

THE RELATIONSHIP BETWEEN NEURAL CONTENT AND VISUAL

PHENOMENOLOGY By

MAUREEN CASSIDY TURNER

A Thesis Submitted to The Honors College

In Partial Fulfillment of the Bachelors Degree

With Honors in

Philosophy

THE UNIVERSITY OF ARIZONA

May 2010

Approved by:



Dr. Uriah Kriegel
Department of Philosophy

**The University of Arizona Electronic Theses and Dissertations
Reproduction and Distribution Rights Form**

Name: Turner, Maureen Cassidy	
Degree title: BA	
Honors area: Philosophy	
Date thesis submitted to Honors College: Title of Honors thesis: The relationship between neural content and visual phenomenology	
:The University of Arizona Library Release	<p>I hereby grant to the University of Arizona Library the nonexclusive worldwide right to reproduce and distribute my dissertation or thesis and abstract (herein, the "licensed materials"), in whole or in part, in any and all media of distribution and in any format in existence now or developed in the future. I represent and warrant to the University of Arizona that the licensed materials are my original work, that I am the sole owner of all rights in and to the licensed materials, and that none of the licensed materials infringe or violate the rights of others. I further represent that I have obtained all necessary rights to permit the University of Arizona Library to reproduce and distribute any nonpublic third party software necessary to access, display, run or print my dissertation or thesis. I acknowledge that University of Arizona Library may elect not to distribute my dissertation or thesis in digital format if, in its reasonable judgment, it believes all such rights have not been secured.</p> <p>Signed:  Date: 5 May 2010</p>

I would like to dedicate this thesis to my father, who, in many ways, has been the source of my intellectual curiosity.

Also, many thanks to Dr. Uriah Kriegel for all of the support and guidance you have given me during this process.

I. Introduction and overview (what I am going to say)

In this paper, I will attempt to account for the contents of visual phenomenology. I will suggest that the most useful and sufficient way to account for the information experienced during visual consciousness is to appeal to fundamental, sensory points of data that are organized spatiotemporally. That is, we should understand phenomenal experience as composed solely of tiny flecks of sensory information spread over the entire three-dimensional space of which we are aware. This species of information is both intuitively fundamental to our experience and appears to be actually fundamental neurally, so, it is an excellent candidate for a sufficient account of the contents of visual phenomenology.

For the present purpose, I will stray little from dealing solely with visual consciousness. However, it is my aim that the account should extend smoothly to each other sensory domain of phenomenology. I will argue that this kind of extension is enough to fully describe the contents of all experience. Visual consciousness is in no way separate from other aspects of conscious experience, except in that it has been most extensively studied scientifically. For that reason, I will focus on how my account functions with respect to visual phenomenology.

In sections II and III, I will begin by looking at what a theory of content is and then why we need a theory of content for visual phenomenology. In section IV I will examine the relationship between content and representation and also attempt to provide a set of criteria for selecting a theory of content. In section V I will examine the relationship of the audience to the theory. Section VI will provide an overview of the theory of content that I think follows most

naturally from the preceding considerations and standards. I will close with a survey of several objections in section VII.

II. Definitions (what I am going to say things about)

The brunt of what we are dealing with here is visual phenomenology. Visual phenomenology, simply put, is the experience of sightedness. Visual phenomenology is the experience, the *what-it's-like*, to have sight. It is the conscious array of information that has been fed to the visual stream from both internal and external sources. Beyond this, there is little we can say about visual phenomenology without committing ourselves to some theory of the content of visual phenomenology.

A theory of content, in the case of the visual modality, is a theory describing what sorts of things are experienced in phenomenology. We do not simply have experiences, rather, we have experiences of things – our experiences have meaning to us and we are able to derive a great deal of information from them. Theories of content are fundamentally linked to theories of phenomenology. This point can be understood most readily by observing that, if one is not having an experience of something, then one is not having an experience. I tend to think that this is a really important observation. In the course of this paper, I will suggest that theories of content are equal to, or more important than, theories of the phenomenon of experience proper to solving the problem of consciousness in a satisfying manner.

My particular theory of conscious content appeals to the sensory building blocks that make up experience. Throughout the paper, I will avoid referring to these as “sense data,” because my theory has little in common with classical sense-data theories. Instead, I will call

them “sensory values.” A sensory value is meant to be the smallest unit of information that is available to the human perceptual system in a given modality. Here again my account will seem limited, to the point of chauvinism, in that it deals only with the human perceptual system. However, the account can be altered, with few changes to fit the perceptual systems of other animals. The basic spirit of my sensory value theory is that we ought to appeal to units of content that are physically, anatomically basic. My theory is then lent strength by its intimate link to our ontology. Sensory values, in the case of the visual modality, are units essentially of light, which we experience as colors in various shades and luminosities. How these units function to comprise a seamless phenomenology will be the focus of the paper from here on out.

