

Copyright  
by  
Ting Fung Ho  
2022

**The Dissertation Committee for Ting Fung Ho Certifies that this is the approved  
version of the following Dissertation:**

**Kinds of Awareness**

**Committee:**

Michelle Montague, Co-Supervisor

David Sosa, Co-Supervisor

Mark Sainsbury

Susanna Schellenberg

Galen Strawson

Michael Tye

**Kinds of Awareness**

**by**

**Ting Fung Ho**

**Dissertation**

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

**Doctor of Philosophy**

**The University of Texas at Austin**

**August 2022**

## **Dedication**

To my mother

## **Acknowledgements**

I would have to, first and foremost, thank my dissertation committee for their support over the years. Collectively, they have helped me learn to think and write with rigor without sacrificing originality. Individually, I would like to thank Michelle Montague for the encouragement and reading my work with a careful eye, Mark Sainsbury for the consistently high-quality feedback, Susanna Schellenberg for helping me find my voice in philosophy of perception, David Sosa for all the insights, Galen Strawson for showing me how deep philosophy can be, and Michael Tye for the invaluable and critical input. I have also benefitted from the very helpful written feedback on this and related projects from Ray Buchanan, Jonathan Dancy, Sinan Dogramaci, Daniel Drucker, Katherine Dunlop, Gary Hatfield, Kathleen Higgins, Cory Juhl, Rob Koons, Robbie Kubala, Errol Lord, Lisa Miracchi, Ian Proops, Bronwyn Stippa, and Matthew Vermaire.

Finally, a huge part of this dissertation is written amidst personal, social and global turmoil. I am grateful to my partner, Wenjin Liu, for being with me along the way.

## **Abstract**

### **Kinds of Awareness**

Ting Fung Ho, PhD

The University of Texas at Austin, 2022

Supervisor: Michelle Montague and David Sosa

There are various ways to be aware of things around us. When I gaze at a tree as an anchor point of meditation, the tree appears differently from the way it typically does. While typical perceptual awareness presents objects as external to the subject, meditative awareness – under certain Buddhist traditions at least – is marked by a diminished sense of boundaries between the subject and the external world. This dissertation is to understand how different forms of awareness are similar to one another. More specifically, it investigates the phenomenal, metaphysical and neural differences among three kinds of awareness: perceptual awareness, meditative awareness, and aesthetic awareness.

A unified picture that will emerge in this dissertation is that differences among these three kinds of awareness are fundamentally differences in *salience distribution*. These kinds of awareness share a basic structure that consists of (1) the subject element, (2) the subjective aspect of experience element, and (3) the object presentation element. Each element can vary in degrees of salience. Typical perceptual awareness is characterized by

the domination of the object presentation element; meditative awareness is characterized by the domination of the subjective aspect of experience element; and aesthetic awareness is characterized by an increased salience of the subject element.

The first three chapters each focus on a kind of awareness. The first chapter develops a new interpretation of perceptual presentation. The second chapter takes insight from Vipassana meditation to revamp the idea of introspection as an inward-looking activity. The third chapter investigates the role of perceptual surprise in aesthetic appreciation.

The last two chapters address the question: How does the structure of awareness help explain the value of awareness? These chapters mark the beginning of my exploration of this question, focusing on perceptual awareness. Chapter four argues that outward-directed attention, which is characteristic of perceptual awareness, is essential for sophisticated forms of agency that involves complex action-decision planning processes. Chapter five addresses the phenomenon that perceivers are apt to receive moral and epistemic blame for attending in inappropriate ways, arguing that inappropriate attention is blameworthy because it reflects the vicious character of the subject.

## Table of Contents

Chapter 0. Introduction .....	10
Chapter 1. What is Perceptual Presentation? A Lesson from Metaphysics of Action.....	13
0. Introduction.....	14
1. Hornsby on Metaphysics of Trying and Action.....	18
2. New Development of the World-Presenting Thesis .....	22
3. Conclusion .....	30
Chapter 2. Rediscovering Introspection Through Buddhist Meditation.....	34
0. Introduction.....	35
1. How Introspection is Lost.....	40
2. Rediscovering Introspection Through Vipassana Meditation .....	47
3. Looking Inward Without Inner Sense.....	71
Chapter 3. Perceptual Surprise and Aesthetic Appreciation .....	74
0. Introduction.....	75
1. What is Perceptual Surprise? .....	78
2. Perceptual Surprise in Aesthetic Appreciation .....	87
Chapter 4: Attention, Perceptual Judgment, and Perceptual Science .....	103
0. Introduction.....	104
1. What is Perceptual Judgment?.....	105
2. Conscious Attention and Perceptual Judgment .....	108
3. Implications for the Debate about Unconscious Perception .....	114
4. Conclusion .....	118



Chapter 5: Looking, Responsibility, and Rationality.....	122
0. Introduction.....	123
1. Looking, Seeing, and Perceptual Experience .....	124
2. Moral Responsibility of Looking.....	127
3. Explaining the Normative Significance in Terms of Looking.....	128
References.....	137

## Chapter 0. Introduction

There are various ways to be aware of things around us. For instance, when I gaze at a tree as an anchor point of meditation, the tree appears differently from the way it typically does. While typical perceptual awareness presents objects as external to the subject, meditative awareness – under certain Buddhist traditions at least – is marked by a diminished sense of boundaries between the subject and the external world. This dissertation is to understand how different forms of awareness are similar to one another. More specifically, it investigates the phenomenal, metaphysical and neural differences among three kinds of awareness: perceptual awareness, meditative awareness, and aesthetic awareness.

A unified picture that will emerge in this dissertation is that differences among these three kinds of awareness are fundamentally differences in *salience distribution*. These kinds of awareness share a basic structure that consists of (1) the subject element, (2) the subjective aspect of experience element, and (3) the object presentation element. Each element can vary in degrees of salience. Typical perceptual awareness is characterized by the domination of the object presentation element; meditative awareness is characterized by the domination of the subjective aspect of experience element; and aesthetic awareness is characterized by an increased salience of the subject element.

The first three chapters each focus on a kind of awareness. The first chapter, “What is Perceptual Presentation? A Lesson from Metaphysics of Action,” develops a new interpretation of what it means for perceptual awareness to present external objects. In the

literature, the idea of perceptual presentation is often cashed out in metaphorical terms (e.g., “openness” to the world). Drawing on Jennifer Hornsby’s metaphysics of action, this chapter develops a non-metaphorical account of perceptual presentation. The second chapter, “Rediscovering Introspection Through Buddhist Meditation,” takes insight from Vipassana meditation to revamp the idea of introspection as an inward-looking activity. It argues that Vipassana meditation involves a *sui generis* form of attention, that I call “anti-Gestalt attention,” and that this form of attention offers an apt description of what it means to “look inward” when one introspects. The third chapter, “Perceptual Surprise and Aesthetic Appreciation,” explores the role of perceptual surprise in aesthetic appreciation. It develops a new account of perceptual surprise, according to which perceptual surprise arises from perceiving features that are contra-standard relative to the categories that the subject attribute to a determined object and its features. On the basis of this account, it argues that perceptual surprise contributes to aesthetic appreciation by motivating the subject to discover more contra-standard features or more ways in which the features already observed are contra-standard.

The last two chapters address the question: How does the structure of awareness help explain the value of awareness? These chapters mark the beginning of my exploration of this question, focusing on perceptual awareness. I argue that the object presentation element is accentuated because perceptual awareness is for successful interactions with the external world. This view is partly established in chapter four, “Attention, Perceptual Judgment, and Perceptual Science”, where I argue that outward-directed attention, which is characteristic of perceptual awareness, is essential for sophisticated forms of agency

that involves complex action-decision planning processes. In chapter five, “Looking, Responsibility, and Rationality”, I address the phenomenon that perceivers are apt to receive moral and epistemic blame for attending in inappropriate ways. I argue that inappropriate attention is blameworthy because it reflects the vicious character of the subject.

## Chapter 1. What is Perceptual Presentation? A Lesson from Metaphysics of Action

*Abstract.* It is natural to think that perceptual experiences are direct *perceptual presentations* of external items. However, the nature of perceptual presentation is poorly understood. Oftentimes, perceptual presentation is characterized either in metaphorical terms (e.g., “openness” to the external world, “acquaintance” with external objects) or in negative terms (i.e., not via the mediation of sense-data). More recently, Michael Martin (2002, 2004) and John Campbell (2002, 2010), among others, develop a positive, non-metaphorical way to characterize perceptual presentation. Using the notion of *parthood*, they maintain that the external items perceived are *part* of veridical perceptual experience. But this way of characterizing perceptual presentation is not only counter-intuitive, but also inconsistent with contemporary perceptual science (Burge 2005; Pautz 2020; Beck 2021). In this paper, I develop a new way to characterize perceptual presentation. Instead of *parthood*, I make use of a metaphysical relation which I call *essentiality for a kind*. Jennifer Hornsby (1980, 1997) uses this metaphysical relation to theorize about action. I argue that applying the metaphysical structure of Hornsby’s account to perception leads to a better interpretation of perceptual presentation: (1) it easily makes sense of perfect hallucination; (2) it dissolves the longstanding dichotomy between common-kinds views and disjunctivism; and (3) it is compatible with contemporary perceptual science, which usually takes perceptual experience to be internal representation.

## 0. Introduction

Consider Jane who is looking at an apple in front of her. Jane’s perceptual experience *presents an apple to her*.<sup>1</sup> It is a way for her to engage with an apple – an item in the external world – a way that is [1] *primordial* and, more importantly, [2] *immediate*. It is primordial in the sense that it underlies other more sophisticated ways to engage with the apple. On the basis of her experience, Jane can, for instance, pick up and eat the apple, or mentally point at it and form conscious thoughts like “*There it is, an apple!*”.<sup>2</sup> It is immediate in the sense that it makes the apple available in such a way *immediately* – without the need of Jane being first consciously aware of items that are not in the external world (e.g., mental items like a throng of sensations), or forming conscious thoughts like “my experience is appropriately caused by the apple”.<sup>3</sup> Throughout this paper, I refer to the thesis that conscious visual perception has such a world-presenting nature as the *world-presenting thesis*.

---

<sup>1</sup> I use “perceptual experience” sometimes as a countable noun, and sometimes as an uncountable noun. When used as a countable noun, perceptual experience refers to a specific conscious perceptual episode, whereas a conscious perceptual episode can be individuated in terms of the worldly items that the experience is about. For instance, a perceptual experience (as) of an apple is a different conscious perceptual episode from a perceptual experience (as) of a cat. When used as an uncountable noun, “perceptual experience” refers to conscious perceptual state as a kind of mental state.

<sup>2</sup> There is the further question of what sophisticated ways to engage with an apple is enabled by the perceptual experience. There is no definite answer; it depends on factors such as what capacities the subject in question has. For instance, a mouse might as well have a perceptual experience that presents an apple to it. But arguably the presentation would not enable a mouse to form demonstrative thoughts about the apple, because it lacks the conceptual capacity to do so. Also, I do not mean to say that perceptual experience is the *only* primordial way to engage with worldly items. Bumping into something might as well be a primordial way to engage with that thing.

<sup>3</sup> See also Johnston (2019: 152).

The world-presenting thesis is widely accepted in recent theorizing about perceptual experience.<sup>4</sup> Different aspects of it have been employed as important desiderata for theorizing about perceptual experience.<sup>5</sup> The thesis also has strong intuitive appeal. When we reflect on what it is like for us – normal human beings – to visually perceive, phenomenologically it seems that we are presented with worldly items in the sense described. When Jane consciously sees an apple, it doesn't feel as though she has to go through any extra mental activity or non-worldly entity in order to be aware of the apple. Instead, the apple seems available to her *right away*. This phenomenological feature of perceptual experience is widely noted in the literature. For instance, P. F. Strawson famously tells us that “mature sensible experience (in general) presents itself as, in Kantian phrase, an immediate consciousness of the existence of things outside us” (1979: 97).<sup>6</sup>

In this paper, I am going to take the world-presenting thesis seriously, and offer a new way to develop it. To do so, I draw on the metaphysics of trying and action that Jennifer Hornsby developed in *Actions* (1980) and *Simple Mindedness* (1997). Hornsby famously claimed that when a subject tries to perform an action, and when her attempt is a success,

---

<sup>4</sup> But note that the world-presenting thesis is – like any other philosophical thesis – not universally accepted in the history of philosophy. An infamous theory that conflicts with the world-presenting thesis is the *sense-datum theory*. In this paper, my primary goal is not to defend the world-presenting thesis. What I'm interested in is rather: *Given* the world-presenting thesis, what should the ontological structure of perceptual experience be like? So, I'm going to set such theories aside.

<sup>5</sup> For recent representative works, see Brewer (2011), Campbell (2002, 2010), Fish (2009), Johnston (2019), Martin (2002), McDowell (2013), Schellenberg (2018), Searle (2015), Travis (2013), and Tye (2014, 2015).

<sup>6</sup> See also Crane and French (2015) and Johnston (2009: 142). More recently, Johnston (2019) calls it *the phenomenology of the presence of the external* and uses it as a support for Disclosive Realism, his variation of the world-presenting thesis.

“the attempt is the same as the action” (1997: 94-5). This claim can be confusing. An action depends for its existence on (some of) its effects. For instance, the action of raising one’s arm doesn’t take place without one’s arm’s going up. But one can try to raise one’s arm without one’s arm’s going up. Alternatively put, the attempt is fallible in a way that the action is not. How can they be the same? As I will explain, Hornsby makes use of the distinction between *parthood* and *essentiality for a kind*. The attempt is the same as the action in the sense that they are composed of the same event, thereby having identical parts. The effect of the action, such as one’s arm’s going up, is not part of that event. Yet, it is essential for the event’s belonging to the kind *action*.

I apply this distinction to the perceptual case, by making *perceptual experience* an analog of *trying/attempt*, and *conscious perception* an analog of *action*. I present a view according to which the worldly items presented are essential for an experience’s belonging to the kind *conscious perception*, but they are not part of the experience itself. My interpretation is different from a common interpretation of the world-presenting thesis that has drawn wide attention in the recent literature, usually known as *relationalism*.<sup>7</sup> Relationalism maintains that the internal occurrence within the subject, the relation to the worldly items, and the worldly items themselves are all *parts* of the experience.<sup>8</sup> This view, which literally takes perceptual consciousness to be out there in the external world, could be found radical. Relationalism also has a number of serious

---

<sup>7</sup> Representative works include Brewer (2011), Campbell (2002, 2010), Fish (2009), and Martin (2002).

<sup>8</sup> Searle (2015) describes the view as follows: “A favorite saying of theirs is that in the veridical case the object is *literally* part of the perceptual experience” (174).



problems: it entails an untenable version of disjunctivism (Johnston 2019); it fails to offer a plausible explanation of perfect hallucinatory experience (Searle 2015 Ch. 1 & 6); and it is incompatible with the science (Burge 2005; Pautz 2021). In contrast, the metaphysics I offer provides a modest interpretation of the world-presenting thesis that has none of these problems.

Two caveats before I begin. First, by making an analogy with action theory, I am not saying that perceptual experience is trying of any sort, nor that conscious perception is action of any sort. What I'm getting out of Hornsby's theory is at a very high level of abstraction. I'm abstracting the metaphysical relation between trying and action out of Hornsby's account, and applying it to the perceptual case.<sup>9</sup> Second, my discussion will focus exclusively on vision. I suspect that many of the things that I am about to say will hold for other modalities. But whether it does is not a concern of this paper.

Here is the plan. In section 1, I offer a critical discussion of Hornsby's account of trying and action, carefully laying out the metaphysical structure involved. In section 2, I apply the metaphysical structure to the perceptual case. I argue that this new development of the world-presenting thesis is better than the common interpretations. In section 3, I explore the broader implications of my view on philosophy of perception. I argue that my view dissolves the longstanding dichotomy between disjunctivism and common-kind

---

<sup>9</sup> For an account that takes conscious perception to be some sort of action, see Noë (2006).

views, and that it revitalizes internal representationalism against the charge of failing to respect the world-presenting thesis.

### **1. Hornsby on Metaphysics of Trying and Action**

Consider actions such as raising one's arm. Accounts of metaphysics of action can be distinguished in terms of the theoretical choices about how the following three kinds of event are related: *the attempt/trying* to raise one's arm, *the action* of raising one's arm, and *the resultant state* of one's arm's rising or going up. Before getting into Hornsby's specific choices that mark her account, let me make some general remarks about these three kinds.

The attempt is a fallible kind: When I try to raise my arm, my attempt can either be successful or unsuccessful. The success condition is determined by whether the resultant state of the arm's going up obtains. Trying is a common kind shared by both successful and unsuccessful attempts: I can try to raise my arm even though my arm fails to go up (e.g., when I am unknowingly tied up). That trying is a common kind does not imply that it is neutral about whether it is successful or not. Attempts are geared toward success: we can conceptualize trying to raise one's arm as the kind of event that results in one's arm's rising in the good case. Unlike trying, the action is a success kind. I don't raise my arm unless my arm indeed goes up. In other words, the action entails the resultant state.

When it comes to the metaphysics of action, it might be natural to think that the action is composed of the resultant state: that the action of raising my arm is composed of my arm's going up. This view could be appealing because intuitively bodily actions like

raising one's arm are observable. It seems rather natural to think that in seeing my arm's going up, you are thereby observing my action of raising my arm. Nevertheless, this is not the view that Hornsby opts for.

Hornsby maintains that the resultant state is not even part of the action. She takes the resultant state to be a distinct event that bears a causal relation to the action: the action of raising my arm causes my arm's going up. She writes, "I take the view that the relation between an action of someone's moving her body and her body's moving is a causal one and that a person's body moving is no part of her action of moving her body" (1997:94).

Still, Hornsby does not deny that the resultant state has a special role for what the action is. In the case of bodily actions such as raising one's arm, bearing a causal relation to the bodily movement is essential for an event to be the action. It is not a contingent fact, but a matter of necessity, that when I perform the action of raising my arm, my arm goes up. In other words, the relation between the action and the bodily movement is unlike that between the wind's blowing and the leaves' falling. The wind's blowing causes the leaves' falling, but not necessarily so; it is possible that the wind blows but no leaf falls. On the other hand, bearing a causal relation to the resultant state is essential for an event to be the action: no action of raising the arm without causing the arm's rising.

In sum, the relation between the action and the resultant state can be understood as an application of (what I call) the Principle of Metaphysical Dependence without Parthood:

For entities  $x$  and  $z$ , and a kind  $y$ , it may be that  $x$  is essential to  $z$ 's belonging to  $y$ , even though  $x$  is external to (/not part of/separate from)  $z$ .

To apply this principle to the action of raising one's arm, let  $x$  be the event of one's arm's going up,  $z$  be the placeholder of the event that composes the action of raising one's arm (I'll explain what the event is in a moment.), and  $y$  be the kind *the action of raising one's arm*. With the principle, one's arm's going up is not part of the event that composes the action of raising one's arm, but it is essential for that event's belonging to the kind *the action of raising one's arm*.

This principle can also be found in Davidson (1987). Davidson applies it to the nature of sunburn. The sunburn is an injury of the skin. The existence of the sun, and the causal relation between the injury and the sun are essential for a skin injury's being *sunburn*. But neither the sun, nor the causal relation to the sun, is part of the injury itself: If I have a sunburn, and I'm asked to locate where my injury is, I point at a particular area of my skin. Pointing at the sky would be a mistake. The fact that it is a mistake to locate the injury where the sun is shows that the sun is not part of the injury, even though the sun is pertinent to the nature of sunburn. A different way to establish the same point: It is possible to have a tissue damage that is an exact duplicate of a sunburn but is not a sunburn, when the skin damage is caused by an ultraviolet tanning machine for instance. The injury is compositionally indistinguishable from a sunburn, but it is not a sunburn.<sup>10</sup>

---

<sup>10</sup> A structural difference between sunburn and the action of raising the arm might be worth mentioning. While the non-composing essential event for the action of raising the arm is what the action *causes* (i.e., the arm's rising), the non-composing essential event for sunburn is what causes the sunburn (i.e. the UV rays from the sun). This difference is not important for this paper. As we will see in the next section, conscious perception is more like sunburn in this sense.

So, what exactly is the event that composes the action then? Hornsby takes it to be the event of trying. According to Hornsby, trying is an inner event that lies within the body. In the case of trying to raise the arm, we can understand the trying as a mental effort that will lead to the arm's going up if the world cooperates (e.g., that the body is well functioning, that there's nothing outside locking up the arm). Taking the action and the attempt to be identical in composition leads to Hornsby's famous claim that actions lie within the body.

In other words, it is not the case that I *first* try to raise my arm, and then raise my arm. Trying to raise my arm and the action of raising my arm are two different kinds of event to which the same event belongs. To motivate her view, one of the cases that Hornsby (1980:39) asks us to consider is someone who suffers intermittently from partial paralysis. When she tries to raise her arm, her arm may or may not go up. In the scenario where the patient's attempt does result in her arm's going up, there is nothing extra for her to do in comparison with the scenario where the trying fails to lead to the arm's rising. What accounts for the difference between two scenarios is just that the world cooperates in the successful one, but not in the failing one. Since "no further doing is called for from the agent" (1980:38) when her attempt successfully causes her arm's rising, her action and her attempt are identical in composition.

To put the same point in a different way, the action of moving the body is a sub-kind of the attempt to move the body; it is the successful attempt – *the attempt that results in the bodily movement*. While the attempt and the causal relation to the bodily movement

are both essential for the action, the action metaphysically depends on them in different ways. The action isn't the trying plus the causal relation to the bodily movement in the sense that a knife is a blade plus a handle, or serving dinner is serving the main dish plus all the other courses. A blade and a handle are equally *part* of a knife; and serving each course is *part* of serving dinner. But (the causal relation to) the bodily movement is not part of the action. The trying and (the causal relation to) the bodily movement contribute to the nature of the action in different ways: the former being its composition, and the latter being essential for but not part of it.

Let's summarize Hornsby's metaphysics of trying and action, illustrated with arm-raising. The attempt to raise one's arm is an inner event whose token composes the action of raising one's arm; the action has no additional part that the attempt does not have. But a token of the inner event can fail to be the action; one can *merely* try to raise one's arm, with no success. To be the action of raising one's arm, the token has to cause the rising of one's arm. The rising of one's arm is essential for, but not part of, the action.

Hornsby's account, which takes actions to be composed of inner, non-observable events, can be counterintuitive. This problem is known as the *invisible action objection*. While Hornsby in her later work (1997) tries to offer a defense of her view against this objection, how it pans out is not a concern of my paper. For what matters here, I am going to show that adapting the metaphysics above for the perceptual case results in an attractive interpretation of the world-presenting thesis.

## **2. New Development of the World-Presenting Thesis**

In the metaphysics of action there is controversy over the relation between the action and the resultant state. Likewise, in the metaphysics of perception there is controversy over the relation between veridical perceptual experience and the external items perceived. Relationalism occupies one position in conceptual space, holding that the external items perceived are part of veridical perceptual experience. While in the philosophy of action, the natural view is that external events are (at least) part of actions, such that actions are observable, the opposite is true in the perceptual case. No reasonable person would expect to be able to see others' perceptual experiences; saying that publicly observable external items are part of perceptual experience is a radical view to hold. Below, I develop a modest position of the experience-world relation: one that keeps veridical perceptual experience in the head and yet can make sense of the world-presenting thesis.

