

Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford: Oxford University Press.

Zahavi, D. (1999). *Self-Awareness and alterity: A phenomenological investigation*. Evanston, IL: Northwestern University Press.

Zahavi, D. (2005). *Subjectivity and selfhood: Investigating the first-person perspective*. Cambridge, MA: MIT Press.

Downloaded by [Indiana University Libraries] at 14:47 30 September 2013