III. What's the problem?

Visual phenomenology is vast, rich and complicated. Beginning very early in life, each of our individual phenomenologies accompanies us throughout our entire waking (and to various degrees at other times) life. Phenomenology extends temporally across years, and manages to be reasonably continuous despite constant interruption in input. Spatially, our visual phenomenology spreads out from the center of our minds to the blurry edge of our visual input. In small part, the spatiotemporal volume of our visual phenomenology makes its content appear intuitively scientifically intractable.

Visual phenomenology poses (literally) a massively complex problem for philosophers and scientists. The heart of the issue lies in questions about the essential nature of phenomenology: how does the physical brain produce phenomenology? Why do we have a phenomenology instead of factual input without an experiential component, like we imagine

robots do? It is logical to expect that the first step towards answering these questions will be to establish an understanding of the limits of the phenomenon in question. Understanding what phenomenon in the world we are attempting to explicate is difficult in the case of phenomenology owing to the general *transparency of experience*ⁱ. When we begin to pick out a token instance of a certain conscious phenomenon, we invariably end up picking it out in terms of what the experience is of, that is, seemingly, we can only tell our conscious experiences apart insofar as they differ in content.

IV. Theories of content

Taking this in stride, what we will need if we would like to solve the problem of consciousness is a theory of phenomenal content. To properly assign the contents of a particular phenomenal experience is simply to sufficiently characterize that experience. Assigning content is a means of making a phenomenon tractable and poised for explanation by scientific means. What will be left to solve after we have established a solid theory of phenomenal content is a mere instantiation of this theory in the brain.

However, theories of content are in no short supply, there is a question of how we will know when we have established a correct theory and should begin to seek out the neural substrata. Understanding the close link between phenomenal content and phenomenology itself leads naturally to a certain set of criteria for a pragmatically acceptable theory of content.

Crafting a theory of phenomenal content is an essentially representational pursuit. When I say that something has a certain content, I am saying that it has a certain meaning to me, or to a broader audience in many cases. Assigning content is the act of picking out what information can

be gotten from a certain symbol. Representation has a symbol, a meaning (usually defined in terms of information content) and an audience for whom that symbol has meaning. The features of representational accounts at large will define the features of our accounts of conscious content.

We should consider that the search for a best theory of consciousness content is somewhat unlike the vast majority of scientific and a good part of other philosophical pursuits in that there is not likely to be theory that is 'correct' in terms of what the nomological actuality is, there is no natural content to a certain phenomenology. Because content is something that we are assigning, there is no fact of the matter about what the best theory of content is. Simply put, I believe this is true as a result of the arbitrary nature of our assignment of the conventional representation relation. In order to best bring out this point, we will examine cases of content attribution that lie outside of the domain of mental content. Language is a particularly strong example of this feature of conventional representation, because words are highly meaningful completely arbitrary symbols. When I comprehend a certain word as having a particular meaning, there is nothing special about the token of the word that causes it to have that meaning to me. I might just as well have picked a different word instead to have that very meaning. There is no necessary connection between what the physical or acoustic nature of a certain word is and its meaning to me. All conventional representation works in this way. Other symbols are arbitrary in the same way that words are. There are many sounds and colors that we encounter on a daily basis that we have learned to associate with various contents or meanings, and we make signs that represent actions to be taken in traffic. Road signs present an especially perspicuous example of how easy it is to miss the arbitrary nature of representation. It is very easy to think that road signs have meaning in virtue of their visual resemblance to the situation in the world

that they are about. Consider, though, that we might have learned that a different visual symbol on a sign, for instance if yield signs were marked instead with an exclamation point, meant the same thing. There is no change to the meaning of a symbol to be had by altering the symbol itself. This is the fundamental feature of representation.

Why is it that an information content can be represented by any representer? When we speak of the representational power of symbols, we must understand that conventional representation is a product of there being an analyzing audience, in a way that natural representation is not. Natural representations occur in the world as part of a causal chain. Symbols, on the other hand, work as a product of cognitive processes. A given symbol has the meaning that it does if, and only if, it is viewed by an audience that has learned the association between the symbol and its object or meaning. This is hugely important to bear in mind when we are setting out to make claims about phenomena that are the product of the physical brain, like phenomenology. A neuron can be said to represent a certain feature of the world because of the causal interaction between the neuron and the outside world. The actual, extent, mind-independent tie between the neuron and the world is merely a causal one. Some event occurs in the outside physical world this event is followed causally by a physical event in the eye that is followed by a series of physical events in the brain. However, the *content* of a phenomenal experience, to scientists, can be anything. Choosing content is merely a means of divvying up a phenomenon so that we can think about it. Some choices of content will be more meaningful and more useful (not necessarily both though!), but, technically, as long as our chosen content covaries in the appropriate way with the phenomenon in question, then it seems that we will have ‘correctly’ assigned content. Or at least not ‘wrongly.’ That the relationship between content and symbol is arbitrary has further importance. If I would like to represent the number three, for

instance, I can pick any symbol I would like and decide that it represents three. What I am trying to get at here is that we could assign whatever content we liked to a particular phenomenology. If I pick out a content, then there is always going to be a way in which I can understand the symbol (or in this case, the phenomenal experience) as representing it. This is all just to motivate the thought that we should not be looking for something in the experience itself to tell us what our theory of content *should* be. Rather, we need to define other parameters along which we can evaluate the fit of our theories of content.