The two kinds of event that I operate with are *perceptual experience* and *veridical perceptual experience/conscious perception*. I take them to be the analogs of *trying* and *successful attempt/action* respectively. Just like trying, perceptual experience is a fallible kind that is not neutral about whether the success obtains. The success condition for perceptual experience is determined by whether the experience presents the pertinent external items to the subject. As a fallible kind, having a perceptual experience does not imply that its success condition is met. The literature usually uses the term “as of” to indicate its fallibility. For instance, it uses the term – perceptual experience *as of* (say) a cat, rather than *of* a cat – to indicate that a perceptual experience with the success condition of presenting a cat could fail to present a cat to the subject. While perceptual

experience is a common kind shared by veridical and non-veridical perceptual experiences, perceptual experience, just like trying, is geared toward success. This idea has been captured by characterizing perceptual experience as the kind of event that *purportedly* presents external items.<sup>11</sup> The word *purportedly* is introduced to indicate that even though perceptual experience can be non-veridical, veridical perceptual experience is metaphysically prior to non-veridical perceptual experience.<sup>12</sup>

Just like trying, I maintain that perceptual experience is an inner event. To be compatible with science, which tells us that perceptual consciousness is realized in the brain, I take perceptual experience to be composed of neural activity going on in the head.<sup>13</sup> One has a perceptual experience as of a cat in virtue of undergoing neural activity, *not* in virtue of being related to a cat or anything else external.

Unlike perceptual experience, *veridical perceptual experience/conscious perception* metaphysically depends on relations to external items. Analogous to Hornsby's treatment of action, conscious perception is the successful sub-kind of perceptual experience: it is

---

<sup>11</sup> See, for example, Nudds (2009).

<sup>12</sup> Schellenberg (2018) captures the same idea by maintaining that perceptual experience is constituted by the employment of perceptual capacities: perceptual capacities are fallible, but the successful employment is metaphysically prior to the unsuccessful one.

<sup>13</sup> By science, I refer to the *mainstream* perceptual science. I set aside non-mainstream approaches, such as the enactive approach sketched in Noë (2006). Noë maintains that perceptual experience is not just realized in the brain, but also in bodily activities.

Accepting mainstream perceptual science doesn't mean denying the significance of bodily activity for conscious visual perception. It is a pre-theoretical point that we often move around to see things. Think about seeing things that are rich in detail, such as a Van Gogh painting. We move closer to the painting in order to have a better look of it, uncovering more of its details. We could accommodate this fact without making bodily activity part of what realizes perceptual experience, for instance, by maintaining that perceptual experience *causally depends* on bodily activity. See Block (2005) and Burge (2010) for expositions of mainstream perceptual science.



perceptual experience that presents the pertinent external items to the subject. To present external items, the mere having of perceptual experience is not sufficient. It is also essential that the subject is appropriately related to those external items. From the science, we know that this relation is established by a causal process that starts with light reflecting off external objects and ends with the generation of perceptual experience. For instance, consciously perceiving a cat requires (1) having a perceptual experience as of a cat, and (2) that the experience has events such as retinal stimulation from light reflecting off the cat as its distal cause.

While perceptual experience and the relation to the pertinent external items are both essential for conscious perception, conscious perception metaphysically depends on them in different ways. Applying Hornsby's proposed metaphysical relation between trying and action, i.e., composition, to the perceptual case delivers the following view:

The perceptual experience – the inner neural activity – *composes* conscious perception.

The relation to the pertinent external items is *essential for* conscious perception, but it is *not part of* conscious perception.

On this view, conscious perception of a cat has no additional part that perceptual experience as of a cat does not have; conscious perception is composed of neural activity and nothing else. But undergoing the neural activity alone does not entail conscious perception of a cat: a subject can have a perceptual experience that *merely* seems to present a cat to her, with no cat present. To consciously perceive a cat, the neural activity

needs to have events such as retinal stimulation from light reflecting off the cat as its distal cause. Those events are essential for, but not part of, conscious perception of a cat.

The metaphysics of conscious perception developed above is an interpretation of the world-presenting thesis: perceptual awareness of the external world is indeed immediate. This is because there is nothing non-external that the subject has to be aware of prior to her awareness of (say) a cat. The composition of conscious perception is just neural activity. In virtue of undergoing the pertinent neural activity, the subject undergoes a perceptual experience as of a cat, which purportedly presents a cat to her. And the gap between purportedly presenting a cat and presenting a cat is not any additional mental activity. Their difference is simply that for the neural activity to be presenting a cat – not just purportedly so – it has to be appropriately caused by a cat.

Below, I argue that my view is preferable to some common interpretations of the world-presenting thesis. More specifically, those views face the question of how to explain hallucination, whereas my view does not.

### *2.1 Making Sense of Perfect Hallucinatory Experience*

As we've seen, relationalism, a common way to unfold the world-presenting thesis, takes veridical perceptual experience to be composed of the relation between the perceiver and the external items perceived and its relata. The experience has the phenomenal character that it does in virtue of its composition, particularly the external items perceived. Here is a representative quote from John Campbell:

[Q]ualitative properties are in fact characteristics of the world we observe; our experiences have the qualitative characters that they do in virtue of the fact that they are relations to those aspects of the world. So looking for the qualitative character of experience in the nature of a brain state is looking for it in the wrong place; we have to be looking rather at the colors of the objects experienced.

(2010: 206)

Relationalism faces the question of how to explain *perfect hallucination*, hallucination that is phenomenally identical to veridical perceptual experience. When a subject hallucinates a whole scene, she *merely* seems to be aware of external items; there is nothing external that she is aware of. Since the hallucinatory experience is not composed of relations to external items, its phenomenal character cannot be explained by relations to external items. Relationalism, which uses relations to external items to explain phenomenal character, thereby leaves mysterious how perfect hallucination is possible.

A number of prominent theorists have suggested a constraint on the external items that contribute to the phenomenal character of perceptual experience. They propose that it is the *properties*, but not the *instances of those properties*, that contribute to the phenomenal character. These theorists do not deny that when having veridical perceptual experiences, the subject is aware of not only properties, but also instances of those properties. Yet, the instantiation of those properties makes no difference to the phenomenal character.<sup>14</sup> This proposal delivers an explanation of perfect hallucination.

---

<sup>14</sup> See Dretske (2000), Johnston (2004), Pautz (2007), Sosa (2007), Tye (2014).

When a subject hallucinates a cat, the subject is in an awareness relation to external properties, such as *cat-ness*. The hallucinatory experience has the phenomenal character that it does in virtue of those external properties. But unlike veridical perceptual experience, those properties are *uninstantiated*.

This uninstantiated property treatment of hallucination does not do justice to the intuition that when a subject hallucinates a whole scene, she is not aware of anything external; she is at best aware of her experience, which is a state of her own. Furthermore, since uninstantiated properties are not spatio-temporally located and not causally efficacious, it is questionable whether they can be the objects of sensory awareness. Accepting the existence of uninstantiated properties might also require a Platonic ‘two realms’-view, which is inconsistent with a physicalist worldview (Schellenberg 2018: 147-8).<sup>15</sup> On the contrary, I maintain that perceptual experience – veridical or not – is not composed of any peculiar entities, and thus perceptual experience does not have its phenomenal character in virtue of being related to those entities. On my view, a perfect hallucinatory experience has the same phenomenal character as a corresponding veridical experience because both experiences purportedly present the same external items, which is further explained by the fact that these experiences are composed of the same neural pattern.<sup>16</sup> The difference in their nature – that only a veridical perceptual experience

---

<sup>15</sup> See Thompson (2008), Gow (2018) and Schellenberg (2018 ch. 6) for detailed discussion of these problems of the uninstantiated property treatment.

<sup>16</sup> My treatment does not imply that a subject must have perceived a specific sort of worldly item in order to have a hallucinatory experience of it. It is possible to hallucinate a unicorn without having seen one. What is required to hallucinate a unicorn is that the subject undergoes the neural activity that purportedly presents a unicorn to her. Presumably, there’s a limit as for what a subject could hallucinate if she has never had veridical perceptual experience of that sort. For instance, a colorblind subject might not be able to

presents external items – is not involved in the composition of the experience, and thus has no repercussion for the phenomenal character.

I suspect that the uninstantiated property treatment is popular because of the assumption that the phenomenal character of perceptual experience must be explained by some object of awareness. Under this assumption, once we've ruled out sense data as the object of awareness, what else could a hallucinating subject be aware of other than uninstantiated properties?<sup>17</sup> But as we've just seen, this assumption can be challenged.

## 2.2 Making Sense of Phenomenological Particularity

Proponents of relationalism often claim that their views have a distinguished advantage of respecting *phenomenological particularity*. Phenomenological particularity is a phenomenal feature of perceptual experience. The idea is that when I have a perceptual experience as of (say) a cat, it phenomenologically seems to me that a *particular object* – a cat – is presented to me. Some – perhaps most prominently Michael Martin (2002, 2004) – have argued that it is necessary to appeal to the objects in the external world themselves to explain phenomenological particularity. Martin writes:

In order to do justice to the intuition that the very object one is perceiving is an aspect of the phenomenology of one's experience, we must recognize that such

---

hallucinate colors, because her perceptual system could not undertake the activity that is required to perceptually present colors to her. But what the limit is isn't a philosophical question, but an empirical one.  
<sup>17</sup> See also Schellenberg 2018 ch.6 (especially p.139-143) for this dialectic.

objects figure within any adequate specification of the particular phenomenal nature of the experience one has at a time. (2002: 196)

If Martin were right, my view would fail to account for phenomenological particularity, because it does not rely on external objects to explain perceptual consciousness.

However, Martin's proposal does not stand. The reason is that perfect hallucination also exhibits phenomenological particularity (Montague 2016: 134-5). But by stipulation, external objects are not pertinent to the nature of hallucinatory experience. When I have a perfect hallucinatory experience of a cat, it still seems to me that a particular object – a cat – is presented to me, even though there is no cat to present. So, one should not appeal to external objects to explain phenomenological particularity.

To explain phenomenological particularity, what's really needed is just that perceptual experience seems to present *particulars* to the subject. For example, my perceptual experience as of a cat seems to present not only the property *cat-ness*, but also the particular, a cat. Seeming to present particulars does not require that there are indeed particulars to present. Particulars only have to be presented in the good case when the experience is veridical. Hallucinatory experience is the bad case in which the experience *merely* seems to present particulars.

### **3. Conclusion**

Let me end with the broader impacts of my discussion on contemporary theories of perceptual experience.

### 3.1 Dissolving the Dichotomy between Common-Kind Views and Disjunctivism

A longstanding dichotomy in philosophy of perception is the dichotomy between common-kind views and disjunctivism. As I am understanding them, they are theses about the *fundamental kind* of the mental event that occurs when one consciously perceives. According to common-kind views, the mental events that occur when one consciously perceives and when one hallucinates belong to the same fundamental kind. According to disjunctivism, these mental events belong to different fundamental kinds. My discussion suggests that this is a false dichotomy. When one consciously perceives, the mental event belongs to *two fundamental kinds*: it belongs to the kind *perceptual experience*, which is shared with hallucinatory experience, and also to the kind *conscious perception*, which is not shared with hallucinatory experience.

The dissolution of the dichotomy is made possible by my distinction between that which is essential for the kind to which a mental event belongs and that which composes a mental event. Once we recognize that the former does not imply the latter, we are open to accept that conscious perception and the corresponding hallucination have the same composition, and belong to a common kind in virtue of the composition, and also that there are additional elements that distinguish conscious perception from the corresponding hallucination that are not involved in the composition of the mental event.

In reply, defenders of the dichotomy could argue that one or the other of the two kinds to which veridical perceptual experience belongs is *not* a fundamental kind. But on what grounds should we decide which kind is fundamental and which is not? As far as I

can tell, there isn't any clear answer to this question. Introducing fundamentality is important because it is easy to attribute a kind to things due to their superficial character. For instance, I can attribute to my coffee mug the kind *my favorite mug on a sunny day*. With regard to the nature of the coffee mug, this kind is trivial; it is not revealing about what the coffee mug is. Introducing fundamentality can distinguish trivial kinds like this one from the significant ones that are reflective of the nature of the coffee mug, such as the kinds about its function and material constitution perhaps. But as we have seen, both kinds that I attribute to veridical perceptual experience are reflective of its nature. When being asked: What *is* veridical perceptual experience, according to my view? An answer would be: It is the kind of mental event that purportedly present worldly items *and* the kind that presents worldly items. Neither kinds is trivial; they are both what veridical perceptual experience fundamentally is.

A further reply on behalf of the dichotomy defenders is to assume that the mental event that occurs when one consciously perceives belongs to an *exclusive* kind – that it only belongs to *one* fundamental kind. But what justifies this assumption? Exclusive kinds are rare once we reflect on things that are slightly complex. For instance, what's a female tiger? An animal, a mammal, a tiger, a female. What's a heart? A physical object, an organ, a heart. So, to maintain that veridical perceptual experience is an exclusive kind is to make an unjustified *ad hoc* claim.

### *3.2 Revitalizing Internal Representationalism*



Throughout the paper, I have been operating with the idea that perceptual experience is composed of nothing other than neural events. This idea maps well onto *internal representationalism*, which holds that perceptual experience is internal representation of the external world. Internal representationalism is the standard scientific picture of perceptual experience.<sup>18</sup> However, internal representationalism is accused of failing the world-presenting thesis. Theorists either abandon internal representationalism on the grounds that it fails the world-presenting thesis (e.g., Campbell 2002; Martin 2004; Tye 2014), or embrace internal representationalism together with the rejection of the world-presenting thesis (e.g., Hoffman 2019; Pautz 2021). My application of Hornsby’s metaphysical structure reveals an alternative. That a perceptual experience presents external items does not mean that external items are part of the experience. “Interiorizing” perceptual experience does not mean “alienating” us from the external world.

---

<sup>18</sup> For scientists’ endorsement of internal representationalism, see Hoffman (2019) and Tononi & Koch (2015).

## Chapter 2. Rediscovering Introspection Through Buddhist

### Meditation

*Abstract.* It is natural to think of introspection as a unique way of gaining self-knowledge, a way that involves *looking inward*. However, the default interpretation of looking inward – the inner-sense theory – is known to have deep problems. In this paper, I develop a novel interpretation by drawing on Vipassana meditation, a paradigmatically inward-looking activity. I argue that Vipassana meditation involves a *sui generis* form of attention that I call “anti-Gestalt attention”. Whereas in more typical forms of awareness one attends to meaningful wholes, in practicing Vipassana meditation one shifts the attention to the sensations that make up the meaningful wholes. Based on this idea, I develop an account of looking inward as anti-Gestalt attention. And with this account, I introduce a new introspective method to acquire self-knowledge that I call the method of phenomenal profile matching. I argue that this method has various advantages over the existing ones. It aligns with our intuitive, pre-theoretical conception of introspection as an inward-looking activity without committing to the inner-sense theory. And it offers a unified route that leads to various types of self-knowledge that are naturally considered as epistemic products of introspection, including deep knowledge about one’s own character and values.

## 0. Introduction

Introspection is a unique way of gaining self-knowledge. Pre-theoretically, it is a way that involves *looking inward*, into one's own mind. As William James famously says, "The word introspection need hardly be defined – it means, of course, looking into our own minds" (1890: 185). Also, we often introspect to learn about something deep and elusive about ourselves, such as our character, values and emotions. Nevertheless, many contemporary philosophical theories of introspection deviate significantly from this pre-theoretical conception of introspection. The locus of discussion is often on 'trivial' self-knowledge – self-knowledge that does not require any effort of searching within and is not about anything deep about oneself, such as knowing that one believes that there will be a third world war, and knowing that one is seeing a red ball.<sup>19</sup> And the method of acquiring such self-knowledge does not involve any sense of inward-looking. For example, in his influential discussion of how to come to know whether one believes that there will be a third world war, Gareth Evans (1982) claims that one "must attend... to precisely the same outward phenomena as [one] would attend to if [one] were answering the question "Will there be a third world war?"" (225).

The goal of this paper is to revamp the pre-theoretical conception of introspection. Traditionally, the idea of looking inward has been associated with the inner-sense theory, which has been under attack in contemporary discussion, for reasons such as that there is

---

<sup>19</sup> See Cassam (2014) for more on the distinction between trivial and substantial self-knowledge.

no direct empirical evidence of a dedicated mechanism of inner sense. I develop a new theory of introspection that does not postulate the existence of such a mechanism.

To do so, I draw on meditation, a paradigmatically inward-looking activity.

Meditation has historically received detailed characterizations in ancient Buddhist texts, and recently received booming interests in cognitive neuroscience. There are many different schools in Buddhism, and many different styles of meditation. The school that I focus on is Vipassana meditation. Ancient Buddhist texts such as Abhidharma characterize Vipassana meditation as the practice that cultivates the awareness of *dhammas*, fundamental phenomenal units that compose the complex phenomenology of our physical and mental states. In the case of seeing, dharma-awareness would mean that the subject – rather than being aware of unitary and discrete objects that endure over time – is aware of the transitory phenomenal units that make up object-awareness.

I argue that Vipassana meditation essentially involves a *sui generis* form of attention, which I call “anti-Gestalt attention”. In practicing Vipassana meditation, one decouples sensations from the meaningful wholes that those sensations constitute, and attends to the decoupled sensations themselves. This form of attention cannot be captured by the existing scientific distinctions about attention (e.g., exogenous *versus* endogenous). But it is beginning to receive scientific understanding, such as being associated with decreased sensory gating, the neural processes that filter out irrelevant sensory stimuli.

The new theory of introspection that I introduce in this paper is *inward-looking as anti-Gestalt attention*. When one introspects in the way I propose, one steps back from

the meanings and categories that are associated with one's sensations, and attends to the sensations themselves. This theory does not amount to saying that introspection is or involves Vipassana meditation. Among other differences that I will explain, anti-Gestalt attention has different levels of skill; and Vipassana meditation requires a very high level of skill, such that the subject is able to, for example, attend to a multitude of fine-grained details of her sensations. Introspection in many everyday cases does not require that level of skill.

I propose a route of how introspection thus characterized leads to substantial self-knowledge. Emotions and character, among other deep conditions that one can learn about introspectively, have distinctive phenomenal profiles related to them. For example, fear is related to sensations of increased heart rate, shortness of breath, and tightening of muscles in the arms and legs when threat approaches; generosity is related to the feeling of warmth when helping others even if one receives nothing in return. Introspection, when practiced well, brings to one's awareness the subtle details of one's inner life that are central to the phenomenal profile of the deep condition that one tries to learn about. And one acquires substantial self-knowledge when successfully matching those details with the right phenomenal profile. I call this route to self-knowledge *the method of phenomenal profile matching*.

I believe that the method of phenomenal profile matching ameliorates the epistemology of self-knowledge acquired through introspection. Representative theories of introspection (which I will canvass in the following section) are often tailored to trivial

self-knowledge. What is believed to justify their setting aside of substantial self-knowledge is the conviction that there is nothing epistemologically distinctive about the acquisition of substantial self-knowledge: one uncovers deep things about oneself through *inference*, which is essentially the same as the way in which one comes to know about the others.<sup>20</sup> However, this predominant narrative has been put under pressure by Cassam (2014) and Lawlor (2009), who point out that even if inference is the way to know deep things about oneself, it is still epistemologically distinct from how one comes to know about the others in the sense that the *bases* of the inference in acquiring substantial self-knowledge is one's own internal goings-on. I accord with Cassam and Lawlor in challenging the predominant narrative. And I take their challenge a step further. My method of phenomenal profile matching, which consists of comparison of phenomenal qualities, is first-person through and through. So, it is a completely different way from how one comes to know about the others. Furthermore, I will show that this method can lead not only to substantial self-knowledge, but also to trivial self-knowledge, with the difference being a matter of the *level of difficulty* in leading to the self-knowledge. Therefore, my theory delivers a more comprehensive and unified picture of how introspection leads to both trivial and substantial self-knowledge that also preserves the epistemological distinctiveness of introspection.

Before I begin, let me make a few clarifications about the scope of this paper. I am not trying to say that the method of phenomenal profile matching is the only route to self-

---

<sup>20</sup> See Cassam (2014) for this narrative.

knowledge, not even that it is the most common route to self-knowledge. My point is simply that it is an *introspective* way of getting self-knowledge. And in comparison to the other “introspective” ways proposed by the other authors, this method aligns better with our intuitive, pre-theoretical conception of introspection. Now, one might wonder about the *value* of introspection, asking: What is the point of introspection (in the way I propose) when there are other ways to acquire self-knowledge? It would take a separate paper to give an adequate answer to this question. But my initial answer is as follows. Unlike the quick and immediate ways of getting self-knowledge emphasized in the literature, introspection (in the way I propose) is particularly useful for bringing to one’s awareness the usually unnoticed details of one’s inner life. It is important for gaining self-knowledge that is somehow “hidden” from us and yet discoverable through careful attention to those details.

The plan of the paper is as follows. In section 1, I begin with an overview of how introspection as an inward-looking act is lost in recent philosophical discussion. I describe the inner-sense theory and its problems, and some major alternatives that abandon the idea of looking inward. I also explain how those theories fail to address substantial self-knowledge. In section 2, I explore the nature of Vipassana meditation, explaining how it involves a *sui generis* form of attention that I call “anti-Gestalt attention”. And based on this form of attention, I develop the method of phenomenal profile matching, and show how this method can lead to trivial and substantial self-knowledge. Finally, in section 3, I contrast my theory with the inner-sense theory, showing how my theory does not require postulating an inner sense.

Here is the plan of the paper. In section 1, I begin with an overview of how introspection as an inward-looking act is lost in recent philosophical discussion. I describe the inner-sense theory and its problems, and some major alternatives that abandon the idea of looking inward. In section 2, I explain the nature of Vipassana meditation, and then derive from it the Meditative Account of Introspection. I show how this account addresses both trivial and substantial self-knowledge. Finally, in section 3, I contrast the Meditative Account with the inner-sense theory, showing how this account does not require postulating an inner sense.

### **1. How Introspection is Lost**

Introspection is something that we naturally turn to in order to learn about our own conditions. The kind of paradigm examples that are the locus of discussion in the literature are mental states that one can learn about without any conscious effort of looking within. For instance, a subject who is consciously seeing a red ball right in front under good lighting condition and has the relevant concepts (e.g., ‘red ball’) can immediately come to know that she is having the visual experience of a red ball without trying to look inward. But naturally, we also rely on introspection to (try to) learn about more elusive conditions about ourselves. Those conditions include psychological conditions like values and character (e.g., ‘Do I value my work over my family?’ ‘Am I a generous person?’), and even conditions that are not strictly speaking psychological (e.g., ‘Am I injured?’). Consider the following examples.