There is no natural theory of neural content. Neural representation seems to be natural representation of the events occurring in the outside world, largely in virtue of the causal connection between outer events and neural events. However, I think that there is no relationship of natural representation between the subjective phenomenal scene and our theory of the content of that scene. Given that natural representation is non-arbitrary, we might more readily agree to a theory of content that approximates the natural representational relations that do exist. This consideration is a vote in favor of a theory of content that is closely tied to neural actuality.

Bearing that in mind, we can begin to consider what the criteria should be for a theory of visual phenomenal content. The relationship between content and symbol, between representation and what is represented, is an interesting one. The meaning of a symbol emerges only under analysis or when someone assigns the meaning, there is nothing in the symbol itself that causes it to have that certain meaning. Meaning is mind dependent. So, when we go about attempting to assign content to visual phenomenal experiences, we should remember that there is no correct answer. Phenomenology does not have any intrinsic content. Probably it is the case that there is not even one type of content that can truly be said to be better than another. We

simply need an account of content that carves phenomenal experience at its joints in sufficient detail.

We should consider what our theory of content must do in order to provide a sufficient, scientifically useful theory of phenomenology. In order to give a satisfactory account of the content of particular bit, say one point in time, of visual phenomenology we must be able to describe every aspect of what is experienced by the subject whose phenomenology it is at that time point. Our theory must pick out that slice of phenomenal life uniquely such that only it (and possibly states that are identical to it) are described by the information content that our theory produces. We need the content to completely describe *what-it's-like* for the conscious experiencing agent. Description of *what-it's-like* is both the most difficult and most important task for a theory of phenomenal content. It is easy to make a mistake and think that this means that a theory of content must give us, as scientists and philosophers, some kind of access to the subjective state of the agent. I think this is an unreasonable expectation. We merely need a theory of content that yields the subjective information content of the state they are experiencing.

There are several criteria that we might use to compare account that seems sufficient to us for a best fit over all. Given that the eventual goal of those who study consciousness is to come to an understanding of the phenomenon in question, we should look for an account of content that will make a meaningful contribution to the pursuit of an understanding of the whole phenomenon. There are a number of types of theories of content, each appealing to various different types of information and conceptual distinctions. Many among these, are able to account for most of the vast array of phenomenal events that we experience.

Some accounts of content are more neurally viable than others. The study of consciousness itself has become a jointly philosophical and scientific pursuit. In fact, we have

good reason to suspect that the final answers to questions about consciousness will be made in neuroscientific terms. When we are assigning content, we are defining a phenomenal event so that it is tractable for study. We are defining the borders of the phenomenon so that we are able to further study it by scientific means. The best account of neural content will define phenomenal events in a way that makes them highly amenable to neuroscientific explanation.

Consciousness content is an output. What I mean by this is not what the traditional modularist means when he says that consciousness is the finished product of a certain module or chain or closed unconscious module. That sort of ‘consciousness-as-the-final-product-of-the-brain’ view has always seemed a little bit puzzling to me because it seems dangerously susceptible to epiphenomenalist arguments. Rather, I think it is important that what we experience consciously is a finished product in most cases. While the input from the outside world is nothing if not ambiguous and the top down influences that the brain exerts on perception in light of goals and past experience may constantly conflict, what we get in visual consciousness is highly unified stable display. Saccade induced blindness is an excellent example of this. It has been known for some time that during eye movements people are essentially blind, that is, they are taking in no new information. However, even though each of us is constantly readjusting the direction of our eyes and making tiny micro-saccades, we never experience any kind of gaps in experience that ‘should’ correspond to the gaps in input. A variety of other observed effects and perceptual phenomena (including object persistence over time, sctomic filling-in, etc) point to the importance of continuity and unification in visual conscious content. But, continuity does not seem likely to be an end in and of itself in the case of visual consciousness – the purpose of consciousness, so to speak, is probably not to unify all percepts. Seemingly, that sort of unification could be accomplished without ever giving rise to experience. Instead, we should use

unity as a pointer. We know, given unity's huge role in conscious processing, that whatever the purpose of consciousness may be, it is likely that it necessarily requires unity of percepts. Also, we can say that theories of consciousness and conscious content that are able to account for unity and give it a place of proper importance should be preferred over those theories that do not. We need a theory of content that will allow for the totally seamless and massively rich nature of each phenomenal display.

The primary feature of visual phenomenology is that it is subjective. Subjectivity underlies the intuition that we cannot give an adequate, yet scientifically meaningful account of conscious content. Any account of content will have to explain why the conscious agent experiences that content subjectively, from the first person point of view, during a phenomenal event.

V. Who is the audience of a theory of content? (a little more on subjectivity)

At first glance it may seem trivial to ask who the audience of a theory of content ought to be. Phenomenal consciousness is a phenomenon that is totally subjective, so we often take it for granted that the audience of our theory should be the agent who is having the experience. This causes a lot of trouble for a theory of content. When we try to understand the content of an experience as it is experienced by the agent, we are not able to say anything about subjectivity because the content remains within the subjective realm of the agent. It is an attempt to sidestep subjectivity as a problem to be addressed. This does not yield a theory that allows for insightful objective analysis. Instead, we should attempt to define phenomenal content from the point of view of an outside agent. Then we have to face up to subjectivity and account for it in our theory.