*Injured Tom.* Tom feels some discomfort in his chest while working on a paper. He wants to understand what it is. So, Tom turns his attention away from the outward-directed task of writing a paper, and toward the feeling of discomfort happening inside his body. In attempt to understand the nature of the discomfort, Tom entertains various hypotheses. He wonders whether it is associated with negative emotions, or something purely physical. Drawing on his knowledge that intense negative emotions like grief and anxiety can induce the feeling of tension in the chest, he attends more closely to the fine-grained character of the discomfort: Does it feel like a dull tightness? Or does it rather feel like a sharp pain at a particular spot? If it is the former, Tom will then run through a list of possible events that might lead to negative emotions (e.g., being rejected by a journal, pet's passing away), and evaluate whether those events are bad enough to induce the discomfort. If it is the latter, Tom will then try to recollect the physical activities that he has done in the past few days; maybe he has pulled his chest muscles in his gym session. To further test this latter hypothesis, Tom moves his chest muscle, and observes whether the pain intensifies with movement. In fact, Tom finds out that the pain gets sharper as he moves, and is located at a particular spot. He thereby comes to know that he is physically injured.

*Generous Sally.* Sally has always considered herself as a generous person. In college, she spent extra hours helping her classmates who were struggling with their assignments. And once she started having a decent income, she helped pay for her nephew's college even when she could have spent the money on herself.

One day, Sally wonders if she is really the generous person that she thinks she is. She starts entertaining alternatives: maybe she has been helping others out of a deep desire of being wanted, and maybe she just succumbs too easily to the pressure that others exert on her. She recalls the number of times when she bought things that she didn't like, things that she bought because of the soft pressure that the salesperson exerted on her, and the guilt she has felt when she ended up not buying those things. She recalls those feelings of uneasiness and compares them with what she would have felt if she had chosen not to help her classmates and her nephew. Are those feelings the same, or are they different? To her relief, she finds out that those feelings are different; she helped her family and friends not because she felt pressured to do so, but because she genuinely cared about those people.

In the cases above, the subjects engage in introspection to learn something about themselves that is not immediately obvious to them – his injury in *Injured Tom* and her character in *Generous Sally*. What makes it the case that introspection is involved in these cases is first and foremost that the subjects look within; they focus on their internal goings-on, such as the way the pain feels to him as he moves his chest muscle for Tom and the way it feels to be uneasy for Sally.

This idea of introspection as an attempt to acquire self-knowledge on the basis of attending to one's internal feelings and sensations (, i.e., internal goings-on) should be a natural starting point for theorizing about introspection. An adequate theory of introspection should additionally address the following issues. First, it needs to say more

about the mechanism involved in attending to the internal goings-on. Second, it needs to be mindful of the fact that not every case of attending to the feelings amount to introspection. For example, when I am looking at an intensely beautiful sunset, I do not need to turn inward in order to be attentively aware of the awe. An adequate theory of introspection has to tell us why introspection is involved in some cases of attending to one's feelings but not the others. Third, the theory should be telling about the connection between the internal goings-on that one is introspectively aware of and the self-knowledge acquired on their basis. For instance, how should we understand the relationship between the feelings that Sally attended to and the self-knowledge about her character (i.e., that she is a generous person)?

While the above is the set of issues that one might expect from a theory of introspection, what we actually see in the philosophical discussion is often very different. Many representative theories of introspection are not concerned with attending to internal goings-on. And their reason is not that there is something else internal that are the objects of introspection. Rather, those theories give up the idea of looking inward altogether. Below are two such theories.

1. *The transparency theory* (e.g., Dretske 1995, Tye 2000, Byrne 2005) maintains that introspective knowledge (, i.e., self-knowledge acquired through introspection) of one's mental states is acquired through inference from attending to outward phenomena. Here are two examples given by the proponents. One has the introspective knowledge that one is seeing a red ball in virtue of inferring from one's seeing a red ball. One has the

introspective knowledge that one believes that there will be a third world war in virtue of inferring from one's affirmation that there will be a third world war.

2. *The self-fulfillment theory* (e.g., Burge 1988, Gertler 2000, 2001) maintains that introspective knowledge of one's mental state is automatically acquired when one makes a self-attributive judgment, because the mental state is "embedded" or "involved" in the judgment. For instance, in judging that I am thinking of a banana, I have the introspective knowledge that I am thinking of a banana, because my thought of a banana is contained in the judgment as its part.

It is puzzling why these are theories of introspection, rather than theories of some non-introspective means to acquire self-knowledge. Neither of these theories has looking inward as an essential component of the mechanisms to acquire self-knowledge. Inference from attending outward on the transparency theory is an explicit denial of looking inward; since knowledge of one's mental state is already embedded in the self-attributive judgment, one does not need to make the effort of looking inward to learn about one's state of mind on the self-fulfillment theory. Perhaps what is thought to justify them to be theories of introspection is that the mechanisms proposed lead to answers to *inward-directed questions* (Cassam 2014). These theories provide mechanisms that give answers to questions such as "Do I believe that P?", and "Do I have a red experience?". But this should not suffice for making a theory a theory of introspection. As a familiar example, I can come to know whether I have a mental state  $x$  by reading the outputs of a neuroimaging machine and applying my knowledge about how those outputs relate to  $x$ .

While this process leads to an answer to an inward-directed question, it is not a case of introspection by any stretch. Proponents of those theories might add further qualifications to rule this case out, such as by maintaining that the process involved in answering the inward-directed questions must be “unique to the first-person case” (Schwitzgebel, 2012: 42-3). But there are other easy counter examples. Suppose that there is a red object right in front of me, and I am focusing on it. I can answer the question “Do I have a red experience?” right away without introspection. I know that I have a red experience *in my having of a red experience*. This observation has received philosophical development by Michelle Montague (2016), who argues that conscious experiences constitutively involve “awareness of awareness”. My visual experience of a red object constitutively involves my awareness of myself being aware of a red object. The awareness of awareness, which is unique to the first-person case, provides the non-introspective ground for my knowledge that I am having a red visual experience. Given these difficulties to come up with an alternative plausible definition of introspection, we have to ask: why not simply adopt the idea of looking inward?

An important reason, I believe, is that the idea of looking inward is typically – if not exclusively – associated with the *inner-sense theory* (e.g., Armstrong 1968, Goldman 2006), which has had a bad philosophical press in recent years (Cassam 2014). The inner-sense theory maintains that there is a sense modality that is dedicated to detecting one’s own mental states, just like there are sense modalities (e.g., vision) dedicated to detecting the external world. On this theory, I come to know that I believe that P in virtue of the operation of an ‘inner scanner’ that directly gives me the information that I believe that P.

There are all sorts of reason why the inner-sense theory is unpopular. An empirically driven one is that this theory requires an expensive empirical commitment that there is a dedicated introspective mechanism such as an ‘inner scanner’ that functions to detect first-order states and issue in higher-order beliefs about those states, which we have no direct empirical evidence of (Paul 2014: 297-8).<sup>21</sup>

An epistemic one is that the inner-sense theory has trouble explaining how introspection can lead to deep, substantial self-knowledge, such as knowledge about one’s character. To illustrate, think about *Generous Sally*. Sally does not acquire the introspective knowledge that she is a generous person through a direct, perception-like process. The information <I am a generous person> is not available to her in a way that she can tell right away by scanning her mind. Rather, she has to go through a complicated process involving attending to various feelings, comparing them, and inferring what that means with regard to her character.<sup>22</sup>

Finally, a metaphysical one is that it is hard to see how the inner-sense theory squares with externalism about phenomenology. Campbell (2002) and Tye (2014), among others, argue that the phenomenology of perceptual experience, such as what it is like to see a red object, is an external feature. They motivate this view with the observation that in our typical perceptual awareness, phenomenal features like colors are experienced as features of external items. This view seems incompatible with the inner-sense theory. Given that

---

<sup>21</sup> See Shoemaker (1994) for other challenges to the inner-sense theory.

<sup>22</sup> This is a common problem of all three kinds of accounts discussed in this section. See also Cassam (2014) for discussion.

those phenomenal features do not reside in the mind, information such as <I am having a red experience> should not be made available through scanning the mind.<sup>23</sup>

With all the deep problems that the inner-sense theory suffers from, I am going to propose a new theory of looking inward. To do so, I will look into Buddhist meditation, in particular Vipassana meditation. Vipassana meditation is a good place to look into, because it is a paradigmatically inward-looking activity, and it has received detailed characterizations in Buddhist texts and is recently gaining attention in the scientific community. From it, I will derive a new theory of introspection that does not require postulating an ‘inner scanner’, that can explain the acquisition of deep, substantial self-knowledge, and that can make sense of the externalization observed in the phenomenal features of perceptual experience.

## **2. Rediscovering Introspection Through Vipassana Meditation**

### *2.1 Typical Awareness*

There is a plethora of sensations that are accessible to us every moment, but we are consciously aware of only a small portion of it. Consider vision, one of the most extensively studied area in science. There are billions of bits of information on the retina, and processing every bit of information would overload the visual system. Instead of trying to retrieve every sensible detail of the surrounding environment, our visual system “is exquisitely adapted to the task of extracting conceptual information from visual input”

---

<sup>23</sup> In a similar vein, Boghossian (1989) argues that the inner-sense theory is incompatible with externalist theories of content.

(Potter et al. 2014). To illustrate, consider looking at a lawn. Rather than generating high fidelity representations of every blade of grass (e.g., the individual components of the textures), the visual system prioritizes categorizing the visual input as a lawn, enabling us to see the lawn as a lawn at a glance. The findings coming out of Mary Potter's laboratory show that extracting meaning from a visual image – detecting and categorizing a specified target from visual images – can take place as quickly as 13 ms, which is around one-eighth of the shortest recorded blink times (Mandelbaum 2017: 6).

Visual scientists have identified various mechanisms underlying the quick construction of a meaningful scene. One of these mechanisms is sensory gating. After visual inputs are received by the retina, and before they reach the cortex, the visual inputs travel to a subcortical area called the *thalamus*. Traditionally thought to be a relay station that sends the sensory inputs to the appropriate cortical area for further processing, the thalamus is more recently believed to have the function of a gatekeeper. It filters out irrelevant and uninformative sensory inputs, suppressing them from being further processed. For instance, in a lesion study, it is found that patients with lesions in pulvinar, the largest nucleus of the thalamus, show deficits in filtering distracter information. When these patients are given the task of discriminating target stimuli, they have no difficulty when the target stimuli are shown alone, but their discrimination performance is impaired when salient distracters are present, which suggests that the lesion leads to difficulty in filtering out unwanted information present in the visual display (Danziger et al. 2004; Snow et al. 2009).



Accumulated amount of evidence shows that the thalamus is reciprocally connected to different parts of the brain, and is able to receive feedback signals from cortical areas that are important for object recognition. And based on the feedback signals, the thalamus can adjust which sensory inputs to gate out (Zhou et al. 2016; Wimmer et al. 2015; Sillito et al. 2006). In a recent study coming from the laboratory of Biyu He, Levinson et al. (2021) showed subjects images of various kinds (e.g., animal, face) presented at different levels of contrast, and asked the subjects whether they recognized the image or not. Levinson et al. found a positive difference in response magnitudes in the thalamus in the recognized versus unrecognized trials, which suggests that the thalamus is involved in object recognition. They also found that while object category in the recognized trials could be decoded from the activity pattern of cortical areas involved in object recognition, it could not be decoded from the thalamus. This suggests that the thalamus is not sensitive to the specific content of the visual representation. Rather, the thalamus contributes to object recognition by receiving feedback signals from cortical areas involved in object recognition to filter out aspects of the visual display that are inconducive to object recognition.

## *2.2 Vipassana Meditation and Mindful Awareness*

While our sensory system typically prioritizes the construction of a scene with meaning over revealing the full details of the sensory signals, Buddhist meditation cultivates a different style of dealing with sensations, and is mechanistically different in terms of how sensory signals are processed. There are many different styles of Buddhist

meditation practices. What I am focusing on is Vipassana, the oldest Buddhist meditation practice that is being widely practiced and studied today.

The kind of awareness that Vipassana meditation cultivates, and how it differs from typical awareness, is carefully described in Bhante Gunaratana's influential book, *Mindfulness in Plain English*. He writes:

The goal of our practice is to become fully aware of all facets of our experience in an unbroken, moment-to-moment flow. Much of what we do and experience is completely unconscious in the sense that we do it with little or no attention. Our minds are on something else entirely. We spend most of our time running on automatic pilot, lost in the fog of day-dreams and preoccupations. (1991: 97)

To illustrate with seeing: In our typical experience of seeing, we are only fully aware of a small portion of our visual sensations – the portion that is conducive to the instantaneous construction of a meaningful scene. In contrast, Vipassana meditation cultivates awareness of *all* details of our visual sensations. While this might just be an ideal, given the limited capacity of the visual system, what this points at is essentially a different processing style. Our visual system is by default tuned to extract meanings from visual inputs, and set aside inputs that are not conducive to the extraction of meanings. This idea that our typical visual awareness is dominated by external-world meanings is well recognized not only by visual scientists, but also by philosophers of perception who endorse the view known as *representationalism*, according to which there are no “qualia”, meaningless phenomenal qualities that are intrinsic to our visual experience,

and that the phenomenal character of visual experience is metaphysically correlated with certain representational content. On this view, when I see a dog, I am not aware of the phenomenal character as qualities of my experience; what I am aware of is rather the external-world content (e.g., *a dog*) that the phenomenal character is metaphysically correlated with. When practicing Vipassana meditation, on the contrary, the focus is placed on the visual sensations themselves, not on the small portion of the visual sensations that are conducive to the construction of a meaningful scene, nor on the meanings extracted from the visual sensations. Below is Bhante Gunaratana's characterization of the kind of seeing experience that Vipassana meditation cultivates.

When you first become aware of something, there is a fleeting instant of pure awareness just before you conceptualize the thing, before you identify it. That is a state of mindfulness. Ordinarily, this state is short-lived. It is that flashing split second just as you focus your eyes on the thing, just as you focus your mind on the thing, just before you objectify it, clamp down on it mentally, and segregate it from the rest of existence. It takes place just before you start thinking about it—before your mind says, 'Oh, it's a dog.' That flowing, soft-focused moment of pure awareness is mindfulness. In that brief flashing mind-moment you experience a thing as an un-thing. You experience a softly flowing moment of pure experience that is interlocked with the rest of reality, not separate from it. Mindfulness is very much like what you see with your peripheral vision as opposed to the hard focus of normal or central vision. [Y]et this moment of soft, unfocused awareness contains a very deep sort of knowing that is lost as soon as

you focus your mind and objectify the object into a thing. In the process of ordinary perception, the Mindfulness step is so fleeting as to be unobservable. We have developed the habit of squandering our attention on all the remaining steps, focusing on the perception, recognizing the perception, labeling it, and most of all, getting involved in a long string of symbolic thought about it. That original moment of mindfulness is rapidly passed over. It is the purpose of vipassana meditation to train us to prolong that moment of awareness. (1991: 82)

The passage above not only shows how the kind of awareness that Vipassana meditation cultivates – which Bhante Gunaratana calls *mindfulness* – lacks external-world content, but also tells us why representationalism is such a natural view to hold. As Bhante Gunaratana puts it, mindful seeing – the seeing sub-kind of mindfulness – is hard to achieve. It demands getting over the habit of the mind, which is to put things that we see into object categories, see things *as* such-and-such external items.

Gunaratana's characterization, which stresses the non-conceptual nature of mindfulness, has been challenged for its lack of emphasis on the role of concepts. For instance, the Buddhist scholar Bhikkhu Bodhi argues that Gunaratana's account is inadequate because "the meditator deliberately uses thought and concepts to keep the object before the mind" (2011: 28), and that "precisely labeling the constituents of one's experience" plays an important role in Vipassana meditation (29).

There is no real conflict between Bodhi's concern and Gunaratana's characterization. To see this, we have to be careful about the distinction between the *practice* of Vipassana

meditation and the mindful awareness that the practice cultivates. That the former relies on the use of concepts does not imply that the latter is conceptual in nature. What Bodhi brings up is the well-recognized fact that Vipassana meditation often involves labeling, making soft, mental notes of the sensations that the meditator is aware of. For instance, the meditator may choose to use descriptive terms like “warmth” to label bodily sensations that she is feeling. However, these concepts – rather than being constitutive of mindful awareness – are tools to cultivate mindful awareness, such as (as Bodhi himself points out) keeping the mind from being carried away from the sensations themselves. The instrumental role of concepts is nicely described by Sayadaw U Pandita, one of the foremost masters of Vipassana meditation. He writes, “[I]n meditation it is important to penetrate behind that conventional concept.... Labeling helps us to perceive clearly the actual qualities of our experience, without getting immersed in the content” (2020).

Furthermore, even mindful awareness itself may not have to be concept-free. Mindful seeing, as described by Gunaratana in the passage above, is devoid of categories and concepts about *enduring, external objects* like dogs and tables. This is compatible with mindful seeing being constituted by categories and concepts about fleeting visual sensations, such as demonstrative concepts about the shades that one is seeing.<sup>24</sup> This point is strengthened when we consider the Abhidharma, the canonical sources of mindfulness. The Abhidharma offers a systematic account of sentient experience and the

---

<sup>24</sup> Admittedly though, Gunaratana’s characterization of mindfulness as “pure” awareness is potentially misleading. This is because the notion of pure awareness is usually thought of as a state that is empty of all experiential content and phenomenal qualities (Chadha 2015). Mindfulness, on the contrary, is rich with non-external-object content and phenomenal qualities.

world given in such experience. According to the Abhidharma, Vipassana meditation has the epistemic goal to enable us to see the reality as it truly is. On the Abhidharma worldview, the reality is ultimately constituted by *dharmas*, basic entities that only exist in a moment, “basic in the sense that they cannot be further analyzed into more basic components.” (Chadha 2015: 69). The “medium-sized” physical objects that our ordinary perceptual experience presents to us, such as cats and cars, are nothing but composites or aggregates of dharmas; they “are real only with reference to conception” (Chadha 2015: 74). Mindful awareness, as the culminating state of Vipassana meditation in which one is aware of the reality as it truly is, is awareness of dharmas. This Abhidharma-Buddhist interpretation of mindful awareness is not in conflict with all forms of conceptualization. While it resists the idea that mindful awareness is constituted by concepts about enduring unified objects that last over time, it seems compatible with fine-grained concepts about the momentary nuanced phenomenal qualities, or *dharmas*.

Now, the question is: How exactly does Vipassana meditation work to cultivate mindful awareness, such as mindful seeing? To answer this question, we have to look into the *Satipathanna Sutta* (*The Discourse of the Establishing of Mindfulness*), a discourse attributed to the Buddha himself. The *Satipathanna Sutta* describes how to practice Vipassana meditation, and sets the foundation for subsequent developments of Vipassana meditation, such as Bhante Gunaratana’s teachings.

The *Satipathanna Sutta* describes Vipassana meditation as having a number of stages. The first stage is to focus on the sensations of breath. During breathing meditation, one

makes an effort to pay full attention to every nuanced sensation of breath. One focuses on the tactile sensations of air passing in and out of the nostrils, the expansion of the chest, the rising and falling of the abdomen, etc.. The second stage is to focus on the sensations of bodily activities like walking. During walking meditation, one pays attention to sensations of the walking process, the complex series of subtle motions that make up the simple action of walking. One tries to be aware of the actual sensations throughout the leg area when one is lifting the leg, the tactile sensations in the foot area when one is placing of each foot, etc..

Focusing on visual sensations comes in the later stage of the practice, after one has mastered focusing on the sensations of breath and of bodily activities. While the *Satipathanna Sutta* does not explain why Vipassana meditation takes the stages that it does, it is not hard to see the rationale behind. In our everyday experience, we are typically unaware of our breath and bodily activities like walking. Usually, we just breathe in and out without noticing ourselves doing anything, let alone conceptualizing it in any way. And when asked to attend to the breath, it is natural to attend to the components that make up the breath. The same goes for walking. Usually, we are on the “automatic pilot” mode when we walk. And when asked to attend to my walking, I struggle to find anything to attend to except for the subtle motions that make up my walking. In short, there is no habit of attending in some other way that we have to get over in order to be aware of the sensations of breathing and walking. Learning to be aware of the constituent sensations of breath and walking is just learning to be aware of breath and walking. On the other hand, we tend to bypass the visual sensations

themselves in our typical visual awareness. What makes attending to visual sensations difficult is exactly that we have the natural tendency to be aware of external-world meanings of the visual sensations, rather than of visual sensations themselves. This is the habit that we need to learn to get rid of in order to be aware of the visual sensations themselves.

### *2.2.1 Anti-Gestalt Attentional Strategy in Vipassana Meditation*

The attentional strategy that Vipassana meditation adopts, which enables mindful seeing, is “anti-Gestalt”. While we are typically visually aware of things as categorized wholes (e.g., aware of a dog as a dog), Vipassana meditation calls for decoupling the individual sensations from the categorized whole that they make up, and attending to the individual sensations themselves.

This style of attention cannot be sufficiently captured by the various sets of distinctions that psychologists have made about attention. The first set concerns the unit of attention – what the attention is directed at. Psychologists have distinguished between spatial, object-based, and feature-based attention. Spatial attention means tracking whatever stimuli are present at a specific spatial location, whereas object-based attention and feature-based attention means tracking specific objects and features respectively regardless of their spatial location. This set of distinction fails to capture what is unique about anti-Gestalt attention. When practicing Vipassana meditation, the subject is not selective about the kind of unit that the attention tracks. In fact, Vipassana meditation is sometimes described in the empirical literature as “open monitoring”, meaning that it



openly observes every sensation that arises without focusing specifically on any of them.<sup>25</sup>

The second set of distinctions concerns the manner in which one enters into the attentional state. Endogenous attention is goal-directed and voluntary, whereas exogenous attention is stimulus-driven and involuntary. Since Vipassana meditation is something that the subject intentionally does, it is a kind of endogenous attention. But this characteristic of being endogenous attention alone does not capture the uniqueness of Vipassana meditation. For instance, when I voluntarily attend to a cat as it is running toward me, my attention is endogenous even though it is obviously different from the kind of attentional strategy used in Vipassana meditation. In fact, the two sets of distinctions covered above are primarily used to understand the typical ways of seeing, ways in which we are aware of things as categorized, external things presented to us.

The third set of distinctions more directly concerns the way in which we are aware of things, such as how we look. It is the distinction between diffuse attention and focused attention. When one attends diffusely over a group of items, one does not focus on any of the items in particular. A good example of diffuse attention is when one is sitting on a park bench on a beautiful afternoon, just chilling. While there are lots of items that one is aware of (e.g., children running around, birds, and trees), one is not focusing on any of them in particular. There are admittedly similarities between diffuse attention and the anti-Gestalt attentional strategy that is definitive of Vipassana meditation. When one

---

<sup>25</sup> For example, see Millière et al. (2018).

practices Vipassana meditation, one's attention is distributed over various sensations without being captured by any of them. But an important difference is that while for diffuse attention the subject is not aware of the individual details of the items, for Vipassana meditation the subject is trying to be fully aware of the details of the sensations – in fact, trying to focus on every one of them. So, it is more accurate to describe Vipassana meditation as a combination of focused and diffuse attention.

Bence Nanay (2016) proposes that distributed attention and focused attention could be practiced at the same time. His idea is that one can attend to an object in one manner, and to the object's features in the other manner. He uses this hybrid to explain an atypical way of seeing – the way one sees when one finds something aesthetically pleasing. He argues that seeing a painting aesthetically is crucially a matter of having the attention focused with regards to the object (i.e., the painting) and distributed with regards to the features (e.g., the shades, the spatial configuration).

This is not the place to evaluate Nanay's proposal about aesthetic seeing. What I would like to note is rather that the kind of hybrid that Nanay proposes is not applicable to the kind of attention we find in Vipassana meditation. What is unique about Vipassana meditation is that the distinction between the object and the features of the object becomes moot. When one meditates, one decouples sensations from object categories. Mindful awareness is "one-layered" in the sense that the object of awareness disappears into a flow of sensations, and thus there is no object as distinct from its features in one's awareness.

This unique attentional strategy adopted in Vipassana meditation is beginning to receive neuroscientific understanding. Vipassana meditation is found to be correlated with the deactivation of the thalamus (Fox et al. 2016). As we have seen, the thalamus is involved in filtering out sensory stimuli that are irrelevant for the construction of a meaningful scene. The deactivation of the thalamus maps well onto the idea that the goal of Vipassana meditation is to become aware of all details of sensations.

The neuroscience of meditation is a booming field, and there are a lot more that we can expect to learn from it in the future. One further testable hypothesis that is yet to be studied is that deactivation of the cortical areas involved in object recognition would be found when practicing Vipassana meditation. Since mindful seeing entails seeing things in a way that is de-coupled from object categories, we would expect that when a subject, say, sees a tree while practicing Vipassana meditation, the neural pattern correlated with tree recognition is not activated.

### *2.2.2 The Epistemic Value of Mindful Awareness*

Now, one may wonder: What is the point of Vipassana meditation practice? This question is particularly apt when we consider the time and effort it takes to cultivate mindful awareness. Why bother to practice attending in the anti-Gestalt way? When we look at the ancient Buddhist texts, answers to these questions focus on ethics, concerning how Vipassana meditation changes the way we lead our lives. The ancient Buddhists maintain that the source of suffering is attachment, such as clinging to attractive and desirable objects. Vipassana meditation puts an end to suffering by bringing about the

insight that everything is impermanent. This insight follows from the sort of awareness that Vipassana meditation cultivates, the awareness of things as composites of *dhammas*, fleeting basic entities that exist only in an instant. Presumably, this ethical doctrine could be a turnoff to those who are not into the Buddhist way of living. But contemporary empirical research shows that Vipassana meditation has non-ethical values too.

Contemporary empirical research shows that Vipassana meditation practice has the epistemic value of enhancing the senses. In a seminal study, Lazar et al. (2005) found that long-term meditators have an increased cortical thickness in brain regions involved in attention and sensory processing, where cortical thickness is correlated with performance. The neuroimaging results translate well into behavioral results. For instance, Hodgins et al. (2010) found that meditators perform significantly better than non-meditators in change blindness flickering tasks, which show subjects scenes flashed between two photographs that were identical except for one change. They found that meditators identified a greater number of changes in flickering scenes, and noticed changes more quickly. What these findings collectively show is that meditators have increased perceptual acuity. The increased perceptual acuity can be plausibly explained by the repeated practice of anti-Gestalt attention. Through long-term practice of attending to the subtle details of sensations, meditators become more capable of noticing more sensible features – including those that are not relevant to the quick categorization of the scene – even when they are not meditating.

### *2.3 Back to Introspection*

The anti-Gestalt attentional style offers a fitting description of what Tom and Sally are doing when they are focusing on their feelings and sensations. When Tom introspects, what he is doing essentially involves attending to the nuanced sensations that make up his pain experience, such as the sharpness or dullness of the pain as he moves his chest muscle. Similarly, when Sally introspects, she directs her attention to the subtle details of her mental life, such as the fine-grained phenomenal qualities of what she feels when considering the past possibility of not having helped her friends and nephew.

Despite employing the same attentional strategy, introspection and Vipassana meditation are not identical activities. While the official goal of Vipassana meditation is to lead to the insight that everything is impermanent and to end suffering, this is obviously not what we are trying to achieve when we introspect. Introspection is paradigmatically led by the goal to acquire self-knowledge; and the facets of experience that the subject becomes aware of through anti-Gestalt attention serve as the means to acquire self-knowledge. To illustrate, before Tom starts introspecting, he was not aware of the individual details of the sensations in his chest; his awareness of the pain was too coarse for him to know whether he is physically injured or not. Introspection reveals to him the individual details of the sensations, which enable him to come to know that he is injured.

Given the different goals that Vipassana meditation and introspection are trying to achieve, these two kinds of activities usually require different levels of *skill* in anti-Gestalt attention. Since Vipassana meditation aims to cultivate the awareness of things as

composite of *dharmas*, it requires an extremely high level of skill in anti-Gestalt attention, such that one is able to attend to a multitude of basic phenomenal units. On the other hand, for introspection the required level of skill in anti-Gestalt attention varies depending on the self-knowledge that one tries to acquire. And in most cases, introspection requires a much lower skill level than Vipassana meditation. This is because acquiring self-knowledge introspectively rarely requires one to have the sensations decomposed to the basic level; the sensations that the introspecting subject attends to are usually more limited in scope and less nuanced in depth. For instance, for Sally to learn about her generosity, what she needs is just the level of details that allows her to identify her feelings about helping her nephew and classmates as different from feelings of uneasiness. This is, of course, not to say that introspection will never get to the basic phenomenal units. How far down the decomposing goes depends on the introspective task in hand; and it is possible that in some rare case the introspective task demands awareness of the basic phenomenal units.

I think that it is best to think of the relationship between Vipassana meditation and introspection as having a significant constitutive overlap as well as instrumental. There is the overlap in that they are both constituted by anti-Gestalt attention. Vipassana meditation is instrumental to introspection in that Vipassana meditation practice improves one's introspective ability. This point follows from the epistemic value of Vipassana meditation discussed in the previous sub-section. Since Vipassana meditation practice enhances our senses, it increases the range of sensations that we can be aware of introspectively.

Anti-Gestalt attention is not the whole story of Tom's and Sally's search for self-knowledge. What it does is to bring to Tom's and Sally's awareness the nuanced sensations that are central to the conditions that they want to learn about. To acquire knowledge of those conditions, they have to take the additional step of associating the sensations attended with those conditions. So, how do the associations take place? A careful look at the cases suggests that they adopt a first-person method that I call *the method of phenomenal profile matching*. Let me illustrate with Tom. Recall that upon being aware of the fine-grained phenomenal character of his pain, Tom entertains various hypotheses about the possible conditions that he is in. In entertaining those hypotheses, a significant part of what Tom is doing is to consider the *phenomenal profiles* of the conditions that he hypothesizes himself to be in. For instance, when considering the possibility of having pulled his chest muscle, he thinks, "Is this what it feels like to have chest muscle strains?" He runs through a pattern of phenomenal qualities that is representative of chest muscle strains: the sensations of sharp pain at a particular spot, the increase in intensity of the pain with movement, etc.. And with this phenomenal profile of chest muscle strains in mind, Tom learns that he in fact has chest muscle strains by finding out that the sensations that he is having *matches* with the phenomenal profile of chest muscle strains.

Now, one might wonder: Should we call the whole method of phenomenal profile matching – not just the anti-Gestalt attention part of it – *introspection*? I think that there is good reason to do so. After all, this method consists in comparisons among internal goings-on; and it is a distinctively first-person way to acquire self-knowledge. For

instance, no one else other than Tom himself could adopt this method to come to know that he is injured, since it relies on comparison among phenomenal data, which are only available to Tom himself.

My point above must not be mistaken as saying that every means to acquire self-knowledge is introspection. Even in *Injured Tom*, Tom adopts both introspective and non-introspective means in his search for self-knowledge. An example of the non-introspective part is his recollection of the physical activities done in the past few days in trying to confirm that he is fact physically injured. This part is not introspective, because this way of knowing about his condition is also available to others. For example, those who know Tom well, upon hearing that Tom is feeling pain in his chest, could as well recall that Tom was at the gym a few days ago and speculate from this that Tom might have a physical injury.

The above is the sketch of my theory of introspection. Below, I will spell out more of its details by explaining how it deals with introspective error and considering some possible objections. After that, I will show how it intervenes in the existing debate.

### *2.3.1 Types of Introspective Error*

Corresponding to the two elements of introspection (i.e., anti-Gestalt attention and matching the attended sensations with the right phenomenal profile associated with one's own condition), there are two main ways in which introspection can fail to lead to self-knowledge. First, one can be bad at attending to the details of one's sensations. As we have seen, it can take lots of practice to be skillful at attending to those details. Those



who are not skilled enough may overshoot certain sensations, and undershoot those that are pertinent to the condition that one tries to learn about. An example of this kind of introspective error is the “recency effect”. When one is asked to evaluate how one feels about a relationship, one might have the tendency to overshoot one’s feelings about the most recent interactions with one’s partner. If one just had an argument with one’s partner, one could overdraw the bad feeling about the quarrel and end up under-evaluating what one really feels about a relationship. This kind of introspective error can be avoided through the practice of Vipassana meditation. Through Vipassana meditation, one learns to not overshoot any specific feeling and sensation and be aware of more facets of one’s experience.

But even when one is capable of being aware of all the sensations that are pertinent to the condition that one tries to learn about, such as by being an expert in Vipassana meditation, one is not immune to introspective error. Matching the sensations with the right phenomenal profile requires understanding of the relationship between sensations and the condition of oneself that one wants to learn about. The second kind of introspective error is attributable to the lack of such understanding. To illustrate, Tom’s success in matching his pain sensations with the phenomenal profile of chest muscle strains is based on his understanding of what kinds of sensations are associated with chest muscle strains. If Tom did not understand, for instance, the relationship between sharp pain at a particular spot and chest muscle strains, he would not be able to learn that he is physically injured through his awareness of bodily feelings.

I am not saying that these are the only ways in which introspection can fail to lead to self-knowledge. For example, one's overall conditions like being drunk could also lead to introspective error. But the two sorts of error noted are more unique to introspection.

### 2.3.2.1 *The Generalizability Objection*

Thus far, I have been developing a new introspective way to get self-knowledge based on anti-Gestalt attention that is definitive of Vipassana meditation. My opponents might wonder how common this way of getting self-knowledge is. They might say, "This is not what I seem to be doing when I introspect! Attending in the anti-Gestalt manner is just one among many other ways to introspect."<sup>26</sup> Let me call this worry the *generalizability objection*.

To begin my reply, I would like to note that the generalizability objection is strictly speaking not a problem for my theory. This is because the point that I am making is not that this is *the only* way to introspect, but rather that this is *a* way to introspect.

Furthermore, I think that it is actually more common than it might seem to introspect in the way I propose. An important source of resistance against my proposal could be that it seems too difficult to do. My opponents might say, "When I try to look inside, my mind is just clouded, and I am not able to attend to my sensations in the way that Tom and Sally and the Vipassana meditators do." To reconcile with my opponents' observation, we just have to be aware of the fact that anti-Gestalt attention has *skilled* and *unskilled*

---

<sup>26</sup> See, for example, Schwitzgebel (2012) for arguing that introspection is pluralistic in nature, such that there is no overlapping mechanism shared by all cases of introspection.

forms. In its skilled form, the subject is able to decompose her sensations into some basic units, and be aware of the fine details of her internal goings-on. And it is indeed difficult and takes lots of practice to be skilled at anti-Gestalt attention. But this is the ideal rather than the baseline of what it means to engage in anti-Gestalt attention. Insofar as one is stepping back from the meanings and categories that are pre-introspectively associated with the sensations, and is attending to the sensations themselves, one is already attending in the anti-Gestalt manner. In its unskilled form, the subject could just be aware of a concoction of unclear sensations, as my opponents describe.

“But don’t we sometimes learn about ourselves by attending to the Gestalt?” my opponents might ask. And this is the other source of resistance against my proposal: that self-knowledge can be acquired through awareness of the experience as a whole rather than the constituents of the experience. My opponents might continue, “For example, I know that I am tired by simply being aware of my overall feeling. I don’t have to decompose my feeling to learn about this.”

I think that this observation says nothing about how common it is to introspect in the way I propose. To see this, it is important to remember that not every case of acquiring self-knowledge relies on introspection. As I pointed out, we are not self-blind to our feelings and sensations until the moment when we introspect. When I feel tired, I can usually tell right in my feeling of tiredness that I am tired. With regard to learning about my own tiredness, I have to introspect only when it is not clear to me whether I am tired or not. And in those occasions, what I do is exactly breaking down the Gestalt, attending

to the nuanced constituents of my feelings, and checking if the attended constituents match the phenomenal profile of what it is like to be tired.

#### 2.3.2.2 *The Absence of Distinctive Phenomenology Objection*

A different kind of worry to my proposal is that there is no distinctive phenomenal quality associated with each kind of condition of oneself. As such, my opponents might argue that one cannot learn about what condition one is in by the kind of match between phenomenal qualities that I am proposing. Let me call this worry the *absence of distinctive phenomenology objection*.

In reply, I would like to note that my proposal does not require the existence of a distinctive phenomenal quality for each kind of condition of oneself. What it needs is a distinctive phenomenal *profile*, which is a *pattern* of phenomenal qualities. It is perfectly reasonable that each of the phenomenal qualities that composes a phenomenal profile is shared with many other phenomenal profiles.

Furthermore, I am not trying to defend the thesis that every self-knowledge can be acquired through introspection. It is perfectly plausible that there is a limit as to how much introspection can tell us about ourselves. More specifically, I leave open 1) the possibility that some condition of oneself has no distinctive phenomenal profile associated with it, and thus cannot be known through my proposed method; and 2) the possibility that some conditions of oneself have overlapping phenomenal profiles, and thus the proposed method alone does not suffice for acquiring the knowledge of the specific condition that one is in.

In the above, I have developed my new introspective method of acquiring self-knowledge. Below, I will argue for ways in which it ameliorates the epistemology of introspective self-knowledge.

### *2.3.3 Deep, Substantial Versus Trivial Self-knowledge*

An advantage of my theory over the existing ones is that it offers a unified route that can lead to both deep, substantial self-knowledge that takes conscious effort to acquire as well as trivial self-knowledge that is easy to come by. With my proposed method, the crucial difference between acquiring these two types of self-knowledge lies in the different degrees of complexity of the introspective analysis, such as how much skill is needed to attend to the pertinent details of the sensations, and how difficult it is to grasp the phenomenal profiles of the conditions that one wants to learn about. For instance, in her quest to know that she is a generous person, Sally has to notice some very subtle details of her feelings, such as the difference between the feeling of uneasiness and the feeling that she would have if she had chosen not to help her classmates and her nephew. She also needs to understand the relationship between those feelings and generosity. On the other hand, trivial self-knowledge, such as the self-knowledge that one is having a red visual experience (when there is a red thing right in front under normal viewing condition), requires minimal skill to attend to details of sensations, and the link between the sensations attended and the condition of oneself having a red visual experience is obvious.

### *2.3.4 From Knowledge to Action*

My theory is different from the existing theories in what they take to be fundamental to introspection. Existing theories often maintain that introspection is fundamentally led by the epistemic end of learning about one's mental life. The Stanford Encyclopedia of Philosophy's entry on "Introspection" begins with: "Introspection, as the term is used in contemporary philosophy of mind, is a means of learning about one's own currently ongoing, or perhaps very recently past, mental states or processes" (Schwitzgebel 2019). On the contrary, my theory is not primarily concerned with the kind of knowledge that introspection brings about. It is primarily concerned with what it means for introspection to be an inward-looking activity, leaving open what epistemic ends this activity leads to.

This shift of primary focus to the characteristic activity of introspection, and the flexibility in epistemic ends that comes with it, is a theoretical advantage of my theory. As we have seen, not all paradigm cases of introspection are led by the goal to learn about one's mental life. For instance, Tom introspects to learn about his physical condition, that he is physically injured. We can also imagine an alternative version of *Injured Tom*, where Tom attends to the nuances of his sensations in the same way without trying to learn anything about himself. While Tom does not have the goal to acquire self-knowledge, it seems natural to think that he is introspecting. What these show is that it is overly restrictive to constrain introspection by the epistemic goal to learn about one's own mental life. And since what makes all these cases of introspection is that subject engages in anti-Gestalt attention, my approach offers a more promising direction to understand what introspection fundamentally is.

### **3. Looking Inward Without Inner Sense**

Throughout this paper, I have developed and defended the method of phenomenal profile matching as a new way to acquire introspective self-knowledge by drawing on the anti-Gestalt form of attention that is definitive of Vipassana meditation. To conclude, I explain how this theory offers a new interpretation of looking inward that does not postulate an inner sense.

On my theory, introspection is inward-looking in the sense that it directs attention away from the external-world meanings of internal goings-on (e.g., sensations), toward the internal goings-on themselves regardless of their relevance to external-world meanings, and leads to self-knowledge through comparisons among internal goings-on, matching the internal goings-on that one is having with those associated with the internal condition that one is in.

This interpretation of looking inward does not require the costly empirical commitment of an “inner scanner”. The inner-sense theory maintains that introspection is inward-looking in the sense that it is the exercise of an “inner scanner” that takes first-order mental states (e.g., the visual experience as of a red thing) as inputs to generate second-order beliefs with the content that one is in a certain mental state (e.g., the belief that I am seeing a red thing). According to my theory, on the other hand, introspection does not tell us about our state of mind by detecting our mental states right away. Rather, it works by re-distributing attention to the more basic elements that make up our mental states, and we learn about our state of mind through matching those elements with the

phenomenal profile associated with our state of mind. As we have seen, science is beginning to pin down the mechanistic details underlying this form of attention re-distribution; it involves the modulation of thalamic activity, and relatedly of sensory gating.

My discussion undermines externalism about perceptual phenomenology, the metaphysical view that conflicts with inward-looking theories of introspection. Recall that this view is motivated by the observation that perceptual phenomenology is experienced as external. It stands as a metaphysical explanation of this observation: that the phenomenal features are in the external world. My discussion undercuts this motivation by offering an alternative, psychological explanation: that it is the habit of the mind to put things that we sense into external categories like cats and cars, and allocate attention predominantly to the external-world meanings. Given the goal of this paper, I would not try to defend my psychological explanation against the metaphysical one.<sup>27</sup> For my purpose here, it suffices to show that my theory has more resources to address this challenge to inward-looking theories than the inner-sense theory.

Epistemically, my theory has greater flexibility than the inner-sense theory; it offers a unified route that can lead to both self-knowledge that is easy to acquire and self-knowledge that is not immediately available in a perception-like manner. For self-knowledge that is easy to come by, the execution of the method of phenomenal profile

---

<sup>27</sup> See also Matthen (2010) for psychological explanations of the externalization of perceptual phenomenology.



matching is simple. For instance, to acquire the self-knowledge that one is having a pain in the left knee, one does not need to attend to the nuances of the pain experience (e.g., the specific intensity of the pain experience), and the relationship between the sensations attended and the condition of having a pain in the left knee is easy to grasp. On the other hand, gaining deep self-knowledge through introspection, like Sally's knowledge of her generosity, is achieved through skillfully attending to the nuanced details of one's inner life and with deep understanding of what those details mean with regard to oneself.

### **Chapter 3. Perceptual Surprise and Aesthetic Appreciation**

*Abstract.* Works of art that we find aesthetically pleasing often involve perceptible features that surprise us. However, there is as yet no established account of how perceptual surprise contributes to our aesthetic experience. In this paper, I develop and defend the *Surprise After Surprise Account*, according to which perceptual surprise contributes to aesthetic experience by motivating the subject to discover more perceptual surprises within a determined object. Inspired by the distinctions that Kendall Walton made in “Categories of Art,” I argue that surprising perceptible features are contra-standard relative to the categories that the subject attribute to the object and its features. So, aesthetically significant perceptual surprises are those that motivate the subject to discover more contra-standard features or more ways in which the features already observed are contra-standard. Along the way of developing and defending this account, I offer a novel defense of relationalism about aesthetic qualities (which my account entails), argue that recent exchange between aesthetics and philosophy of perception has neglected the important role of the subject in aesthetic experience, and make progress in resolving a puzzle in recent empirical work on aesthetics.

## 0. Introduction

Works of art that we find aesthetically pleasing often involve perceptible features that surprise us. And our aesthetic experience of those works of art is often constituted by the surprises that we experience. Consider Kandinsky's *Yellow-Red-Blue*. If you experience the painting in the way that I do, you will find your aesthetic experience being constituted by series of perceptual surprises. When you first look at the painting, your attention will be drawn to the face-like shape on the left side. But once you start seeing those lines as a face, the yellow rectangle and the smaller light-blue rectangle with a bright white circle inside will strike you as out of place. Those shapes resemble a giant door below the sun; and seeing them this way evokes the feeling of surprise – how can a giant door be on top of a face! And right at the moment when you see the white circle as the sun, you will be doubly surprised by the larger circle above, which looks even more like a sun – the bright yellow surrounding the circle looks like the sun's corona and the grey circle looks like the sun during a solar eclipse. But how can there be two suns? The prominent of perceptual surprises in aesthetically appreciating Kandinsky's *Yellow-Red-Blue* is nothing exceptional. Stravinsky's *Rite of Spring*, Kubrick's *2001: A Space Odyssey* – pieces that we consider as masterworks – are masterworks (in part) because they are irremediably rich with features that surprise us. Because of their surprising features – such as the unique tonality and rhythm of *Rite of Spring* and the shocking gigantic black monolith in the pre-human world in *2001: A Space Odyssey* – those works are described as revolutionary, absolutely new and original.

Despite the significant role that surprise plays in aesthetic appreciation, not much has been said about this topic.<sup>28</sup> While surprise is influentially taken to be one of the six basic emotions by the psychologist Paul Ekman, unlike the other basic emotions such as anger and sadness, it is relatively unexplored in both philosophical and psychological literature on emotions.<sup>29</sup> A naïve view of surprise, which is often taken for granted, maintains that surprise occurs in virtue of being caused by events that are *unexpected*. But the expectation-based view has obvious problems. Baras and Na’aman (2021) point out that not all unexpected events elicit surprise; and Judge (2018) argues that not all cases of surprise arise out of the unexpected. In the context of aesthetics, the expectation-based view faces extra difficulties. For instance, once one has viewed an artwork several times, the perceptible features of the artwork plausibly become expected. But this does not seem to take away the surprise that one experiences when aesthetically appreciating the artwork.

My primary goal in this paper is to understand the nature of perceptual surprises – surprises that arise from perceiving surprising features – in the context of aesthetic appreciation. More specifically, I explore what sort of perceptual surprise constitutes our aesthetic experience, and how it does so. My investigation has two parts. Since the expectation-based view is unsatisfactory, the first part develops a new theory of surprise that works in the aesthetic domain. To do so, I draw on Kendall Walton’s seminal work “Categories of Art” (1970). In this paper, Walton distinguishes perceptible features in an

---

<sup>28</sup> For exceptions, see Declos (2014) and Judge (2018).

<sup>29</sup> For a recent philosophical attempt to break the silence, see Baras and Na’aman (2021).

artwork into three kinds: features that are *standard*, *variable*, and *contra-standard* relative to the category that one attributes to the artwork. Drawing on this set of distinctions, I develop the *category-based theory of surprise*, according to which surprising perceptible features are contra-standard relative to the categories that the subject attribute to the object and its features. In contrast to the expectation-based view, the category-based theory does not regard the unexpected as essential for surprise. As I will show, expectations play a more indirect role in perceptual surprise than assumed, by permeating the subject's categorization of objects and its features.