The purpose of a theory of content is to define the information that is gained by the audience of the symbol. Here, we must consider the purpose with respect to the audience. The content of a word in a language is its meaning (likely some combination of sense and reference) to the listener. There is an intended content too, the meaning of the word to the speaker, but the point is that there is content in virtue of there being an interpreter of the symbol. The desires, knowledge, beliefs of the interpreter confine and define the content of the symbol in practical use. The same thing is going to be true for accounts of phenomenal content.

There are a few ways to understand the relationship between the content of visual phenomenology and the audience of the content. At first, it may be intuitive to say that the content of a given phenomenal scene should be related to the experiencer of the scene, the person who is having the subjective experience. With respect to the pursuit of an explanation of the phenomenon of visual consciousness, however, the audience who is assigning content is composed of those who are studying the phenomenon. I think that perhaps it is easy to confound the related, but separable types of phenomenal content. Importantly, when we are trying to ‘figure out’ consciousness, we cannot but address the issue from the second point of view. When we are attempting to assign content to a particular experience, we are not looking for the meaning of the experience to the experiencer. What we want is a complete characterization of the experience (which may necessarily need to take into account the subjective nature of the experience from the point of view of the experiencer). We want to be able to define the phenomenon in a way that is amenable to scientific explanations.

There is a difference between accounts that address content from the point of view of the observer and accounts that address content from the point of view of the conscious agent. I think this difference is important because it may have an impact on our overall theory of

consciousness. When we are examining the phenomenal array objectively, we are relatively unable to account for executive and attentional mechanisms that are at play. We are simply able to say what information is part of the entire phenomenal array. This kind of a view is not at all consistent with theories of consciousness that have very strict limitations on access. For instance, a theory that says that an agent is only properly conscious of the tiny slice of their experience that they are attending at any given moment (attention has been viewed as the gateway to consciousness) will not be consistent with the assertion that our entire phenomenal array composes our phenomenal content. So, the relationship between items in our phenomenology which we have access to and items that we do not have access to may limit how much we want to include in our theory of content. I happen to think that my phenomenal array includes broadly everything that I am experiencing at any given time point, rather than what I have some kind of access to. The theory that I will outline in what follows, though I will present it in its broadest incarnation, is amenable to adjustments for access and can account for executive functions insofar as they are realized phenomenally through actual perceptual changesⁱⁱ.

For example, the experiencer of a particular phenomenal scene is potentially able to extract a vast array of contents from the scene that may in fact never be extracted. Due to the influence of attentional mechanisms, knowledge about the world, past history and a host of other ‘top-down’ influences, different phenomenal arrays might be parsed differently. The most intuitive sort of example here is going to be the case where two viewers are looking at the same scene in the world. We can imagine that the viewers might have perfectly matched sensory organs, such that they have exactly the same input to their brains. Yet, we can imagine that, if we asked these viewers what they saw, they might describe different aspects of the scene in different ways. Do we want to say that these viewers have had different experiences? Yes, most likely we

do. But, do we want to say that the content of their experiences was different? No, I think we should not. They have simply made different verbal reports. Certainly, we should leave room for actual differences in phenomenology based on, for instance, attentional influence such as brightening and contrast changes. But we would not want to say that the actual information content was different. This is very tightly related to the idea that we have chosen for the audience of phenomenology, with respect to the study of consciousness is not the conscious agent, but the scientists who are studying the phenomenon.

Another example, we can imagine two viewers, who again have similar or identical sensory setups, who are looking down a dark alley at a dark shape. One of the viewers experiences the shape as a person crouched in the alleyway. The second viewer sees the dark shape as a shadow cast by a dumpster. Intuitively, the two viewers are seeing the same thing. They are having the same visual experience that is accompanied by different internal thought patterns. Importantly, from this example, we should take away that the same visual phenomenal scene can be parsed in different ways by the experiencer. This process of parsing, does not itself, change the visual phenomenal content. So, we need an account of visual phenomenology that allows one content to withstand breaking up according to multiple parsings. Notably, an object based account of phenomenology would have to claim that the two viewers have different phenomenal contents.

VI. Sensory Value Theory

We can start out by asking, broadly, of what our visual phenomenal experience is composed. Our conscious experience is obviously made up, at least in part, of sensory data.

When asked what we are experiencing, we invariably describe the sensory scene: objects and events that we have taken in and encoded from the outside world with our sensory organs. This seems unproblematically true. We also tend to describe our feelings, our subjective take on the outside objects and or internal emotional and somatic landscape. These events occur to us as spatially and temporally continuous. My experience of my internal state stops at my skin, which is right where my sensory experience of the outside world begins. In fact, the information being taken in via various modalities is all bound relatively seamlessly into one unified conscious representation. There are absolutely not gaps in my phenomenal experience. No matter which way I direct my sensory organs, both internally and externally, my brain organizes the information into a complete array by the time it reaches consciousness. It is this final, rich experiential representation that theories of consciousness broadly and theories of phenomenal content more specifically aim to address.