Based on the category-based theory of surprise, the second part develops an account of how perceptual surprise contributes to aesthetic appreciation. I argue that perceptual surprise does so by motivating the subject to continue engaging with the object for the sake of discovering more contra-standard features or more ways in which the features already observed are contra-standard – in short, more perceptual surprises. The perceptual surprises that constitute our aesthetic experience are those that motivate the discovery of more perceptual surprises. I dub the view the *Surprise After Surprise Account* (or SAS, for short).

On the SAS, aesthetic experience is the experience of a dynamic play between the subject and the object of aesthetic appreciation. With the example of Kandinsky's *Yellow-Red-Blue*, the painting is aesthetically pleasing because it entices the subject to form certain hypotheses about what she sees, surprises her with features that conflict with

those hypotheses, leads her to form new hypotheses about what she sees, and surprises her again.

The SAS has import on various existing discussions in aesthetics and beyond, which I will detail along the way. First, it has philosophical import on the metaphysics of aesthetic qualities. It results in a version of relationalism, maintaining that aesthetic qualities (in the cases discussed) emerge in the interaction between the subject and the object. Second, it sheds light on how some intentionally boring art can be aesthetically pleasing. Third, it complements recent attempts to investigate aesthetic experience through the lens of philosophy of perception (e.g., Nanay 2016; Ransom 2022). Finally, it helps make progress in the empirical study of aesthetics. Neuroscientists are puzzled by the recent findings that aesthetic experience activates the brain's Default Mode Network (DMN), which is primarily associated with self-related processes. The SAS resolves the puzzle by showing that aesthetic experience is as much about perceiving external features as it is about how one reacts to those features.

## **1. What is Perceptual Surprise?**

### *1.1 Surprise as Arising from the Unexpected*

Suppose that you just return home. Once you open the door and look inside, you are surprised by what you see. The wall close to the door, which is typically plain white, has a huge blob of red paint on it. How should we understand the nature of perceptual surprise?

On the naïve view, perceptual surprise is connected with the unexpected. There are various ways to develop this view. On the first version, you are surprised by what you perceive because what you perceive is different from what you expect to perceive. On the second version, you are surprised because you do not expect to perceive what you perceive. There are subtle differences between these two versions. The first version requires that you form an expectation about what you are to perceive. Illustrating with the case above, this version maintains that for you to be surprised by the huge blob of red paint, you have to form the expectation upon entering your home that the wall is plain white. The second version does not have this requirement, since not expecting  $x$  does not entail expecting not  $x$ . For instance, when I am in my hotel room and hear through the walls that the person living next door whom I have never met speaks French, this is something that I do not expect to hear. Nonetheless, I have not formed any expectation about what languages I am to hear through the walls. Since on the naïve view expectations play a crucial role in surprise, I dub the view the *expectation-based view of surprise*.

Both versions of the expectation-based view face easy counterexamples. For the first version, consider yourself roaming around while mind-wandering. As you turn around the corner of a street, you see a pink elephant. Since you are mind-wandering, you make no expectation about what you are to see. And yet, the sight of the pink elephant surprises you. For the second version, consider looking out the window and observing the exact shade of the sky. You do not expect to see that exact shade. But seeing it is unsurprising

(Baras and Na'aman 2021). In other words, you are not surprised by seeing what you do not expect to see.

With regard to these counterexamples, one can – as the proponents of the expectation-based view would do – finetune the relation between surprise and the unexpected, adding extra qualifications with the aim to ruling out the counterexamples.<sup>30</sup> Alternatively, one can give up the centrality of the unexpected altogether and identify other features as essential for surprise. Judge (2018) has recently argued that the second option is the right way to go. Her reason is the expectations are in essence future-oriented, but there are cases of surprises that are based solely on assessment of the situation at the present moment. So, any account of surprise that takes expectations to be central would in principle fail to account for those cases.

It is not my purpose to consider ways to mend the holes of the expectation-based view of surprise and evaluate how well they fare. Rather, I will simply go for the second option, and develop a new theory of surprise that gives the central role to *category* instead of expectation. But before that, I would like to bring up a problem in the context of aesthetics that both versions of the expectation-based view above face.

### *1.1.1 The Problem of Repeated Viewings*

Consider one of the early scenes in Stanley Kubrick's *2001: A Space Odyssey* – the scene where the tribe of primitive apes who were beaten up by a different tribe discover

---

<sup>30</sup> See Baras and Na'aman (2021) for taking this approach.



the gigantic black monolith outside their dwelling, touch it and marvel at it. For the fans at least, this is a truly remarkable scene that remains surprising even after watching the movie for a hundred of times. This phenomenon of finding something surprising after repeated viewings is common with regard to appreciating one's favorite pieces of art. But at the same time it can hardly be explained in terms of the unexpected. After watching the movie closely for a couple of times, the audience could have and might even have formed expectations about what they are about to see. So, the surprise that the audience experience upon re-watching the movie cannot be due to the unexpected.

In reply to this objection, proponents of the expectation-based view might argue that the scene is surprising because the audience have *forgotten* what they are about to see upon re-watching the movie. And since they could not remember what they are about to see, the scene that they see is unexpected even upon re-watching.

This reply is flawed once we consider the experience of a diehard fan (re-)appreciating her favorite artwork. That opening scene of *2001: A Space Odyssey* is just too remarkable to be forgotten by the true fans of the movie. A diehard fan could even recall in vivid detail the execution of the scene – the eerie music in the background as the apes are dancing around the monolith, and the exact moves that the apes make when marveling at it.

A different and more interesting reply that the proponents of the expectation-based view could make is to argue that the emotion that the audience experience during repeated viewings is not surprise but something else. But what emotion could that be?

The most plausible candidate seems to be *amazement*. They might argue that when the audience are re-watching the opening scenes of *2001: A Space Odyssey*, they are amazed but not surprised.

This reply is based on the questionable dichotomy between surprise and amazement. On an alternative and more plausible way to think about the relation between these two kinds of emotions, amazement is a *sub-kind of surprise* – a sort of *pleasant surprise*. To support this line of thought, consider the basic emotions – emotions that are identified as the ‘building blocks’ of all other emotions. Paul Ekman (2007) influentially suggests that there are six basic emotions: sadness, fear, anger, disgust, happiness, and surprise. Operating with Ekman’s classifications, the first four emotions – all of which are negatively valenced – can play no major role to play in amazement, which is positively valenced. Surprise can either be positively or negatively valenced. And it is plausible to treat amazement as a positively valenced sub-kind of surprise – perhaps surprise mixed with happiness. That amazement is a sub-kind of surprise is supported by the fact that amazement fits with the characteristic facial expressions associated with surprise, such as raising the brows, widening the eyes, and opening the mouth.

Now we have seen the expectation-based view of surprise and how it fails to deal with the aesthetic domain. Let’s proceed to the development of a new theory of surprise – one that can account for surprises in the aesthetic domain.

### *1.2 Perceptual Surprise as Constituted by Contra-Standard Features*

In “Categories of Art,” Kendall Walton (1970) famously argues that aesthetic appreciation depends on the way we categorize the work. He proposes that the categories that we attribute to the work depends on three types of perceptible nonaesthetic features, which he calls *standard* features, *variable* features, and *contra-standard* features. Standard features are features in virtue of which an artwork belongs to the category that it does. Variable features are features whose presence or absence do not affect the artwork’s belonging to the category that it does. Contra-standard features are features that count against the work’s category membership. To illustrate, Walton gives the example of the category *painting*. The flatness of a painting and the motionlessness of its markings are standard features relative to the category *painting*; the painting’s particular shapes and colors are variable features; and “a protruding three-dimensional object or an electrically driven twitching of the canvas would be contra-standard relative to this category” (340). To put this set of distinctions to work, Walton develops a thought experiment about Picasso’s *Guernica*. He asks us to imagine a society that does not have an established medium of painting, but produces a kind of artwork called *guernicas*. Guernicas have surfaces with the same colors and shapes as Picasso’s *Guernica*, but the surfaces are molded to protrude from the wall. In our society, where *Guernica* is appreciated under the category *painting*, the flatness is a standard feature, which we tend to bypass in appreciating the work. We focus on variable features such as the colors and figures on the surface, and are led by these variable features to react to the artwork as violent, dynamic, vital, and disturbing. In the imagined society, however, the colors and figures are standard features relative to the category *guernicas* and are not remarkable for aesthetic

appreciation. To them, the flatness is a variable feature, and on the basis of it they would find Picasso's *Guernica* cold, stark, lifeless and perhaps dull and boring.

While Walton's goal is to understand the nature of aesthetic appreciation, and the categories that he focuses on are art categories like *painting* and *sculpture*, his distinctions of the three kinds of features are applicable to non-aesthetic perceptual cases. I illustrate this point with vision, the most studied modality by perceptual scientists.

Our perceptual systems are adapted to constructing a categorized, meaningful scene. And our perceptual awareness is constituted by meaningful objects and their features. For instance, when there is a dog in front of me, what I am aware of are not billions of bits of incoming sensory data on the retina. Rather, what I typically see is a categorized whole – a dog. In more detail, there are billions of bits of information on the retina, and processing every bit of information would overload the visual system. Instead of trying to retrieve every sensible detail of the surrounding environment, our visual system “is exquisitely adapted to the task of extracting conceptual information from visual input” (Potter et al. 2014). In the case of seeing a dog, rather than generating high fidelity representations of every detail (e.g., the individual components of the textures of the dog's hair), the visual system prioritizes categorizing the visual input as a dog, enabling us to see the dog as a dog at a glance. The findings coming out of Mary Potter's laboratory show that extracting meaning from a visual image – detecting and categorizing a specified target from visual images – can take place as quickly as 13 ms, which is around one-eighth of the shortest recorded blink times (Mandelbaum 2017: 6).

We can apply Walton's distinctions of three types of features to the categories in visual awareness. Illustrating with the category *dog*, standard features would include a round-ish nose, a rectangular-ish body, and a tail. Variable features would include the dog's specific color and size. Contra-standard features would include horns on the head, and wings on the body.

With these distinctions made, here is my theory of perceptual surprise.

**Category-Based Theory of Perceptual Surprise:** Perceptual surprise is constituted by perceptible features that are contra-standard relative to the categories that the subject attributes to the object and its features.

On this account, when I am surprised by what I see, I am surprised in virtue of perceiving features that conflict with the categories that I (or my perceptual system) attribute to the object and its features.

To put this theory to work, consider being in a dog park, and there is a small, rectangular-ish, four-legged creature quickly running toward you. With the coarse information about the shape and size, and perhaps also taking into account the contextual information of where you are, your perceptual system quickly categorizes the creature as a dog. And you are consciously aware of the creature as a dog. But suppose that you then take a more careful look at the creature, and notice that it has horns on its head and wings on its body. Obviously, you would be surprised by what you see. This is because those features (i.e., the horns and the wings) are contra-standard relative to the category *dog*.

In contrast to the expectation-based view, the category-based theory maintains that perceptual surprise arises out of perceiving things that conflict with the categories one attributes to the scene, not with the expectations that one has about the scene. My theory is free from the worries and counterexamples that expectation-based view faces. To start, let's reconsider the two counterexamples that the expectation-based view faces. When you look out of your window, you are unsurprised by the exact shade of the sky. This is because the exact shade of the sky is a variable feature of the category *sky*. When you turn around the corner of a street and see a pink elephant, you are surprised by what you see. This is because you attribute categories like *typical American neighborhood* to the scene. The pink elephant is a contra-standard feature relative to this category. By shifting the focus from expectation to category, my theory is compatible with Judge's point about the present-orientedness of surprise. The categories attributed to the scene are evaluations about the present situation. So, perceptual surprise is present-oriented in the sense that it is constituted by conflict between the features that one is presently perceiving and one's evaluation about the present situation.

It is worth noting that I am not saying that expectation plays no role in surprise. But its relevance is more indirect. Expectations can affect the way that the subject categorizes the scene. Consider the famous duck-rabbit ambiguous image. Suppose that you originally see the image as a duck. To switch your perception of the image to a rabbit, a helpful way is to be told that the image can be seen as a rabbit, and thereby form the expectation to see it as a rabbit. The expectation re-orientes your attention, shifting the focus from the left side of the painting (which you would have seen as the beak of a

duck) to the right side of the painting (which you should now see as the mouth of the rabbit), thus enabling you to categorize the image as a rabbit.

The category-based theory has a simple solution to the problem of repeated viewings. One remains surprised after viewing an artwork for multiple times because one continues to categorize the features of the artwork in a way that is a misfit with its other features. An example from *2001: A Space Odyssey* would be the futuristic monolith in the pre-human world. Categorizing the scene as *pre-human world*, the futuristic monolith is a contra-standard feature, thus eliciting surprise. No matter how many times one has watched the movie, one is still inclined to categorize the opening scene as *pre-human world*, in which the monolith is unbelonging. So, the surprise remains.

## **2. Perceptual Surprise in Aesthetic Appreciation**

Now that I have derived an account of perceptual surprise, we are ready to return to the key theme of this paper, which is to explore the role of perceptual surprise in aesthetic appreciation. More specifically, when does perceptual surprise contribute to aesthetic appreciation? And how it does so? To start, I would like to take note of Walton's discussion on this topic. In "Categories of Art," Walton admits that contra-standard features can be "shocking, or disconcerting, or startling, or upsetting, just because they are contra-standard for us" (352). While Walton uses expressions that are suggestive of surprise to describe contra-standard features, he does not explore its positive contribution to aesthetic appreciation. Rather, he thinks of contra-standard features as obtrusive to our aesthetic appreciation. His reasoning can be summarized as follows. To aesthetically

appreciate a work of art, we have to put it under certain art categories (e.g., painting, sculpture, sonata). Contra-standard features relative to art categories are obtrusive because they conflict with the art categories which we situate the artwork in and based on which we aesthetically appreciate the work of art.

Walton's discussion of contra-standard features is overly restrictive in two ways. First, he only considers features that are contra-standard relative to *art categories of art objects*. For example, a protruding three-dimensional object on a canvas conflicts with the category *painting*, which is an art category attributed to the whole art object. But perceptual surprise that is significant for aesthetic appreciation often arises from features that are contra-standard relative to *non-art categories attributed to the features within art objects*. Consider the Kandinsky example again. As we have seen, the aesthetic appreciation of *Yellow-Red-Blue* is constituted by conflicts between the categories *face-like shape* and *giant-door-like shape*, and between the categories *sun-like shape* and *larger-sun-with-corona shape*. These are features *within* the painting. And they might not be considered as art categories; these categories can be attributed to things that have nothing to do with aesthetics, such as when one sees a curly wall and attributes the category *face-like shape* to it.

Also, the scope of aesthetic appreciation is plausibly not limited to objects that are intended to be works of art. We can aesthetically appreciate natural objects; and the aesthetic appreciation can as well be constituted by perceptual surprise. Snowflakes, for example, are aesthetically pleasing to see (in part) because its perceptible complexity and



uniformity is surprising. The surprise is due to the fact that the perceived complexity and uniformity of the shape is contra-standard relative to the category *natural object*, which tends to lack the uniform complexity in its appearance.

Second, even features that are contra-standard relative to art categories of art objects, which Walton's discussion focuses on, can contribute to aesthetic appreciation. Avant-garde art, in particular, is appreciated in virtue of having such features. Consider Duchamp's *Fountain*. By treating a readymade urinal as an art object, it challenges the category *sculpture*, which was not meant to be readymade. Duchamp's *Fountain* is aesthetically appreciated particularly because it has this readymade feature that challenges the boundary of what a sculpture can be.

This second point above can be used to make sense of artworks that are well-appreciated but intentionally boring. Elpidorou and Gibson (forthcoming) point out that good boring art is neglected in traditional discussion of aesthetics. An example of good boring art that they give is Andy Warhol's movie *Empire*. This eight-hour-and-five-minute-long movie is shot with the camera pointing at the Empire State Building and nothing else. Elpidorou and Gibson claim that this movie is intentionally boring, and that it is "successful as art precisely because of the manner in which they bore". They argue that good boring art is good in virtue of its place in the artworld. When we situate Warhol's *Empire* in a network of relationships with the historical context, methods of production, etc., it is good in virtue of making an innovative and daring move in the artworld. We appreciate the movie not in virtue of the perceptual features *per se* (e.g., the

shade and shape), but in virtue of how those perceptual features stand against the habits of the artworld.

While Elpidorou and Gibson do not put things in my way, what they are pointing at is essentially that good boring art is surprising. It is surprising in virtue of having features that are contra-standard relative to the well-established art categories of art objects. When the audience enter the cinema to watch *Empire*, they categorize what they are watching as *movie*, whose standard features include at least having a plot, having living creatures as protagonists, etc.. *Empire* is surprising in being contra-standard relative to this category, presenting the audience just the Empire State Building.

My discussion above shows that different sub-kinds of perceptual surprise can contribute to aesthetic appreciation. Now, this might seem to be making no progress as for the question of what sub-kind of perceptual surprise contribute to aesthetic appreciation. If we summarize the discussion thus far, we will get the following:

Perceptual surprises that are involved in aesthetic appreciation are those that arise from the perception of features that are contra-standard relative to both art and non-art categories attributed to the perceived object and its features.

This is literally the same as saying that *any* sub-kind of perceptual surprise can contribute aesthetic appreciation.

But to see the issue in a positive light, I take it as showing that it is wrongheaded to try to single out the perceptual surprise that contributes to aesthetic appreciation in terms of

the sorts of categories and perceived objects and features involved. Rather, we should try to pin down the uniqueness of the aesthetically significant perceptual surprise in some other way. Below, I develop and defend my own proposal, which I call the *Surprise After Surprise Account*, or SAS for short.

**The Surprise After Surprise Account (SAS):** Perceptual surprises that are involved in aesthetic appreciation are those that invite continual engagement with the object for the sake of discovering more contra-standard features or more ways in which the features are contra-standard.

The SAS is in continuum with the long tradition in aesthetics – most famously developed by Kant – that aesthetic pleasure is *disinterested* in the sense that the subject gains aesthetic pleasure simply in contemplating the object, rather than consuming or using it in any other way. The SAS is a unique development of this idea, maintaining that perceptual surprise constitutes aesthetic pleasure in virtue of motivating the subject to keep searching for new ways in which the perceptible features surprise her.

Let me begin illustrating the SAS with the Kandinsky case. The SAS maintains that perceptual surprises contribute to the aesthetic appreciation because they invite the subject to keep looking at the painting in the search of more perceptible features that conflict with the categories attributed to the painting's features. Upon seeing the giant-door-resembling shape on the face-like shape, the elicited surprise that the subject experiences motivates her to keep her attention focused on the painting and to seek more conflicts between categories within the painting. Aesthetic pleasure emerges from this

ongoing process of categorizing perceptible features in one way, being surprised by the features that conflict with those categories, being further surprised by other conflicts between categories, and so on and so forth.

The SAS is open to the possibility that the continual engagement with the object can be more intellectual in nature. The avant-garde artworks discussed – Duchamp’s *Fountain* and Warhol’s *Empire* – are unlike Kandinsky’s *Yellow-Red-Blue* in the sense that they do not have waves of perceptible features to surprise the viewers. But there are many different ways in which the same sets of perceptible features can surprise the viewers. For instance, after being surprised by how Duchamp’s work challenges the category *sculpture* in virtue of being a readymade object, the viewer would then zoom in onto the “R. Mutt” signature that is easily noticeable on the front side of *Fountain*. And they would wonder: Does the fact that the readymade urinal is signed make it less like other readymade objects but more like objects that are traditionally recognized as belonging to the artworld? In other words, would the signature be contra-standard relative to the category *not-exactly-a-sculpture*, a category just assigned to *Fountain* a moment ago?

The SAS is able to distinguish perceptual surprises in the aesthetic domain from those in the non-aesthetic domain. Consider again looking at a wall in your home and being surprised by a huge blob of red paint on it. The perceptual surprise involved here does not invite you to look for further perceptual surprises on the wall; you are probably stunned by what you see and wonder why this is the case. And even if you choose to continue

engaging with red paint and look more closely at it, you do not do so for the sake of engaging with the paint. Your perceptual engagement is *interested* in the sense that you are looking at it for the sake of having more clues about what is going on in your home. Similarly, consider a doctor who is surprised by what she sees in an X-ray image. She might be motivated to keep looking closely at the image. But she does not aim to discover more surprising features from the X-ray. Rather, she is probably forming hypothesis about what is going on with the patient when she saw the surprising feature, and is hoping that the features that she then sees confirms her hypothesis – in other words, she is looking for *unsurprising* features. Also, the doctor’s continual engagement with the X-ray is done for the sake of figuring out what’s wrong with the patient, not for the sake of engaging with the X-ray image *per se*. This last point points at a more general difference between the aesthetic and the non-aesthetic domains. While in the non-aesthetic domains, we typically form hypothesis with the hope that what we perceive confirms the hypothesis, in the aesthetic domain we often search for cues to disconfirm the hypothesis we make about what we are viewing.

To further clarify the SAS, let me make a few remarks about what it does *not* try to achieve. First, the SAS does not try to answer the deeper question of *why* the sub-kind of perceptual surprises that it singles out constitute aesthetic appreciation. But it is compatible with different possible answers. For example, one might claim that the *context* in which the perceptual surprise occurs matters. In typical situations, perceptual surprises tend to signify something alarming – the huge red blob on the wall of one’s home might indicate that a burglary has taken place – and one is drawn to learn more

about the alarming event that the perceptual surprise indicates rather than focusing on the surprising feature themselves. On the contrary, situations in which aesthetic appreciation happens – such as an art museum – provides the kind of environment where the perceptual surprises are not associated with any alarming event, which enables one to keep engaging with the surprising features.

Second, the SAS does not claim that perceptual surprise is involved in *every case* of aesthetic appreciation. For instance, Mark Rothko's paintings, which explore the expressive potential of color contrasts and modulations, may not be aesthetically pleasing because they consist of elements that surprise us. The SAS is an account to explain why we find something aesthetically pleasing in virtue of the surprising perceptible elements that it has.

Finally, while the SAS is concerned with aesthetic appreciation associated with *perceptual* surprises, it is extendable to mediums of art that are more cognitive in nature. Consider reading a novel whose story is easily predictable after reading the first few pages, and there is nothing that surprises you in the execution of the plot. Plausibly, the novel is dull and does not have much aesthetic value. Just like the other mediums of art, an aesthetically valuable novel is often full of surprises – it can be genre-breaking, full of details that surprise the reader, etc.. Deriving from the SAS, it is surprising in virtue of going beyond the categories that the reader attributes to the novel – it might entice the reader to categorize it as a romance story, and feature time-traveling events nicely

weaved into the story – and the surprising content contributes to aesthetic appreciation by keeping the reader search for content that surprises her in new ways.

The above is my characterization of my positive account of what sub-kind of perceptual surprise contribute to aesthetic appreciation, and how it does so. In the rest of this paper, I will further develop and defend it in the light of recent related discussions.

### *2.1 Connections with Neuroaesthetics*

According to the SAS, aesthetic appreciation is a dynamic play between the subject and the object. The object has certain features that surprise the subject given the categories that she attributes to the object; and those surprises motivate the subject to search for more surprising features within the object. Treating aesthetic appreciation as an ongoing process of interactions between the subject and the object fits well with what recent empirical research tells us about aesthetic experience. Empirical studies inform us that aesthetic experience emerges from a complex interplay between the perceptual and the emotional.<sup>31</sup> The question that the scientists study is: “How much of the aesthetic experience resides in a perceptual experience and how much resides in the emotional response to artwork?” (Chatterjee, 2010: 57). Consistent with this, the SAS takes aesthetic experience to be more than just a way of perceiving by emphasizing the role of surprise in some paradigm cases of aesthetic experience.

---

<sup>31</sup> See, for example, Chatterjee and Vartanian (2014) for discussion.

Furthermore, the SAS is able to resolve a puzzle in the neuroscientific studies of aesthetics (a.k.a. neuroaesthetics). Neuroscientists have repeatedly found that the brain's Default Mode Network (DMN) is activated when subjects are having aesthetic experience, such as when viewing high-rated artworks.<sup>32</sup> This result is puzzling, because DMN activation is primarily associated with self-related processes such as thinking about personally concerning events, and DMN is usually *de*-activated when a person is engaged in a task that requires external focus such as visual search. In other words, tasks that involve outward-directed attention are usually associated with de-activation of DMN. Since viewing high-rated artworks involves outward-directed attention at the artworks and their features, it is puzzling why DMN is not de-activated in the aesthetic cases. There is as yet no established interpretation about this finding. The SAS offers an interpretation by showing how aesthetic appreciation involves both outward-directed attention and self-related processes. Better put, it maintains that aesthetic appreciation is an ongoing interaction between these two processes. Outward-directed attention reveals to the subject features that strike her as surprising. The feeling of surprise then motivates the subject to keep attending to the object in the search of more features that surprise her.

My interpretation receives support from recent discoveries about on the DMN, which find correlation between surprise and DMN activation. In Brandman et al. (2021), the experimenters present subjects footages of movies and ask them to rate the level of surprisingness of various events of the movies. They found a positive correlation between

---

<sup>32</sup> See, for example, Belfi et al. (2019) and Chatterjee et al. (2014).



DMN activation and the surprise ratings. This supports my interpretation of the involvement of DMN activation in aesthetic experience as being correlated with the element of surprise.

## *2.2 Connections with Nanay (2016)*

My discussion enriches recent research on the relevance of philosophy of perception to the understanding of aesthetic experience. One prominent example of this research is Bence Nanay's *Aesthetics as Philosophy of Perception* (2016). The goal of Nanay's book is to understand what distinguishes aesthetic experience from non-aesthetic experience. Nanay sees the solution as lying with theories of perception. He proposes that some paradigm cases of aesthetic experience involve a distinctive kind of attention. To introduce this kind of attention, he makes use of the distinction between diffuse/distributed attention and focused attention. When one attends diffusely over a group of items, one does not focus on any of the items in particular. A good example of diffuse attention is when one is sitting on a park bench on a beautiful afternoon, just chilling. While there are lots of items that one is aware of (e.g., children running around, birds, and trees), one is not focusing on any of them in particular. Nanay proposes that diffuse attention and focused attention could be practiced at the same time. His idea is that one can attend to an object (e.g., a painting) in one manner, and to the object's features (e.g., the colors and shapes of a painting) in the other manner. Using this hybrid to explain aesthetic experience, he argues that some paradigm cases of aesthetic experience is crucially a matter of having the attention focused with regards to the object

and distributed with regards to the features. In other words, one is diffusely aware of various features of a painting at the same time while focusing on the painting.

Given the discovery of the involvement of DMN activation in aesthetic experience, Nanay's account, which introduces a style of *outward-directed* attention, cannot be sufficient for explaining aesthetic experience. Nevertheless, this point should not be taken as an objection to Nanay's account, but rather a call for supplement. This is because Nanay does not intend to provide the sufficient condition of aesthetic experience. His goal is just to provide an important element that underlies the difference between aesthetic and non-aesthetic experiences. So, how might the SAS and Nanay's account work together to provide a more complete account of aesthetic experience?

I think that Nanay's account provides a good model as for how to discover surprising features. Since perceptual surprise arises from conflicts between the categories of various features of an object, distributed attention over various features is a good way to be aware of features whose categories conflict with each other, whereas focused attention on the object ensures that the subject is not distracted by the irrelevant features outside the object.

This is not to say that Nanay's attentional strategy is the only way to discover surprising features. Especially in time-sensitive mediums like movies, surprise is often elicited by events that happen across different times. For instance, the subject might have her attention focused on one event after the other, and experiences surprise in virtue of a conflict between categories attributed to these events. Yet, for non-time-sensitive

mediums like paintings, which are a primary concern of Nanay's discussion, the style of attention that Nanay introduces offers an effective strategy to discover surprising elements. For instance, if one experienced the Kandinsky's *Yellow-Red-Blue* by focusing on one feature after the other, one would have to *recall* that the features that she is currently focusing on is contra-standard relative to the features that she previously focused on. This is a less effective way to experience surprise than attending diffusely to many features all at once.

### 2.3 The "Non-perceptual-ness" of Aesthetic Appreciation

The other line of exchange between the philosophy of perception and aesthetics is on whether we can perceptually experience aesthetic properties. This is a branch of the broader discussion of what properties perceptual experience can represent. While it is uncontroversial that our perceptual experience can represent low-level properties, such as shape, size, and color, it is controversial whether our perceptual experience represents any high-level properties. High-level properties are usually defined stipulatively as any properties that go beyond the basic 'building blocks' of perceptual experience such as shape and color (Ransom 2022).

My discussion shows that aesthetic properties are *not* among the things that our perceptual experience can represent. This is because aesthetic qualities, at least in those cases related to perceptual surprises, emerge from the complex interplay between features that we perceptually experience, like shape and color, and our emotional reactions to

those features perceived. So, aesthetics qualities are not, at least in the cases discussed, just ‘built’ from qualities like shape and color.

My treatment of aesthetic qualities is a form of relationalism, the metaphysical view that aesthetic qualities do not reside in the external world, but arise in the relation between us and external features. This consequence may strike some as counter-intuitive. My opponent may say, “We do attribute aesthetic qualities to external objects. For example, it is the flowers that we are looking at that we find beautiful, not the interactions that we have with the flowers. So, my relationalist treatment, which takes the beauty of the flowers to emerge in our interactions with the flowers, is mistaken, or at least an error theory.”

This objection wrongly assumes that being an external quality is required for the quality to be attributed to external objects. To see this, consider the feeling of unfamiliarity. When I’m in New York, I always walk from the 40<sup>th</sup> Street and Broadway to the 29<sup>th</sup> Street and Broadway to get to my favorite coffee shop. After countless times of taking this route, there is finally a day when I decided to walk from the 29<sup>th</sup> back to the 40<sup>th</sup> after the coffee. Taking the opposite direction, the entire street suddenly feels unfamiliar to me. There is a sense of newness to the street which I have been to for so many times. How should we understand this feeling of unfamiliarity? Obviously the feeling of unfamiliarity is attributed to the street. It is the street – neither me nor my interaction with the street – that feels unfamiliar. At the same time, the feeling of unfamiliarity is obviously not a property of the street.

In more detail, we have to make the distinction between *the world-attributing feature* (WA) and *the seemingly world-revealing feature* (SWR). While WA is a characteristic of both aesthetic experience and of the feeling of unfamiliarity, SWR is absent in these experiences. When an experience of x has SWR feature, it seems to reveal x as a feature of the external world. SWR is a characteristic feature of our typical perceptual experience. For instance, P. F. Strawson famously tells us that “mature sensible experience (in general) presents itself as, in Kantian phrase, an immediate consciousness of the existence of things outside us” (1979: 97). Our typical perceptual experience has both features. When I consciously see a cat, my perceptual experience seems to reveal a cat – something in the external world – to me, and the cat is attributed to the external world. There are non-perceptual experiences that lack both features. The feeling of dizziness is an example of this sort. Dizziness makes one feel as though the world is spinning. However, one would not attribute this spinning feature to the world. When being asked, “Is the world spinning?” The answer is obviously no. It does not seem to reveal any property of the external world either. The feeling of dizziness presents itself as a subjective feature, that something is going wrong with the subject herself. Unlike these experiences, aesthetic experience belongs to a different class of experience. It attributes aesthetic qualities to external objects without taking those objects to possess aesthetic qualities.

With this distinction, we are in a position to see how my relationalist treatment of aesthetic qualities is not an error theory. An error theorist maintains that aesthetic qualities are not what they appear to be. If aesthetic experience had SWR, relationalism

about aesthetic qualities would indeed be an error theory. This is because aesthetic experience would seem to reveal aesthetic qualities as features of the external world (due to SWR), whereas relationalism maintains that aesthetic qualities are not external features. Nevertheless, our intuitions only commit us to the view that aesthetic experience has WA, and there is nothing incompatible between attributing aesthetic qualities to external things and taking aesthetic qualities to emerge in the constant search of surprise from the external features perceived.

## **Chapter 4. Attention, Perceptual Judgment, and Perceptual Science**

*Abstract.* In interpreting experimental findings, perceptual scientists often rely on the intuition that: If a subject fails to make accurate perceptual judgment about  $x$ , then the subject does not see  $x$ . This paper argues that once we recognize the role of conscious attention in perceptual judgment formation, this intuition and the interpretations of experimental findings that rely on it become questionable. Focusing on recent debate about unconscious perception, I challenge the popular idea that vision-for-action is unconscious, and argue that unconscious perception (under certain developed definition) is impossible.

## 0. Introduction

Suppose that there are a number of fast-moving objects moving past right in front of you, one after the other. And you are asked to judge the color of each object, say, by speaking out loud the color as each object moves past. Do you see the color? Intuitively, if you accurately judge the color, then you see it; and if you fail, then you don't see it. This pair of intuitions are the basic assumptions in many major experimentation methods of perceptual science. For instance, a classic experimentation method in psychophysics studies the limit of perception of specific types of stimuli through the accuracy and inaccuracy of perceptual judgments. Illustrating with my example above, we can imagine a psychophysics experiment that aims to understand the maximum velocity of moving objects for color perception to be possible. In the actual experiment, the subject would be presented with stimuli moving at speeds that ascend in discrete steps. In each trial, the subject would be presented with a stimulus moving at a specific speed. If the subject accurately judges the color of the stimulus, then a stimulus moving at a higher speed would be presented. The increase in velocity would continue until the subject fails to accurately judge the color of the stimulus. The point where the subject's perceptual judgment is inaccurate is taken to the "sensory threshold", the highest velocity of moving objects where colors can be perceived.

While perceptual judgment plays a crucial role in scientific understanding of perception, little is understood about perceptual judgment – what it is and how it is formed. In this paper, I argue that *conscious attention* is a necessary, but perhaps insufficient, condition for making perceptual judgment. As I will show, recognizing the



role of conscious attention for making perceptual judgment has significant impact on the scientific understanding of perception. It challenges one of the basic assumptions stated above: that if a subject fails to make accurate perceptual judgment about  $x$ , then the subject does not see  $x$ . This in turn calls for re-interpretation of scientific findings that rely on this assumption. To keep the scope of this paper manageable, my discussion will focus on recent debate about unconscious perception. I will show that challenging this assumption prompts us to reconsider the popular idea that vision-for-action is unconscious. Furthermore, I will argue that recognizing the role of conscious attention for making perceptual judgment entails that perception (under certain developed definition) is always conscious.

Here is the plan. In section 1, I begin by developing a definition of perceptual judgment – one that reflects the way perceptual judgment is used in perceptual science. In section 2, I explain the role of conscious attention in making perceptual judgment. In section 3, I explore the implications for perceptual science. A caveat before I begin. Throughout this paper, my discussion will focus on the visual modality. While most of the things that I will say should apply to the other modalities, whether they do is not a concern of this paper.

## **1. What is Perceptual Judgment?**

I define perceptual judgment as follows:

Perceptual judgment =<sub>df</sub> Judgment about perceptible features of the external world that reflects what the judger seems to perceive occurrently or in the very recent past.

This definition is not meant to be controversial; it should reflect both what philosophers and scientists take perceptual judgment to be. Let me clarify this definition bit by bit.

### *1.1 Judgment*

‘Judgment’ is conscious endorsement of proposition: When a subject S judges that P, S consciously commits herself to the truth of P. There are two ways to think about what it takes to consciously endorse a proposition, and thus two ways to think about what it takes to make a judgment. On the *passive* view, consciously endorsing a proposition is a passive endeavor: one automatically accepts a proposition when one consciously entertains it. For instance, in thinking that *it is going to rain tomorrow*, one automatically commits to the truth of the proposition, thereby making a judgment with this propositional content. On the *active* view, conscious endorsement of a proposition requires the act of consciously deciding that the proposition is true. To judge that *it is going to rain tomorrow*, one has to consciously decide that the proposition is true.<sup>33</sup> This paper does not hang on which of the two views is right. Following Mandelbaum (2014), I assume the passive view for the rest of the paper.

### *1.2 Perceptual Judgment*

‘Perceptual judgment’ is conscious endorsement of *perceptual proposition* – proposition about perceptible features of the external world that reflects what the judger seems to perceive occurrently or in the very recent past. An example of perceptual

---

<sup>33</sup> See Mandelbaum (2014) for further discussion and for argument that the passive view is superior.

judgment would be conscious endorsement of the proposition that *there's something red* while looking at that red thing.

There are lots of delicacies surrounding perceptual judgment, especially about what counts as perceptual judgment and what does not. I go over these delicacies below.

First, not every proposition that automatically pops up in the mind when conscious perception occurs qualifies as perceptual proposition. Suppose that I see an old car in Tom's garage. Tom once told me that it is a heritage from his great grandfather. Looking at the car it immediately strikes me that *there's a car that Tom's great grandfather once owned*. While this judgment immediately occurs when I consciously perceive the car, it might not be a perceptual judgment. To be a perceptual judgment, the content of judgment has to be a perceptual proposition; it has to be about features that a subject can perceive. Whether this judgment is perceptual depends on whether the property *being a car that Tom's great grandfather once owned* is a perceivable property. There is an active debate over the limits of perception, about whether complex properties like this can be perceived.<sup>34</sup> I am not going to take a stance on this debate. Below, I will stick with propositions whose status as perceptual propositions are relatively uncontroversial. They include propositions about natural kinds (e.g., that *there's a flower*) and low-level features (e.g., that *this object is red*).

Second, perceptual judgments can concern not only features that a subject is *occurently* perceiving, but also what the subject perceived in the *very recent past*.

---

<sup>34</sup> See, for example, Bayne (2009), Hawley and Macpherson (2011), and Siegel (2006).

Consider the case of seeing a car rush past. Right at the moment when the car goes out of sight, you judge that *there's a car*. This is a perceptual judgment even though you no longer see the car as you make the judgment. This sort of perceptual judgment is often found in vision science. A common experimentation method involves showing subjects briefly presented target stimuli (e.g., a leftward-pointing arrow that is presented for 100 milliseconds), and asking them to judge whether the target stimuli were present or not right after the presentation.<sup>35</sup>

Third, since perceptual judgments are concerned with what the judge *seems* to perceive, perceptual judgments are inaccurate when the judge misperceives. Inaccurate perceptual judgments are as important as accurate perceptual judgments for the scientific understanding of perception. More specifically, perceptual judgments with over fifty percent accuracy (i.e., accuracy above chance) across trials are often taken as an indicator that the kind of entity being studied is perceivable. For instance, suppose that the subject is shown arrows pointing at different directions, one in each trial, and each stimulus is presented for 10 milliseconds. The subject is asked to determine the direction of the arrow in each trial. Suppose further that the accuracy of her judgments is below chance (i.e., the subject misjudges the direction of the arrow in over half of the trials). This is taken to show that the direction of arrows presented at 10 milliseconds is not perceivable.

## **2. Conscious Attention and Perceptual Judgment**

---

<sup>35</sup> There is the further, difficult question of how large the temporal gap between perceiving and judging could be for a judgment to remain a perceptual judgment. This question is not special to this project, and I am not trying to answer it here.

With perceptual judgment clarified, it is time to clarify and motivate the key thesis of this paper, which I call the Judgment-Attention Thesis (JAT for short):

JAT: If a subject makes accurate perceptual judgments about  $x$ , then the subject consciously attends (or has consciously attended in the very recent past) to  $x$ .

To start, let me say a few things about conscious attention. By “conscious attention”, I always mean conscious *perceptual* attention. Arguably, there are non-perceptual forms of attention, when one is attentively doing math in one’s head for instance. Unlike non-perceptual forms of attention, conscious perceptual attention requires the engagement of the senses, such as the eyes.

Scientists have typified conscious attention in various ways. The first pair of distinctions is endogenous and exogenous attention. Endogenous attention is goal-directed and voluntary, whereas exogenous attention is stimulus-driven and involuntary. The second pair of distinctions is overt and covert attention. In the case of vision, overt attention refers to attention that involves moving the eyes to fixate on the unit of attention; covert attention means paying attention to what is in the periphery of the visual field without moving the eyes toward it. In addition to these two pairs of distinctions, there is also discussion about what kinds of entities can be the units of attention. They are usually agreed to be locations in space, objects (e.g., a wallet), and features (e.g., redness).

The JAT is neutral about whether *all* types of conscious attention are relevant to making perceptual judgment. For example, exogenous attention might be too coarse-grained to delineate an object from the rest of the visual scene (Roskies 2010: 125), thus

making perceptual judgment about an object might never entail exogenous attention alone (among other types of attention).

To see the truth of JAT, it would be the most helpful to focus on endogenous, overt attention. Consider a typical case of visual search. Suppose that you are looking for your wallet. You look around at the places where you usually leave your wallet: the top of a cabinet, the bookshelf, your desk, etc.. Finally, you see your wallet on your desk. Your search is thereby complete. Your search is made possible by making various perceptual judgments, such as the judgment *that there're just some books* when you are looking at the bookshelf, and the judgment *that the wallet is over there* which marks the end of your visual search. Making these perceptual judgments is always accompanied by conscious attention. It is hard to see how making these perceptual judgments is possible without consciously attending to those items which you make perceptual judgments about.

To further argue for JAT, consider also the following case:

*Yoga Sue.* Sue is practicing yoga in the front row of the room. There is some distance in between the front edge of her mat and the white wall. As Sue is doing the tree pose, she does not pay any attention to what is in front of her. What falls in the center of her visual field includes the front edge of her mat, part of the white wall and the floor in between the mat and the wall. Deeply absorbed in her practice, Sue does not make any judgment about the distance between the mat and the wall; not until the yoga teacher has her hold the pose for too long and eventually Sue starts feeling annoyed and tries to distract herself with random thoughts like 'How far is the mat away from the wall?'

The case above involves a contrast between two situations: at an earlier time  $t_1$ , certain worldly items are placed right in front of a subject, but the subject does not make perceptual judgments about them. At a later time  $t_2$ , the same items are placed at the same location relative to the visual field of the subject, and the subject makes perceptual judgments about them. What marks the difference between  $t_1$  and  $t_2$  is that the subject attends to those items at  $t_2$  but not at  $t_1$ . When Sue was absorbed in her practice of yoga, she did not pay attention to the distance between the mat and the wall, and did not make perceptual judgment about it. But once she pays attention to the distance, she makes perceptual judgment about it. This suggests that conscious attention is essential for making perceptual judgments.

This phenomenon that I bring up with Yoga Sue is common. When we are absorbed in our thinking or distracted by something else, oftentimes we simply do not make perceptual judgments about what is right in front of our eyes. When I am completely absorbed in my philosophical thinking, if you ask me what was in front of me right at the moment when my eyes move to a different direction without letting me look back, I have no idea about what the answer is. This again suggests that we have to consciously attend to things in order to make perceptual judgments about them.

What JAT shows is that conscious attention directed at  $x$  is an *empirically necessary condition* for making perceptual judgments about  $x$ . It does not speak to any deeper relation between conscious attention and making perceptual judgments, such as whether conscious attention is *metaphysically necessary for* or *constitutes* making perceptual judgments. Figuring out the deeper relation requires empirical understanding of the

functional role of conscious attention that we do not currently have. Also, we can easily imagine super-blindsighters who can make judgments about perceptible features of the external world and yet cannot consciously attend to things. Arguing for the metaphysical necessity of conscious attention for making perceptual judgments might require dismissing super-blindsighters as metaphysically impossible. So, rather than arguing either for or against the metaphysical necessity of conscious attention for making perceptual judgments, I consider a key factor that is at play in settling this issue.

To make perceptual judgment, it is metaphysically necessary that the content of perceptual experience is available in propositional format. So, conscious attention would be metaphysically necessary for making perceptual judgments if its functional role includes propositionalizing the content of perceptual experience. Whether conscious attention has this function is an unsettled question.<sup>36</sup> An inconclusive reason for thinking that conscious attention has the function of propositionalization is to rely on the intuition brought up in my situational contrast argument above, that making perceptual judgments goes hand-in-hand with conscious attention. But this reason is inconclusive, because the intuition is compatible with the possibility that propositionalization is a separate process that takes place involuntarily and unconsciously at the same time when conscious attention takes place. An inconclusive reason for thinking that conscious attention does *not* have the function in question is that many non-human animals can also consciously attend, even though they may not be able to entertain propositions. But this reason is also

---

<sup>36</sup> See Roskies (2010) and Watzl (2018) for discussion.



inconclusive, because what mechanisms underlie conscious attention is still an open empirical question, and it is possible that conscious attention in human beings operates differently and serves different functional roles from conscious attention in non-linguistic animals.

Given that much is unknown about the role that conscious attention plays in cognitive processes,<sup>37</sup> we are not in a position to determine whether conscious attention is *sufficient* – both *metaphysically* and *empirically* – for making perceptual judgments. Being sufficient would require, at the very least, that conscious attention has the function of propositionalizing the content of perceptual experience. For otherwise, making perceptual judgments would require a separate process of propositionalization. But as we've seen, whether conscious attention has this function is unsettled. So, JAT must not be confused with the Failed-Judgment-Failed-Attention Thesis (FJFAT for short), which we are uncertain about:

FJFAT: If a subject fails to make accurate perceptual judgments about  $x$ , then the subject does not consciously attend to  $x$ .

The truth value of FJFAT is undetermined because conscious attention might not suffice for making accurate perceptual judgments. Thus, there might be cases where the subject consciously attends to something but still fails to make accurate perceptual judgments about it.

---

<sup>37</sup> See Wu (2011) on the significance of attention for the selection of items to act on.

In the next section, I show how recognizing that JAT is true and that FJFAT is questionable influences the shape of the debate about unconscious perception.

### **3. Implications for the Debate about Unconscious Perception**

#### *3.1 Vision-for-Action May Not be Unconscious*

DF is a patient who developed visual form agnosia following carbon monoxide poisoning. She showed no problem in reaching for or grasping objects. DF was able to insert a card into the slot with a particular orientation just like normal individuals. But DF failed to make accurate explicit reports about the orientation of the slot: She was nearly at chance in saying what the orientation of the slot was and in turning the card (without inserting it) to the same orientation as the slot, and emphasized that she did not see the slot's orientation.

The DF case has been influentially taken by Milner and Goodale (1995, 2006) to show that vision-for-action is unconscious. Their logic is that if DF's explicit reports about the slot's orientation are wildly inaccurate, then she must not consciously perceive the slot's orientation. Yet, DF successfully engages in vision-guided action directed at the slot's orientation. So, vision-for-action is unconscious.