In our day-to-day life, we are accustomed to understanding that we process our external environment via our senses. No one will deny that, under normal circumstances, when I am having the phenomenal experience as of a flower pot sitting several feet in front of me, I am having that experience somehow in virtue of the fact that I am seeing the flower pot. That is, I am using my biological visual system to apprehend the flowerpot. If I see a flowerpot, it is because I have taken it in through the visual modality. Likewise, if I hear a bird, I have added that bird to my phenomenology through my auditory capacity. There is no reason to think that we need any other capacity to apprehend these entities. I would like to suggest that when we experience our internal life, we do so in a way that is very much like the way that we experience external objects. When I think to myself that something is lovely, I have literally heard myself speak internally. This is supported by studies demonstrating that the neural activation that goes on during thought is very

much like the activation that goes on during actual perception. Or, when I have a pain, I perceive it in a tactile manner in the same way that I might feel a rough surface with my fingertips. For this theory of phenomenal content, we must agree on a perceptual understanding of our internal life.

What we experience, in all modalities, can be completely characterized as sensory values. I understand that this is not in keeping with our immediate intuition. I will address alternate, intuitive views, for instance that our experience is as of objects, in section VII. My visual experience can be sufficiently characterized in terms of low level visual properties (like angle, color, line, etc). The subjectivity derives from my simultaneously representing internal sense values. My experience of my self is that I have certain feelings (tactile, somatic) and think certain things (auditory and visual inner activity). I know where the sensory experience of myself ends and the world begins because I have mapped my self spatially. The bounds of my inner experience are spatially where my skin stops and the outside world begins. Subjectivity is the result of the necessary centrality of my self in my field of perception. Also, I understand that my internal speech and imaginings are not available to those around me. The privacy of my inner life also contributes to the feeling of subjectivity. my account deals with subjectivity as an intrinsic feature of phenomenal content. Because sensory values are inherently spatial, they are closely tied to the experiencer. The experiencer has a self in that they are filling in sensory values for the space that is occupied by their own body, so subjectivity is able to be derived from the obvious spatial difference between each observer. As a consideration, this means that my account of phenomenal content is committed to the possibility of two people being in an identical phenomenal state. However, I think that this is not worth addressing as a possible objection, because it cannot actually occur (the timing could never be the same). The odds that two people

would have histories such that they would have the same ‘inner’ sensory life are also so low as to not be worth our time.

There is some room for confusion about the conversion from objectivity to subjectivity. We can imagine a number of circumstances where the experience that a person is having is actually a representation of some events that are occurring remotely to them. Consider the case where a person has on virtual reality goggles. Their brain will have represented the world as a spatial scene that, in fact, is nothing but a flat display. Our intuition here is to say that the subjective experience of the perceiver has nothing to do with the objective reality. However, this is not the sense in which I have used ‘objective.’ We are able to account for subjectivity in terms of objective spatial location because we experience our bodies as spatially located. Our bodily center, so to speak, is the center of our spatial experience. The experience is objective in the sense that there is nothing inherently personal about any of the information that is represented. The experience is subjective in that it radiates outward from the person’s spatial center, which makes the experience unique to them. There can be no other person at the same point at the same time.

Importantly, our experience is not literally objectively spatial from there. Instead, our experience is constructed out from our experience of our bodily center. This does not map onto the objective world. I am suggesting that we have a sort of mental space for constructing our experience, like a Cartesian system in the brain. This is not to suggest that there is something in that brain that we should be able to pick out that is a correlate of this construct. Our entire brain operates along spatial organization, and so, it is natural to say that the ‘whole’ representation would be comprehensible only as a sum total of a wide variety brain activity.

Given that the entirety of what we experience is sensory in nature, the most efficient and useful way of characterizing our phenomenology is via sensory information contents, what I will call sensory values. There is certainly nothing that we experience externally and (though perhaps it is less intuitive) there is nothing that we experience internally that cannot be sufficiently described in terms of sensory values.

Additionally, whatever content it is that we are representing in our experience, it is fundamentally spatial and temporal in nature. It is inconceivable that we might represent the *same* information that constitutes our experience in a non-spatiotemporal manner. Without spatiotemporal features, the information would be obviously different. It is not clear what it would be to represent spatial or temporal data non-spatiotemporally and still retain the full meaning of the information. For instance, say that the brain were to represent time in terms of numbers, like on a number line or an axis of a Cartesian coordinate system. When we say that the brain might represent some event A as having occurred just prior to event B as the numbers one and two respectively, it is not clear how we might distinguish the numbers in a way that does not appeal to numbers of things or at least some kind of spatial ordering of the numbers along a line. Without spatiality, numbers cease to be meaningful. So, we can say that the units that make up experience are necessarily represented spatiotemporally. To represent spatiotemporally just is to represent in a way that is experiential. There is no non-experiential understanding of spatiality. Likewise for temporality. I mean temporality in the sense that things must occur in an order. For each thing that is experienced there are things that were experienced just preceding and experiences that immediately follow. In no way do I mean to imply that the sensory values that compose conscious content are importantly ordered in certain histories. Rather, that they are temporal in nature. In this way, the content that we represent experientially cannot but be