My discussion puts their logic in question. Their argument is based on the following assumptions: (1) that DF's inaccurate explicit reports about orientations indicate inaccurate perceptual judgments about orientations; and (2) that if DF's perceptual judgments about orientations are inaccurate, then DF does not consciously perceive orientations. I will grant Milner and Goodale (1). My discussion shows that (2) is questionable. Here's my argument. Since FJFAT is questionable, DF might be able to

consciously attend to the slot's orientation even though her perceptual judgments about the slot's orientation were inaccurate. If DF was indeed able to consciously attend to the slot's orientation, then DF was able to consciously see the slot's orientation. It follows from the consequent that the DF case does not provide evidence for vision-for-action being unconscious.

To put in more speculative detail of my interpretation, DF's brain damage might have cost her the ability to propositionalize perceptual content about orientations. She might still be able to consciously see orientations, albeit the perceptual content about orientations was merely available in non-propositional format. Her conscious perception of orientations might be similar to non-linguistic animals that can consciously attend to things but cannot entertain propositions about them. Under this interpretation, DF would indeed fail to make accurate perceptual judgments about orientations, even though she could consciously see orientations.

Furthermore, there is good independent reason for thinking that DF was capable of consciously perceiving orientations. Inserting a card to the slot is an *intentional action directed at an external target*. As an *intentional action*, its performance is triggered and guided by an intention – in this case, the intention of inserting the card to the slot with a specific orientation. As an action *directed at an external target*, it culminates with some sort of interaction with specific external entities. In this case, the action culminates with the card's being in the slot. Another example of intentional action directed at an external target would be shooting an arrow to the bullseye. To clarify, not all intentional actions are directed at an external target, and not all external-target-directed acts are intentional.

An example of the former would be doing ten push-ups in a row. An example of the latter would be ducking to avoid a projectile.

It is uncontroversial that performing an external-target-directed intentional action requires the subject to consciously focus on the external target at the beginning of the performance. Just think about what it would be like for you to perform the task the DF was asked to do. At the beginning, you consciously focus on the orientation of the slot (i.e., the target). Then, you turn the card to the right orientation, and insert it to the slot. Recently, some have argued that conscious perception is not necessary for the fine-tuning of action during the performance.<sup>38</sup> Metaphorically put, it is the “automatic pilot” that is in charge of modifying the hand movements. But accepting that fine-grained visuomotor control is unconscious does not affect the point that I am making. The automatic pilot hypothesis is a hypothesis about *action execution*. *Target selection* takes place *before* action execution. No reasonable person would deny that target selection requires conscious perception,<sup>39</sup> and accepting this uncontroversial point suffices for showing that DF consciously saw the orientation of the slot.

My interpretation above does not address the deeper question of what the cognitive value of perceptual judgments is. For instance, if we can consciously attend to things and perform actions directed at them without the mediation of perceptual judgment, why do we need to make perceptual judgments? I suggest that perceptual judgments are

---

<sup>38</sup> For discussion, see for example Fridland (2017) and Montero (2016).

<sup>39</sup> For instance, even Ian Phillips who accepts the automatic pilot hypothesis claims that “the only representations attributable to the individual are those associated with target and action-type selection, and which are associated with consciousness” (2018: 501).

significant for sophisticated forms of agency, such as deliberately weighing the pros and cons before pursuing an action. Without the ability to make perceptual judgments, we cannot ask questions such as: Is this thing really what it looks like? This thing looks poisonous; should I eat it or not? I will say a bit more at the end of this paper.

### *3.2 Ruling out Unconscious Perception (under Some Definition) Conceptually*

Peters and Lau (2015, 2017) recently argue that there is no unconscious perception, by putting the following operational definition of ‘perception’ to empirical test:

To demonstrate that perception has indeed occurred in a perceptual decision-making task, an observer must be able to make some identification or direct discrimination decision about a stimulus better than just guessing. (2017, p. 2)

In the experiment, each trial consisted of two intervals. Only one of the two intervals contained a masked Gabor patch that was either left-tilted or right-tilted.<sup>40</sup> Subjects were asked to make decisions about the orientation of the patch in both intervals (even though there was a Gabor patch only in one), and to bet on which interval they thought that their decisions were accurate. Their idea was that if there was unconscious perception, then subjects should have made some accurate discriminatory decisions without making the right bet. But they found no evidence for unconscious perception: Whenever the subjects discriminated the orientation above chance, they made the right bet.

I agree with Peters and Lau that unconscious perception does not exist under their definition of perception. But the problem is that, once we recognize the truth of JAT,

---

<sup>40</sup> Masking is a technique that is used to prevent the subjects’ judgments to be due to perceptual memory rather than occurrent perception.

their definition of perception already implies that perception is conscious; that perception is conscious is a truth that does not await empirical confirmation. Here's why: Insofar as a subject can make an accurate perceptual judgment about a stimulus, with JAT, she consciously attends to the stimulus. Consciously attending to a stimulus implies consciously perceiving the stimulus. Therefore, perception is conscious under their definition of perception.<sup>41</sup>

#### **4. Conclusion**

Throughout this paper, I developed and defended the view that conscious perceptual attention is at least empirically necessary for making perceptual judgments. Focusing on debate about unconscious perception, I explored its implications for perceptual science. To end, I point out two other areas that my discussion might have consequences for as topics for future investigation.

##### *4.1 Implications on the condition for conscious perception*

What is the condition for consciously perceiving a worldly item  $x$ ? Some maintain that a neurotypical, mature human subject  $S$  consciously perceives  $x$  only if  $S$  can make

---

<sup>41</sup> As an objection, one might wonder: What about Ned Block's (1995) imaginary case of superblindsighters? Superblindsighters cannot consciously attend to things, and yet they are able to make correct 'guesses' about what is in the blind field. There are at least two things I can say with regard to superblindsight. First, conceivability of superblindsighters might not imply metaphysical possibility. Second, it is questionable whether those guesses qualify as *perceptual judgments*. Since visual information from the blind field simply pops into thoughts (p. 233), it is legitimate to question if superblindsighters indeed *perceive* the items in the blind field. (See Phillips 2018 for discussion on conditions for perception.). If they didn't, their guesses wouldn't be perceptual judgments, and so superblindsight would be irrelevant to the theme of this paper.

perceptual judgments about x. I call this view *judgmentalism*, and those who endorse this view *judgmentalists*.<sup>42</sup> Here's a representative quote from Tye (2014):

If one is [perceptually] aware (de re) of some entity, one's awareness directly puts one in a position/enables one to form de re cognitive attitudes with respect to that entity. After all, if one cannot even ask "What's that?" with respect to some entity directly on the basis of one's awareness, surely one isn't aware of that entity. (p. 44, emphasis in original)

Cases that have been used to motivate judgmentalism involve conscious attention. Here's one of them. Suppose that there is a moth perfectly camouflaged on a tree trunk. Even if one is focusing on the part of the tree trunk where the moth is, intuitively one does not consciously see the moth. Why? 'Because one's awareness would not have put one directly in a position even to ask "what it that?" with respect to it' (p. 44-5).<sup>43</sup>

My discussion vindicates judgmentalism when it comes to *attentive* conscious perception. But it also raises the question of whether it is possible to consciously perceive without conscious attention. Think about Yoga Sue again. When Sue's focus was dominated by the tree pose, she didn't pay attention to any worldly items that her visual field covers. Did she consciously *see* at least some of them? The answer seems to be yes. Even when Sue was focusing on her practice of yoga, she did not seem to be drained of

---

<sup>42</sup> Judgmentalism is different from the position according to which S consciously perceives x insofar as S *perceptually differentiates* x from its immediate environment (Dretske 1969). Being in a position to make perceptual judgments about x presupposes perceptual differentiation of x. But perceptual differentiation can take place without getting in a position to make perceptual judgments about the object. Very few kinds of animals can get in a position to make perceptual judgments, whereas much more kinds can engage in perceptual differentiation.

<sup>43</sup> I also attribute judgmentalism to at least McDowell (1994) and Montague (2013).

visual phenomenology. The presence of visual phenomenology might be a good enough indicator that conscious perception occurs.<sup>44</sup>

#### *4.2 The Cognitive Significance of Perceptual Judgments*

The other topic for future exploration is the significance of perceptual judgments for our cognitive lives: What are perceptual judgments *for*? I pinpoint two specific areas of interest.

##### *4.2.1 Perceptual Learning*

Perceptual judgments seem to be significant for perceptual learning – the development of perceptual expertise through practice.<sup>45</sup> Compare two individuals listening to Beethoven’s Symphony No.9. One of them is a connoisseur of classical music, whereas the other is a novice. The connoisseur is able to differentiate many more subtle changes of timbre, dynamics, etc. throughout the piece, and much more capable of unitizing the overall feeling that this piece conveys. The connoisseur’s expertise is acquired through practice. Practice is not just repeated exposure to the music, being passively struck by the sound of the symphony over and over again. It also involves actively attending to different details in the symphony – focusing sometimes on the timbre, the violin, and sometimes more on the overall ambient of the music.<sup>46</sup> As the subject focuses on these different elements of the symphony, she makes perceptual judgments about them;

---

<sup>44</sup> Dretske (1979) has referred to a similar phenomenon as *simple seeing*.

<sup>45</sup> Not every instance of perceptual development is an instance of perceptual learning. For instance, Kellman & Garrigan (2009) point out that the development of infants’ abilities to perceive objects, faces, three-dimensional space, etc. is ascribable primarily to innate or early-maturing mechanisms. For an exploration of this topic, see Connolly (2017, 2019).

<sup>46</sup> See Matthen (2015) for more on the effortful development of perceptual skills for aesthetic appreciation.



sometimes she might even make judgments about how a specific part of the symphony sounds in different versions of the symphony. This suggests that perceptual judgments are an integral part of the development of perceptual skills.

#### *4.2.2 Sophisticated Agency*

Perceptual judgments seem to enable sophisticated forms of agency. Suppose that there is a dish of French fries in front of me. I'm a French fries lover, but I'm not sure if I should eat it. I weigh the pros and cons, measuring the predicted satisfaction that I will gain after eating it against the high calories intake that comes with eating it. Finally, I arrive at an all-things-considered decision about whether to eat it or not, and act it out. Complex action-decision planning processes like this seem possible only after I have made perceptual judgments – committing to the truth of the propositions – such as that *there's a dish of French fries*. Perceptual judgment formation puts perception into the same format as non-perceptual thoughts, thereby enabling measuring perceptual intake against other considerations.

## Chapter 5. Looking, Responsibility, and Rationality

*Abstract.* Negative traits can make us worse perceivers; and when they do, we are apt to receive negative normative appraisals. But why? In *The Rationality of Perception*, Susanna Siegel presented a radical answer: Irrationality is transmitted by negative traits to perceptual experiences; and perceivers are irrational in virtue of having irrational perceptual experiences. In this paper, I present a non-radical answer: Negative traits can make us look carelessly, whereas carelessness is an epistemic vice. Depending on the consequences that the careless looking results in, we can be subjected to other negative normative appraisals, such as being blameworthy and morally responsible. My answer is non-radical because, as I argue, looking is an intentional action, and it is uncontroversial that we are apt to receive normative appraisals from intentional actions. I argue that my answer has significant advantages over Siegel's.

## 0. Introduction

As perceivers, we engage in intentional perceptual activities from time to time. Just like any other intentional actions, we are apt to receive normative appraisals from them. Take the case of voyeurism. When Peeping Tom looks through the window into someone else's home and intently watches the occupants doing a private act, he is morally blameworthy for his perceptual activities. In many parts of the world, this sort of behavior is a criminal offence, and the subject who engages in it is subjected to imprisonment and fines. Intentional perceptual activities are ubiquitous. In the visual modality, which is my focus in this paper, they are described by verbs such as looking for, watching, observing, staring at, scrutinizing, and inspecting. For simplicity's sake, I will use the generic term "looking" to refer to this category of perceptual activities – visual perceptual acts that are intentional.

Looking presents a contrast to *perceptual experience*, a category that is at the center of recent philosophy of perception literature. According to the traditional view of the normativity of perceptual experience, the subject of experience is not apt to receive normative appraisals from her perceptual experiences. For instance, having the perceptual experience of an immoral scene does not affect my moral standing. Presumably, the rationale behind the traditional view has to do with the common way to evaluate perceptual experiences. Perceptual experiences are often evaluated in a way similar to evaluating thermometers. Thermometers are evaluated in terms of whether they *accurately* reflect the temperature of the surrounding environment. Whether accurate or

not, the evaluation does not redound on the normative standing of the subject who owns the thermometer.

Siegel (2017) challenges the traditional view, by arguing that the traditional view fails to account for certain normative phenomena. In light of those cases, she argues that perceptual experiences can be evaluated in terms of *rationality*; and perceivers can be *irrational* in virtue of having irrational perceptual experiences.

In this paper, I argue that Siegel's challenge to the traditional view is undermotivated. The normative phenomena that Siegel uses in her favor can be better explained in terms of looking. Furthermore, I argue that (ir-)rationality might not be the appropriate normative standard to use to evaluate perceivers.

Here is the plan. I will begin by saying more about looking, and explore ways in which looking might be related to perceptual experiences (section 1). Then, I will argue that perceivers are apt to receive normative appraisals from looking by turning to two major theories of moral responsibility. I will show that we bear moral responsibility for looking according to both theories (section 2). Finally, I will address Siegel (2017), arguing that there's no need to expand subject-level normative evaluations to perceptual experiences (section 3).

## **1. Looking, Seeing, and Perceptual Experience**

What does it mean, fundamentally, to engage in intentional visual activities – watching, staring at, scrutinizing, etc. – which I generically call *looking*? The

fundamental feature that is shared across variants of looking seem to be that they all involve intentional maintenance of conscious perceptual attention on what is looked at. Staring at a red cube involves intentionally maintaining conscious visual attention on the cube. Some variants of looking are more complex, but the intentional maintenance of conscious visual attention seems crucial. Consider looking for the car key. It usually involves consciously attending to multiple objects, and what objects the subject maintains conscious perceptual attention on and for how long are determined by whether those objects have any key-like features. If the subject looks at the top of the coffee table and fails to see any key-like objects (or sees the absence of key-like objects), she will immediately look elsewhere to where she thinks the car key is likely to be. If she sees some key-like objects upon looking at her office desk, her conscious visual attention will linger longer there, until she can decide whether her car key is there or not.

As *intentional* visual activities, the various perceptual activities described above are performed *at will*. For instance, looking for the car key is something that the subject performs out of (say) her intention to use her car. And she is free to stop engaging in this perceptual activity at any point.

The relation between looking and seeing is a matter of debate. Consider Matthen's (2019) definition of seeing as "veridical experience of things outside the perceiver brought about by looking". According to this definition, looking is essential for seeing in virtue of being in a causal relation with veridical perceptual experiences. While Matthen

does not offer any detailed characterization of looking,<sup>47</sup> his definition of seeing seems problematic. Consider the case of a bright flash of light suddenly appearing in front of you. Because of its salience, the light catches your attention. This case is naturally described as *seeing without looking at the light*. My characterization of looking fits with this natural description: Looking does not take place because seeing the light does not involve any intentional perceptual activities. Nevertheless, there is something right with Matthen's definition. I take it to be uncontroversial, or perhaps even an *a priori* truth, that looking is an *aspect* or *component* of seeing *when seeing involves looking*. If I look at you and see you, my looking at you is an aspect or a component of seeing you. I see you in virtue of looking at you, but not vice versa.

Matthen's definition of seeing also touches on the relation between looking and veridical perceptual experiences. Matthen adopts the *means-end view* about the relation between looking and veridical perceptual experiences: looking is the means to achieve certain particular sorts of veridical perceptual experience as its end. The relation between looking and certain particular sorts of veridical perceptual experience might be like cooking and the dish cooked; cooking is the means to bring about certain particular sorts of dish. But there is at least one alternative view which also seems plausible. According to the *involvement view*, perceptual experiences are part of looking. Looking cannot be reduced to perceptual experiences because looking involves more. On this view, looking and perceptual experiences might be like inhaling and taking in the particles in the air.

---

<sup>47</sup> He only describes it in sparse terms, such as "looking is something that we do, and of which we can, therefore, be self-aware".

Taking in the particles in the air is part of inhaling. While inhaling is an activity that we can do at will, not every part of it is up to us. The taking in of the particles depends on what's in the air, which is not up to us.

These two views have different consequences. According to the means-end view, one could look without having any perceptual experiences. Since perceptual experiences emerge later in the causal chain, something – perhaps a sudden failure of the visual system – could happen in between looking and the generation of perceptual experiences to preclude the perceptual experiences from being generated. On the contrary, the involvement view maintains that this scenario is impossible. Since perceptual experiences are part of looking, it is impossible to look without having perceptual experiences (at least in the fullest, non-defective, form of looking). Adjudicating between the two views is a complex and delicate issue. It is beyond the scope of this paper to pursue this task. I leave this issue with the following provisional story: Seeing might not always involve looking. But when it does, looking is an aspect or component of seeing. Looking either is a means to achieve certain particular sorts of veridical perceptual experience as an end, or involves veridical perceptual experiences as its part.

## **2. Moral Responsibility of Looking**

Once we have arrived at the category of looking as something that we perform at will, it easily follows that the subject who looks can be subjected to normative appraisals. In this section, I establish this point by focusing on one kind of normative appraisal – moral

responsibility. I show that we bear moral responsibility for looking according to two major theories of moral responsibility.

According to the control-based view, we are only morally responsible for acts that we perform out of our own choices, acts that we have control over whether to do them or not. If I punch you in the face because I am mad at you, I bear moral responsibility for what I do. But if I do so out of a seizure, I am not morally responsible. Why? Because in the former, punching you is a choice that I have made, whereas in the latter, punching you is not within my control.

According to the self-expression view, we are morally responsible only for acts that reflect the quality of our will: what we care about, our values, moral outlook, etc.. With this view, I am morally responsible for intentionally punching you because it expresses certain bad qualities of my will – perhaps my brutality, lack of temperance or a lack of care for you.

Looking is something that we bear moral responsibility for according to both views. Being an activity that is performed at will, looking is done out of one's own choice, and reflects the quality of one's will. Consider Peeping Tom again. Watching the occupants is something that he has control over; he could have easily looked away. And engaging in such profane perceptual activities reveals his character as a pervert, an invader of others' privacy.

### **3. Explaining the Normative Phenomena in Terms of Looking**



In the above, I argued that we are apt to receive normative appraisals for looking. In this section, on the basis of this I derive straightforward explanations of some normative phenomena that have received philosophical attention recently.

Siegel (2017) points out a certain sort of inappropriateness in our perceptual lives. Here is one of her favorite examples:

**Vain Vivek.** Vivek is a vain performer. When he looks at the audiences, their faces appear more pleased than they actually are. Vivek's vanity influences how pleased they appear to him.

According to Siegel, subjects like Vivek deserve negative normative appraisals for their perceptual experiences. But she claims that normative appraisals like *responsibility* and *blame* do not apply. The appropriate standard of normative evaluation is *rationality*: Vivek's perceptual experiences are *irrational*. And given that the perceptual experiences are irrational, the perceptual beliefs formed on their basis (e.g., that the audience is pleased by the performance) are irrational, and thereby unjustified. In the same way, if I believe that the world is going to end soon on the basis of irrationally associating consecutive cloudy days with the end of the world, my belief is irrational and unjustified.

As a number of notable philosophers<sup>48</sup> have pointed out (and as Siegel herself recognizes), Siegel's proposal is radical. It conflicts with traditional approaches to the epistemic role of perceptual experience, according to which perceptual experiences are *a-*

---

<sup>48</sup> See Lord (2020) and Peacocke (2018).

rational, *unjustified justifiers*.<sup>49</sup> How could perceptual experiences be irrational? To further elucidate her approach, Siegel claims that there is an inferential process that connects psychological precursors (e.g., vanity) to perceptual experiences. And the irrationality of those precursors can be passed down to the perceptual experiences. So, Vivek's perceptual experiences are irrational because his vanity is irrational, and – in Siegel's terms – it *hijacks* Vivek's perceptual experiences; his vanity downgrades the epistemic power of his perceptual experiences by inferentially transmitting its irrationality to the perceptual experiences.

I am on board with Siegel that there's something inappropriate about Vivek *qua* *perceiver*, but Siegel's development rests on many highly controversial ideas. They include: that irrationality can be transmitted from the background outlook of the world to perceptual experiences, that there can be genuine inferences from background views to perceptual experiences, and that perceptual experiences can be irrational.

My discussion on looking, on the other hand, offers a non-radical way to establish the inappropriateness of Vivek's perceptual life. Once we recognize that our perceptual lives involve intentional perceptual activities, we can attribute the inappropriateness of Vivek's perceptual life to these intentional activities, which are indisputably objects of normative evaluation. How exactly does that work? Here's my proposal. Vivek's problem fundamentally lies in his *careless looking*, which originates from his vanity. What makes

---

<sup>49</sup> Traditional approaches need not deny that Vivek's perceptual beliefs are unjustified. For instance, they can maintain that those beliefs are unjustified because Vivek's perceptual experiences are unreliable. Accepting this does not imply that the perceptual experiences are irrational.

his looking *careless*? With my earlier discussion about looking, this can be understood in terms of Vivek's not focusing on the right features of the face, and not maintaining the conscious perceptual attention on the faces long enough.

My proposal does not deny that Vivek's vanity influences his perceptual experiences. It can maintain that Vivek's vanity affects both aspects of perception: the way he looks (which is within his control), and his perceptual experiences. But *contra*-Siegel, my proposal attributes the negative normative appraisals to the former. Contrast Vivek with an uncomplacent performer who has an objective grasp about the standard of her artistry. There might as well be moments that her visual experiences are the same as Vivek. Visual systems are noisy, and perhaps her desire to deliver a wonderful performance somehow affects her visual experiences. But being uncomplacent, she would pause and doubt, "Wait. Really?" She would take a more careful look at the audiences and find out that her previous visual experiences are misleading. It should be obvious that there isn't any point in time that the uncomplacent performer deserves any negative normative appraisals. In other words, the fact that the audiences appear more pleased than they actually are does not in itself make Vivek subjected to negative normative appraisals. He is negatively appraised because he does not look at the audiences more carefully when having such experiences.

So, the epistemic fault of Vivek qua perceiver isn't necessarily irrationality. Rather, it is his carelessness. Carelessness – together with traits like prejudice and closed-mindedness – is an epistemic vice (Cassam 2016). On my account, this is because being

careless precludes subjects from appropriately collecting evidence that is readily available to her. Being a careless looker, Vivek fails to appropriately collect the perceptual evidence that is readily available to him (i.e., the audiences' displeased faces right in front of him). Similarly, consider a careless newspaper reader. She's trying to learn about the news. But she reads the newspaper carelessly, skipping words here and there, and ending up with a highly inaccurate picture about the way the world is as reported by the newspaper. Her carelessness is an epistemic fault because even though evidence about what's going on in the world is readily available to her in the newspaper that she's reading, she does not collect the evidence appropriately.

Careless looking can make the subject apt for other negative normative appraisals depending on the consequences that the careless looking results in. Consider the following two cases.