phenomenal. Considering Blindsight allows us to deconstruct the elements that make up a full sensory value scene. A Blindsighted patient does not have access (or does not have at all) to the visual properties of their experiential scene. These patients are fascinating, because they are able to perform well above chance on a variety of spatial tasks. In the case of Blindsight, a person has constructed a spatial scene that does not have visual sensory data. They are able to function, but not well because their experiential layout lacks visual properties. All of the information that the brain is able to collect is built up into the same spatial experiential framework that normally sighted viewers use. However, the information is impoverished. They have experience as of the spatial and temporal aspects of the visual scene but lack visual sensory values.

As we begin to imagine representations of spatiality and temporality, the non-experiential accounts that we come up with are typically numerical. In order to make these numerical accounts have the same sense as our actual experience, however, we need to construct accounts that become more like the phenomenology that we are trying to explain and less like a separate explanation of it. For this reason, spatiotemporal info cannot be fully represented propositionally or ‘pseudopositionally’ as is often suggested when we consider why we have phenomenology at all.

The only understanding, aside from propositional representation, of spatial information that is not a direct spatial representational one itself seems to be some kind of numerical coordinate system, such as the common Cartesian one. This genre of representation is strictly mathematical. The existence of objects or the movement of entities through space can be understood via a series of mathematical equations with suitable parameters. To the trained eye (mathematicians and modelers), browsing through one of these models will allow the reader to

form in their own mind a spatial reconstruction of what is meant by the model. There are only a few ways that spatial information exists: as actual phenomenology or as mathematical models.

What is it to say that something is represented? Commonly put, it is to say that an agent understands a certain symbolic object to have a meaning that is beyond what meaning may be inherent to the object itself. In the case at hand, visual phenomenology, we (those who are thinking about the phenomenon) are the audience who understands neural activity to be an intricate set of symbols that represent phenomenology, among other things. We can parse phenomenology into units of sensory values occurring at certain spatiotemporal locations. Whether we are able to construct a 'one-to-one' understanding of how the sensory value recipe for a given phenomenological display (as it is chopped up into sense units) maps onto neural activity (as it is chopped up into single cell activity) at this point is not the question. What we need to see here is simply that there is no mystery about whether or not it is possible in theory for the sensory components of the phenomenology to be represented in a one-to-one or pseudo-one-to-one manner in the brain. By this I mean we know how brain cells work, we know that each cell or colony of cells or area of cells fires to encode certain properties for which it is best, though each unit may also fire to varying degrees for other properties. We also know that it is not quite this simple, among other things, units of cells change their firing over time, tune to other properties, change firing as a result of overall chemical changes and decay with age. Which is only to say that, while we may not be able to point to a part of the brain that represents each unit of sensory value required by my theory, we can understand that, taking into account dynamic factors like those just listed, we can tract representation of individual points of data in the brain with a reasonable expectation that we are not over or under simplifying what is represented. Maybe we cannot do this with perfect (or even near perfect) accuracy at this point in time. But I

think it is not a great stretch to say that the current program in neuroscience will lead us to be able to tease out, say, the yellow activity that is localized to the rectangle (that is my wall) in front of me at a certain location. Presumably we will also be able to pick out that this certain yellow occurred before other things after other things, that is, we will be able to indicate temporal mnemonic tagging in the brain. None of this really seems all that fantastical with respect to how far we have already come.

If we can say what information must be represented in order to explain consciousness (sensory values organized spatiotemporally) and we can say how that representation occurs (a normal dynamic neural model should do the trick nicely here), then we have a full, satisfactory account of the phenomenon. This would be totally surprising given the reputation that consciousness has for being utterly intractable by regular, scientific means of explanation. A short examination of the usual course of explanation is in order to make sure that I truly have done what needs to be done to offer satisfactory explanation. We can use the example of any physical phenomenon here. In the case of digestion, one can imagine that we (the scientific community) first noticed that something was going on: when a man puts food in their mouth and chews and swallows, not only does the person reap all manner of magical health and well-being benefits, but also something other than the exact food as it was initially inputted comes out as a result. The explanation that is required here is an examination of the phenomenon that is going on, unseen, that is producing all of these various outcomes. First, we notice and describe the circumstances surrounding a phenomenon. The description of the phenomenon is then parsed in terms of what we know about the physical body. Perhaps, during this step, further study of the human body may need to be executed. As findings about the physicality are made and mapped to our observations about the phenomenon, we may need to adjust our understanding of what is a

part of the phenomenon in question to have a more satisfactory account. When we feel that we are able to give a physical characterization of each tiny observed aspect of the initial phenomenon, we can say that we have fully explained it. The beauty of this is that, because we have understood what makes certain contents conscious in terms of what information is represented, we need only make sure that these contents are in line with what we know about neural activity and then we will have a tolerably complete, if not detailed account of how visual (and other modal) consciousness might arise.