**Student Sally.** Sally is a college student. She got to the course website to check when the paper is due. The webpage clearly states that the due date is April 4<sup>th</sup>. But due to her carelessness, she misread it as April 14<sup>th</sup>. Sally ended up missing the deadline. When Sally tells her instructor why she missed the deadline, it is reasonable that the instructor does not excuse her for missing the deadline.

The reason why it is reasonable not to excuse Sally for missing the deadline is that she deserves *epistemic blame*.<sup>50</sup> She bears *responsibility* for not looking at the website carefully enough.

**Driver Tim.** Tim is a negligent driver. He is not paying enough attention to the condition of the road, and hits some pedestrians.

Tim bears *moral responsibility* for not looking at the road as carefully as he should, since his careless looking results in the terrible consequence of injuring the pedestrians. The cases above show that, *contra*-Siegel, there isn't a unique sort of normative standard that is attributable to perceivers.

Note that I'm not saying that every negative normative appraisal that we attribute to perceivers is rooted in looking *carelessly*. Think about Peeping Tom again. He is morally responsible for his looking, but not because he is looking carelessly. On the contrary, the problem is that he is looking too carefully when he should not.

My proposal has various advantages over Siegel's. First, it attributes the normative inappropriateness of one's perceptual life to intentional actions, which are uncontroversially objects of normative evaluation. Second, it does not rely on the puzzling idea that our perceptual lives can be irrational. The normative appraisals that my proposal attributes to perceivers – such as carelessness, and finding a perceiver blameworthy for looking carelessly – are commonsensical. Third, my proposal does not

---

<sup>50</sup> See Brown (2020) for discussion of epistemic blame. But she focuses on beliefs.

involve the mysterious idea of passing down an epistemic fault (e.g., irrationality) from a psychological precursor (e.g., vanity) to a perceptual state. Vivek is epistemically at fault because he is looking carelessly. He is looking carelessly because he is vain, but his vanity itself does not constitute an epistemic fault.<sup>51</sup> Fourth, my proposal offers a commonsensical diagnosis of what's wrong with Vivek. Consider the following plausible follow-up scenario of Vain Vivek.

**Vivek's Realistic Boss.** After finishing yet another mediocre performance, Vivek gets off stage. There, Vivek sees his boss. Vivek greets his boss happily, telling his boss, "Did you see how pleased the audiences are?" The boss, pissed off by how Vivek is reacting to the audiences' tepid reaction, says to Vivek, "Did you even *look at* the audiences? No one is pleased!"

The blame is on the looking.<sup>52</sup> On the contrary, asking Vivek, "Did you even have the *perceptual experience* of the displeased audiences?" seems odd as a way to blame Vivek.

One desideratum for Siegel is to distinguish cases like Vain Vivek from illusions like the Muller-Lyer. According to Siegel, the illusory nature of the Muller-Lyer experience does not downgrade the epistemic power of the experience; purely on the basis of the

---

<sup>51</sup> It is uncontroversial that one's actions can *originate* from one's outlook of the world. For instance, my treating you well originates from my love for you; Sam's action of stealing unnecessary luxuries originates from his greed. How exactly we should understand the originating-relation is a matter of debate. Possible options include the causal relation (e.g., my love causes me to treat you well), manifestation of a disposition (e.g., my love manifests in my actions of treating you well), and parthood (e.g., my treating you well is part of my love for you). My point does not hinge on taking a stance on this debate.

<sup>52</sup> One might point out that Vivek's boss complaint might as well be made by asking Vivek, "Did you even *see* the audience?" This would not be an objection to me, since seeing is a broader activity of which looking is a component or aspect.

experience (i.e., setting aside the knowledge one might have about the illusion), it is rational to believe that the two lines are not of the same lengths. Assuming that Siegel is right, my proposal can also pinpoint the difference between Vain Vivek and the Muller-Lyer illusion. The reason why a perceiver is not apt to receive any negative normative appraisals when subjected to the Muller-Lyer illusion is that she does not make any mistake *qua perceiver*. She has done nothing wrong in her life as a perceiver that shapes her into receiving the illusion. The Muller-Lyer illusion can be fully explained by how our visual system works without pointing to the subject's character, interest, care, etc.. The problem is not with one's not living one's life as a perceiver well enough. No matter how hard one looks at the two lines, one is still going to suffer from the illusion. Since the perceptual error is just a matter of how the sub-system works, it does not redound on the subject's normative standing.

To end, I consider an objection on behalf of Siegel. Siegel might argue that I have gotten her Vivek case wrong. She might argue that Vivek's perceptual error is not to be resolved by looking more carefully. No matter how hard Vivek looks, the audience would still appear pleased even though they are not. And so, the problem of Vivek cannot be located at how well he is looking at the audience.

In response, this characterization of Vain Vivek is either unrealistic or irrelevant. Let me start with the former. This characterization of the case literally says that even if Vivek is standing right in front of an audience member who shows obvious signs of displeasure – frowning, pressed-together lips, nose scrunch – and paying close attention to these

details of the audience member's face, the face would still appear pleasant to him. I just can't see how this is possible. Perceptual experiences are responsive to the sensory input. While psychological precursors can influence perceptual experiences – for instance, a banana-shaped object might appear yellower than it actually is (Hansen et al. 2006) – the degree to which the psychological precursors can alter experiences is very limited. In other words, the psychological precursors cannot affect perceptual experiences in a way that make those experiences *drastically* conflict with the sensory input. Insofar as Vivek is paying close attention to the features of the faces that indicate displeasure, he cannot have a perceptual experience that indicates the complete opposite. Alternatively, Vivek does not pay close attention to those features; it would be possible for the faces to appear pleased to Vivek. But then, the problem would again be Vivek not looking carefully enough.

Finally, let's move on with my charge of irrelevance. That the psychological precursors block off the sensitivity to the incoming evidence seems possible at the level of *beliefs*. Consider someone who is extremely recalcitrant. No matter how much evidence against her opinion is presented to her, she sets the evidence aside and insists on her original opinion. Entertaining this possibility, Vivek's psychological precursors bias his belief formation: In forming beliefs about audience's reaction to his performance, too little weight is given to the perceptual experience of displeased faces. This interpretation, although plausible, defeats Siegel's original purpose, which is to pinpoint something wrong at the level of perception. Since the mistake is at the level of belief formation, Siegel could no longer say that Vivek's perception is irrational.



## References

- Armstrong, D. M. (1968) *A Materialist Theory of the Mind*. London: Routledge.
- Baras, D. and Na'aman, O. (2021) "What Makes Something Surprising?" *Philosophy and Phenomenological Research*, 00:1-21.
- Bayne, T. (2009) 'Perception and the Reach of Phenomenal Content,' *The Philosophical Quarterly*, 59(236): 395-404.
- Beck, J. (2021) "Mundane hallucinations and new wave relationalism," *Nous*, 1-23.
- Belfi, A., Vessel, E. Briemann, A., Isik, A., Chatterjee, A., Leder, H., Pelli, D., and Starr, G. (2019) "Dynamics of aesthetic experience are reflected in the default-mode network," *Neuroimage* 188: 584–597.
- Block, N. (1995) 'On a Confusion about a Function of Consciousness,' *Behavioral and Brain Sciences* 18: 227-287.
- Block, N. (2005) "Review of Alva Noë, *Action in Perception*." *The Journal of Philosophy*, 102, no. 3: 259-72.
- Bodhi, B. (2011) "What does mindfulness really mean? A canonical perspective," *Contemporary Buddhism*, 12:1, 19-39.
- Boghossian, P. (1989) "Content and self-knowledge," *Philosophical Topics*, 17: 5–26.
- Brandman, T., Malach, R., and Simony, E. (2021) "The surprising role of the default mode network in naturalistic perception," *Communications Biology* 4: 79.

- Brewer, B. (2011) *Perception and its Objects*. Oxford: Oxford University Press.
- Brown, J. (2020) “What is epistemic blame?” *Nous* 54(2):389-407.
- Burge, T. (1988) “Individualism and Self-knowledge”, *Journal of Philosophy*, 85: 649–663.
- Burge, T. (2005) “Disjunctivism and Perceptual Psychology.” *Philosophical Topics*, 33, no. 1: 1-78.
- Burge, T. (2010) *Origins of Objectivity*. Oxford: Oxford University Press.
- Byrne, A. (2005) “Introspection,” *Philosophical Topics*, Vol. 33 No.1, 79-104.
- Campbell, J. (2002) *Reference and Consciousness*. Oxford: Oxford University Press.
- Campbell, J. (2010) “Demonstrative Reference, the Relational View of Experience and the Proximality Principle.” In *New Essays on Singular Thought*, edited by Robin Jeshion, 193-212. Oxford: Oxford University Press.
- Cassam, Q. (2014) *Self-knowledge for Humans*. Oxford: Oxford University Press.
- Cassam, Q. (2016) “Vice Epistemology,” *The Monist* 99(2):159-180.
- Chadha, M. (2013) “Meditation and Unity of Consciousness: A Perspective from Buddhist Epistemology,” *Phenomenology and the Cognitive Sciences*, 14:111–127.
- Chadha, M. (2015) “A Buddhist Epistemological Framework for Mindfulness Meditation,” *Asian Philosophy*, 25:1, 65-80.

Chatterjee, A. (2010) “Neuroaesthetics: A Coming of Age Story,” *Journal of Cognitive Neuroscience*, 23(1): 53-62.

Chatterjee, A and Vartanian, O. (2014) “Neuroaesthetics,” *Trends in Cognitive Science*, 18(7), 370–375.

Connolly, K. (2017) ‘Perceptual Learning’, *The Stanford Encyclopedia of Philosophy* (Summer 2017 Edition), Edward N. Zalta (ed.), URL = [<https://plato.stanford.edu/archives/sum2017/entries/perceptual-learning/>](https://plato.stanford.edu/archives/sum2017/entries/perceptual-learning/).

Connolly, K. (2019) *Perceptual Learning: The Flexibility of the Senses*. Oxford University Press.

Crane, T, and French, C. (2015) “The Problem of Perception.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, URL = [<https://plato.stanford.edu/archives/spr2017/entries/perception-problem/>](https://plato.stanford.edu/archives/spr2017/entries/perception-problem/).

Danziger, S., Ward, R., Owen, V. and Rafal, R. (2004) “Contributions of the human pulvinar to linking vision and action,” *Cognitive, Affective, & Behavioral Neuroscience* 4 (1), 89-99

Davidson, D. (1987) “Knowing One’s Own Mind.” *Proceedings and Address of the American Philosophical Association*, 60, no. 3: 441-58.

Declos, A. (2014) “The Aesthetic and Cognitive Value of Surprise,” *Proceedings of the European Society for Aesthetics*, 6: 1-18.

Dretske, F. (1969) *Seeing and Knowing*. Chicago: University Of Chicago Press.

Dretske, F. (1979) 'Simple Seeing,' in D. F. Gustafson and B. L. Tapscott (eds.) *Body, Mind, and Method*. Kluwer Academic Publishers, pp. 1-15.

Dretske, F. (1995) *Naturalizing the Mind*. Cambridge, MA: MIT.

Dretske, F. (2000) *Perception, Knowledge and Belief*. Cambridge: Cambridge University Press.

Dunne, J. D. (2015) "Buddhist Styles of Mindfulness: A Heuristic Approach," in *Handbook of Mindfulness and Self-Regulation*, Brian D. Ostafin, Michael D. Robinson, and Brian P. Meier (eds.), 251-70. New York: Springer.

Ekman, P. (2007) *Emotions Revealed: Recognizing Faces and Feelings to Improve Communication and Emotional Life*. Holt Paperbacks; 2nd edition.

Elpidorou, A. and Gibson, J. (forthcoming) "Really boring art," *Ergo*.

Fridland, E. (2017) "Skill and motor control: intelligence all the way down," *Philosophical Studies*, 174: 1539–1560.

Evans, G. (1982) *The Varieties of Reference*. Oxford: Oxford University Press.

Fish, W. (2009) *Perception, Hallucination, and Illusion*. New York: Oxford University Press.

Fox, K. C. R., Dixon, M. L., Nijeboer, S., Girn, M., Floman, J. L., Lifshitz, M., et al. (2016) "Functional neuroanatomy of meditation: a review and meta-analysis of 78 functional neuroimaging investigations," *Neurosci. Biobeh. Rev.* 65, 208–228.

- Gertler, B. (2000) "The mechanics of self-knowledge," *Philosophical Topics*, 28: 125–146.
- Gertler, B. (2001) "Introspecting phenomenal states," *Philosophy and Phenomenological Research*, 63: 305–328.
- Goldman, A. (2006) *Simulating Minds*. Oxford: Oxford.
- Gow, L. (2018) "Why Externalist Representationalism is a Form of Disjunctivism." *Ratio*, 31, no. 1: 35-50.
- Gunaratana, B. (1991) *Mindfulness in Plain English*. Wisdom Publications.
- Hansen, T., Olkkonen, M., Walter, S. and Gegenfurtner, K. R. (2006) "Memory modulates color appearance," *Nature Neuroscience* 9(11):1367-1368.
- Hawley, K. & Macpherson, F. (eds.) (2011) *The Admissible Contents of Experience*.
- Hodgins, H. S., and Adair, K. C. (2010) "Attentional Processes and Meditation," *Conscious Cognition* Dec;19(4):872-8.
- Hoffman, D. (2019) *The Case Against Reality*. New York: W. W. Norton Company.
- Hornsby, J. (1980) *Actions*. London: Routledge & Kegan Paul.
- Hornsby, J. (1997) *Simple-Mindedness: In Defense of Naïve Naturalism in the Philosophy of Mind*. Cambridge, MA: Harvard University Press.
- James, W. (1890) *The Principles of Psychology*, Vol. 1. New York: Dover Publishing, Inc.

- Johnston, M. (2004) "The Obscure Object of Hallucination." *Philosophical Studies*, 120: 113-83.
- Johnston, M. (2009) *Saving God*. Princeton University Press.
- Johnston, M. (2019) "Sensory Disclosure: Neither a Propositional, Nor a Factive, Attitude." In *Non-Propositional Intentionality*, edited by Alex Grzankowski and Michelle Montague, 152-191. Oxford: Oxford University Press.
- Judge, J. (2018) "The surprising thing about musical surprise," *Analysis*, 78(2): 225-234.
- Kellman, P. J. & Garrigan, P. (2009), 'Perceptual Learning and Human Expertise,' *Physics of Life Reviews*, 6(2): 53–84.
- Lazar, S. W., Kerr, C. E., Wasserman, R. H., Gray, J. R., Greve, D. N., Treadway, M. T., McFarvey, M., Quinn, B. T., Dusek, J. A., Benson, H., Rauch, S. L., Moore, C. I., and Fischl, B. (2005) "Meditation experience is associated with increased cortical thickness," *Neuroreport* Nov 28; 16(17): 1893–1897.
- Lawlor, K. (2009) "Knowing What One Wants," *Philosophy and Phenomenological Research* 79(1): 47-75.
- Levinson, M., Podvalny, E., Baete, S. H., and He, B. J. (2021) "Cortical and subcortical signatures of conscious object recognition," *Natural Communications* 12:2930.
- Lord, E. (2020) "The Vices of Perception," *Philosophy and Phenomenological Research* 101: 727-734.
- Mandelbaum, E. (2014) 'Thinking is Believing,' *Inquiry* 57(1): 55-96.

- Mandelbaum, E. (2017) "Seeing and Conceptualizing: Modularity and the Shallow Contents of Perception," *Philosophy and Phenomenological Research*, 1-17.
- Martin, M. (2002) "Particular Thoughts & Singular Thoughts." In *Logic, Thought and Language*, edited by Anthony O'Hear, 173-214. Cambridge: Cambridge University Press.
- Martin, M. (2004) "The Limits of Self-Awareness." *Philosophical Studies*, 120: 37-89.
- Matthen, M. (2010) "Two visual systems and the feeling of presence," in *Perception, Action, and Consciousness: Sensorimotor Dynamics and Two Visual Systems*, Nivedita Gangopadhyay, Michael Madary, and Finn Spicer (eds.), Oxford: Oxford.
- Matthen, M. (2015) 'Play, Skill, and the Origins of Perceptual Art,' *British Journal of Aesthetics* 55(2): 173-197.
- Matthen, M. (2019) "Objects, seeing, and object-seeing," *Synthese* 198: 3265-3288.
- McDowell, J. (1994) *Mind and World*. Harvard University Press.
- McDowell, J. (2013) "Perceptual Experience: Both Relational and Contentful." *European Journal of Philosophy*, 21, no. 1: 144-157.
- Millière, R., Carhart-Harris, R. L., Roseman, L. Trautwein, F., and Berkovich-Ohana, A. (2018) "Psychedelics, Meditation, and Self-Consciousness," *Frontiers in Psychology* 9:1475.
- Milner, A. & Goodale, M. (1995) *The Visual Brain in Action*. Oxford: Oxford University Press.
- Milner, A. & Goodale, M. (2006) 'Epilogue: Twelve years on,' in A. Milner & M.

Goodale (eds.) *The Visual brain in action*. Oxford: Oxford University Press, 2<sup>nd</sup> ed, pp. 207-252.

Montague, M. (2013) 'The Access Problem,' in U. Kriegel (ed.) *Phenomenal Intentionality*. Oxford: Oxford University Press, pp. 28-46.

Montague, M. (2016) *The Given*. Oxford University Press.

Montero, B. G. (2016) *Thought in Action: Expertise and the Conscious Mind*. Oxford University Press.

Nanay, B. (2016) *Aesthetics as Philosophy of Perception*. Oxford University Press.

Nudds, M. (2009) "Recent work in perception: Naïve realism and its opponents," *Analysis* 69 (2):334-346.

Noë, A. (2006) *Action in Perception*. The MIT Press.

Paul, S. K. (2014) "The Transparency of Mind," *Philosophical Compass*, 9/5: 295-303.

Pautz, A. (2007) "Intentionalism and Perceptual Presence." *Philosophical Perspectives*, 21, no. 1: 495–541.

Pautz, A. (2021) "Naïve Realism and the Science of Sensory Consciousness," *Analytic Philosophy*.

Peacocke, C. (2018) "Are perceptions reached by rational inference? Comments on Susanna Siegel, The Rationality of Perception," *Res Philosophica* Vol. 95(4): 751-760.

Peters, M. A. K., Lau, H. (2015) 'Human observers have optimal introspective access to



perceptual processes even for visually masked stimuli,' *eLife* 4:e09651.

Peters, M. A. K., Kentridge, R. W., Phillips, I., Block, N. (2017) 'Does unconscious perception really exist? Continuing the ASSC20 debate,' *Neuroscience of Consciousness*, 3(1): 1-11.

Phillips, I. (2018) 'Unconscious Perception Reconsidered,' *Analytic Philosophy*, 59(4): 471-514.

Potter, M., Wyble, B., Haggmann, C., & McCourt, E. (2014). "Detecting meaning in RSVP at 13 ms per picture." *Attention, Perception, & Psychophysics*, 76(2), 270–279.

Roskies, A. (2010) "'That' Response doesn't Work: Against a Demonstrative Defense of Conceptualism,' *Nous*, 44(1): 112-134.

Ransom, M. (2022) "Aesthetic Perception and the Puzzle of Training," *Synthesis*, 200(2).

Schellenberg, S. (2018) *The Unity of Perception*. Oxford: Oxford University Press.

Schwitzgebel, E. (2012) "Introspection, What?," in *Introspection and consciousness*, Declan Smithies and Daniel Stoljar (eds.), Oxford: Oxford.

Schwitzgebel, E. "Introspection", *The Stanford Encyclopedia of Philosophy* (Winter 2019 Edition), Edward N. Zalta (ed.), URL = [<https://plato.stanford.edu/archives/win2019/entries/introspection/>](https://plato.stanford.edu/archives/win2019/entries/introspection/).

Searle, J. (2015) *Seeing Things as They are*. Oxford: Oxford University Press.

Shoemaker, S. (1994) "'Self-knowledge and 'inner sense'," *Philosophy and Phenomenological Research*, 54: 249–314.

Siegel, S. (2006) ‘What Properties are Represented in Perception?’, in T. Gendler and J. Hawthorne, J. (eds.) *Perceptual Experience*. Oxford: Oxford University Press, pp. 481-503.

Siegel, S. (2010) *The Contents of Visual Experience*. Oxford: Oxford University Press.

Siegel, S. (2017) *The Rationality of Perception*. Oxford University Press.

Sillito, A. M., Cudeiro, J., and Jones, H. E. (2006) “Always returning: feedback and sensory processing in visual cortex and thalamus,” *Trends in Neurosciences* 29(6): 307-316.

Snow, J. C., Allen, H. A., Rafal, R. D., and Humphreys, G. W. (2009) “Impaired attentional selection following lesions to human pulvinar: Evidence for homology between human and monkey,” *Biological Sciences*, 106(10): 4054-4059.

Sosa, D. (2007) “Perceptual Friction.” *Philosophical Issues*, 17, no. 1: 245–261.

Strawson G. (2015) “Real Direct Realism: Reflections on Perception,” in *Phenomenal Qualities: Sense, Perception and Consciousness*, edited by P. Coates and S. Coleman. Oxford: Oxford University Press.

Strawson, P. F. (1979) “Perception and its Objects.” In *Perception and Identity: Essays Presented to A.J. Ayer with His Replies*, edited by Graham Macdonald. London: Macmillan.

- Thompson, B. (2008) "Representationalism and the Argument from Hallucination." *Pacific Philosophical Quarterly*, 89: 384-412.
- Tononi, G., Koch C. (2015) "Consciousness: here, there and everywhere?," *Phil. Trans. R. Soc. B* 370: 20140167.
- Travis, C. (2013) *Perception: Essays After Frege*. Oxford University Press.
- Tye, M. (2000) *Consciousness, Color, and Content*. Cambridge, MA: MIT.
- Tye, M. (2014) "Transparency, Qualia Realism, and Representationalism," *Philosophical Studies*, 170: 39-57.
- Tye, M. (2014) 'Does Conscious Seeing Have A Finer Grain Than Attention?,' *Thought*, 3: 154-158.
- Tye, M. (2015) "Yes, Phenomenal Character Really is Out There in the World." *Philosophy and Phenomenological Research*, 91, vol. 2: 483-8.
- U Pandita, S. (2020) "What Is Vipassana Meditation and How Do You Practice It?" *Lion's Roar*, URL = < <https://www.lionsroar.com/how-to-practice-vipassana-insight-meditation/>>
- Walton, K. (1970) "Categories of Art," *The Philosophical Review*, 79(3): 334-367.
- Watzl, S. (2018) 'Is Attention a Non-propositional Attitude?,' in A.Grzankowski and M. Montague (eds.) *Non-propositional Intentionality*. Oxford: Oxford University Press, pp. 272-302.

Wimmer, R. D., Schmitt, L. I., Davidson, T. J., Nkajima, M., Deisseroth, K., and Halassa, M. M. (2015) “Thalamic control of sensory selection in divided attention,” *Nature* 526: 705-709.

Wu, W. (2011) ‘Confronting Many-Many Problems: Attention and Agentive Control,’ *Nou<sup>^</sup>s* 45.1: 50–76.

Zhou, H., Schafer, R. J., and Desimone, R. “Pulvinar-Cortex Interactions in Vision and Attention,” *Neuron* 89(1):209-20.