The creation of the conscious intellect is merely a causal chain. This would only be different if, aside from this causal chain, there was also a self, a sort of viewer of the chain. However, I see no reason to posit a separate self from the chain of physical events in the brain that subserves consciousness and reason. The evidence for the self is strictly experiential. I know that I am a self because I can see my self and hear my thoughts. My thoughts are made up of various visual, auditory and tactile sensations that are all simulated together at the same spatial location inside my head. All these events are also mere physical processes. They require no viewer. They constitute the viewer.

Here we have picked out a phenomenon, namely visual consciousness, and used a theory of content to fully describe what the phenomenon is. It is important here that we must feel that our particular theory of content fully, sufficiently, completely describes the phenomenon in question, otherwise it will do no explanatory good to look for the physical instantiation of our definition of the phenomenon if we have picked out a description that does not fit what it is that we want explained. The easy part, then, on my account is coming up with a neural account that is true and fits with the description. I do not think that we have already completed what we need in the agenda of modern neuroscience to give a full exact account of my description of visual

consciousness, only that we are undeniably moving in that direction. And that there is nothing about the account that seems intuitively at odds with the type of neural explanations that we are already giving. This alone is enough to make my account stand out. There is nothing fancy about it or different than the various other phenomena that we have attempted to explain or have explained in neural terms. Consciousness, couched as I have, is just another phenomenon that can be explained by basic brain activity.

On this account, phenomenology, whatever it is that is the hard problem of consciousness, does not in any way emerge as a feature of the manner of representation. It cannot be the case that some information in the brain is represented in one way and so that information is unconsciously had, but some other information is represented in a special conscious way, and so that information is associated with or subserves conscious experience. Further, there are only a few elements in play that are potential candidates that might constitute an explanation of consciousness: the neural actuality of the brain, the outside world, experience itself, or the things that are experienced (the content of experience). Scientists will look for an explanation that exists in the brain alone without looking at the other aspects. They might ask what sorts of cells, when placed together correlate with conscious experience. This kind of an explanation can never lead to the intuitive satisfaction that a causal account can, because it is merely correlational. Those seeking a more mystical explanation might explain conscious experience in terms of the experience itself. These sorts of accounts deal with phenomenology in an essentially dualist way. Consciousness is some kind of element that exists separately from our analysis of the physical world. On these accounts, we cannot explain consciousness because our explanations can only be made in physical, scientific terms that cannot address this separate realm. Many accounts (here I am thinking of sort of theories of embodied cognition) try to explain consciousness in terms of

the interaction between the brain and the outside world. Here we face the same problems that strictly brain oriented accounts face – they are correlational and not satisfactory. There is a common intuition that philosophical zombies are possible. However, I think it is possible to make an account of conscious processing on which it is the nature of the representational content itself, not the mechanism or type of representation and not the nature of consciousness, that causes experience to arise.

I have proposed that we understand the content of visual phenomenology, and experience more broadly, in terms of sense-level units partially because it affords a high degree of unity of the percept. One could argue that an understanding at the object level might also function nicely in terms of unity. After all, each object concept encompasses a multi-modal understanding of that object. When we activate an object concept, the information activated is plausibly very much like the sort of information, with respect to content, that one represents when one experiences that same object. However, the apparent benefit of utilizing object concepts here boils down to the fact that object concepts are built up out of sensory information. I think it is going to be more parsimonious if we simply appeal to sensory values in the first place. Because there is no occurrence of our perception of a sense datum that is not bounded by spatiotemporal constraints, the incoming sensory values that we receive must necessarily be bound during the representation of the data. If I experience a sound as coming from the space that I also experience some color coming from, then the two are bound. Additionally, this account appealing to sensory values

This account is also readily amenable for translation into neural terms for explanation. Because the basic units of the account are sensory, it is easy for us to conceive of how and where they might be represented in the brain. In fact, an account localizing the contents of consciousness to low level visual areas is closely in keeping with literature showing that

phenomenology is impaired if low level visual areas are selectively lesionedⁱⁱⁱ. At any rate, I think that accounts appealing to basic sensory level units will have a much easier time being translated into neural terms than many other accounts of phenomenological content.

VII. Criticisms

Perhaps the major problem that people will have with my theory is one of intuition. When we think of our conscious experience, particularly our visual conscious experience, we think that we are conscious of objects. There is a confound here however, that (necessarily ?) permeates our thinking. When asked what it is that we see, we are likely to respond that we see a particular array of objects with some certain identifying features. We do not, obviously, report that we see certain chunks of sensory data aligned at certain spatial locations, flowing through various changes as time passes. What we have the experience of is objects. So why, it may fairly be asked, would we choose to make sensory level units the primary units of analysis, the units of content, of visual consciousness? In fact, the argument that I have attempted to motivate above in favor of the use of sensor units of conscious content is based, to a large extent on intuition. Unless there is something inherently better about sensory data, perhaps there is something privileged about this type of account in virtue of the nature of our cognition that might make it more explanatorily satisfactory (but I doubt it!), the real reasons to prefer this account are essentially that it makes sense to us to try to understand things in this way.

When we open our eyes, we have the experience of seeing various objects, arrayed together into a coherent world scene. We can, and do, carve this scene up in terms of object classes with very little remainder. In fact, this level of carving often seems to be the most useful

way to us to parse our experience of the world, because it is the level of division that allows us to act in an object-directed way. In so far as cognition occurs for the purpose of active output (pursuit of actions in keeping with evolutionary goals or whatever), understanding our experience as composed of objects is a top contender for 'best' understanding of the world. So, in many ways, it is unsurprising that object-level analyses are a common candidate for explications of phenomenal consciousness. However, while object classes provide a near-seamless, near-sufficient description of the sorts of conscious experiences we have, they are not perfect. For instance, we do not have an object concept for many of the emotional feelings that we experience as completely and inseparably integrated with the rest of the denizens of our consciousness. Even the feelings that we do have name and concepts for are subtly different each time we feel them. It is inadequate as a description of my phenomenology to appeal simply to my feeling happy. For each experience that I might class as happiness is vastly and importantly different from all the others. Using sensory values, on the other hand allows us to pick out our feelings in terms of how they actually physically feel to have. This sort of an embodied account of the phenomenal experience of our feelings is in keeping with trends in current cognitive neuroscience.

I think that the major temptation to resort to analysis at the object level results from our belief that the brain deals with objects. If I experience objects, and I am merely a product of my neurobiological processes, then it must be that my brain is somehow leading up to these objects. We think of an object-based understanding as the goal of the brain, an untouchable modular top of things. Not only is our description of conscious content in terms of object conceptual classes an inadequate description of our phenomenology, but it is not necessarily the case that the brain is object driven in the popular sense.

Expunging our phenomenal content of objects leaves us with an interesting intuitive gap. There are cases where people who have the same input and have the same causal chain leading up to experience none-the-less have quite different final experiences in terms of what objects they experience. Consider the case of two people, one of whom is neurologically normal and the other of whom has prosopagnosia. These two people are looking at a face. The normally sighted viewer experiences a face; the prosopagnosic, swirling colors. We want to say that their experience has different contents – one person can see a face and the other cannot. A theory appealing to sensory values must say that the two viewers have the same experiential content. I think we should deal with cases like this by understanding that the two viewers can have the same content, but derive different information from their experience. For the scientific audience, the two viewers have the same content. If we were taking the perceiver as the audience (I have given several reasons in section V why we should not), then we would have to account for the different information that the two derive.

The content of experience is non-conceptual. Our understanding of the world as made up of various classes of objects and types of things emerges only under analysis. I might have the same visual phenomenal experience whether I have decided that I am seeing a duck or a rabbit. My experiential perception is the same either way. What is different is my perception of what is going on mentally. Perhaps in the case that I am seeing the duck, I say internally to myself ‘duck.’ I literally hear my internal voice saying this. There is ample experimental evidence demonstrating that our internal mental life is played out via the same sensory regions that subserve outer perception.

There is another problem. Representation isn’t really a real thing. That is, representation is just a way of describing certain of the causal processes in the world. This account says that

something in the outside world influences my retinas influences my LGN influences my brain, neuron by neuron. End of story. Representation is just a description of this. So, it seems that it should not matter what content is represented. The account makes clear that it cannot be in the manner of representation that the content becomes conscious.

A causal influence account leaves no room for consciousness. Just because a single neuron is influenced by both spatial and temporal and sensory factors (presumably this is what the having of 'conscious' content is on this picture) does not have any a priori link to its content being conscious. I suspect it does not have any link at all. Perhaps it is that case that a causal link to each of these types of information in a neuron that can be said to give rise to consciousness is necessary for it to give rise to consciousness but not sufficient.

More generally, we can ask: what is representation? Representation on this account is a purely causal chain in the physical world. It becomes a representational process under analysis by a conscious being. A person decides that there is some kind of link between a neural action and something in the outside world that is far out there in a causal chain linking the world event and the neural event. We call this representation. It seems like there are a billion variations on this: sometimes neurons are thought to represent what other neurons are doing, sometimes whatever is out there, all kinds of stuff. Representation happens when a person looks at the neural event and it means something to them. There has to be a viewer.

Perhaps there a problem with all accounts of consciousness that appeal to representation. I mean, we've got accounts like self representationalism or first-order stuff, or all kind of other accounts. But there always has to be a viewer, to make what is represented conscious. The viewer is really special in representation as I mentioned in the preceding. It does not seem to me that having the info be accessible to other parts of the brain or poised in such a way (so as to

make the rest of the brain the 'viewer') does enough. To have representation, it has to have meaning to the viewer. Not just any kind of viewer is capable of finding meaning. I think probably only an intelligent conscious one is.

ⁱ Harman, G. (1990). The Intrinsic Quality of Experience. In *Philosophical Perspectives 4*, edited by J. Tomberlin: Ridgeview Publishing Co.

ⁱⁱ Which is quite far actually. See the work of Dr. Marisa Carasco et al.

ⁱⁱⁱ There is an extensive literature about the phenomenon, Blindsight.