Australasian Journal of Philosophy

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/rajp20

The Higher-Order Map Theory of Consciousness

Joseph Gottlieb

To cite this article: Joseph Gottlieb (2021): The Higher-Order Map Theory of Consciousness, Australasian Journal of Philosophy

To link to this article: https://doi.org/10.1080/00048402.2020.1861038

Published online: 02 Feb 2021.

AUSTRALASIAN JOURNAL OF PHILOSOPHY
https://doi.org/10.1080/00048402.2020.1861038

The Higher-Order Map Theory of Consciousness

Joseph Gottlieb

Texas Tech University

ABSTRACT

I begin by developing a challenge for the Higher-Order Thought variant of Higher-Order representational theories of consciousness. The challenge is to account for the distinctive phenomenal character of visual experience—its *presentational character*. After setting out the challenge, I articulate a novel form of Higher-Order theory that can account for presentational character—the *Map Theory* of consciousness. The theory's distinctive claim is that the relevant higher-order representations have a cartographic format.

KEYWORDS HOT theory; consciousness; maps; iconicity; presentational character

1. Non-Relational HOT Theory

The Higher-Order Thought ('HOT') theory is one variant of the Higher-Order ('HO') theoretic approach to phenomenal consciousness. HOT theory is often framed in *relational* terms, where, for any first-order perceptual state M of a subject S, M is (phenomenally) conscious at *t* iff S harbours a suitable HOT M* that represents M at *t*.[1] Here it is first-order ('FO') perceptual states—those with exclusively first-order or world-directed contents—that are conscious. And their being conscious just consists in their being represented by a suitable HOT.

Theories of consciousness should do at least two things. First, they should give an existence condition: they should tell us the difference between conscious experience and unconscious perception. Second, they should give us an identity condition: they should tell us what individuates one conscious experience from another. Implicit in relational HOT theories is a specific view of how theories should interpret this first task. These theories, along with competing FO theories, treat giving an existence condition as equivalent to answering the following question: in virtue of what is a token FO perceptual state M conscious, rather than unconscious?[2] The relational HOT theorist says that M is unconscious when it is not suitably HOT-targeted, and conscious when it is.

Interestingly, it turns out that few, if any, HOT theorists endorse relational HOT theory. Most, if not all, endorse a *non-relational* HOT theory, in which no relation between M* and M is needed:[3]

[1] I focus on 'ambitious' HO theories [Block 2011]. For brevity, I will largely elide 'phenomenal'. I use 'HO' to avoid assumptions about the nature of the HO representation (M*).

[2] FO theories of consciousness (e.g. Dretske [1995]; Tye [2000]) deny that HO representation is necessary or sufficient for being in a conscious state.

[3] For non-HOT theorists who frame HOT theory as relational HOT theory, see, e.g., Byrne [1997: 107], Wilberg [2010: 618], Block [2011: 421], Edwards and Platchias [2018: 345], and Farrell [2018: 2749]. Rosenthal [2005:

Non-Relational HOT Theory

EXISTENCE For any subject S, S's being in a suitable higher-order mental representation M* (an assertoric and *de se* HOT) is necessary and sufficient for S's being phenomenally conscious.

IDENTITY The content of M* (the HOT) wholly determines phenomenal character—the respect in which S is phenomenally conscious.[4]

On Non-Relational HOT theory, M* represents S as being in a first-order perceptual state M. But M is not a conscious state—at least not in the sense that M is the state in virtue of which there is something it is like for S. That state is M*, although S need not be aware of being in M* for M* to be conscious in this sense. Strictly speaking, the property *being (phenomenally) conscious* is only instantiated by subjects [Berger 2014; Brown 2015]. This property *just is* the property of having the right kind of M*. We can speak as if M is conscious—and, admittedly, non-relational HOT theorists occasionally do—but all that this amounts to is that M is the state, existent or not, in which we *represent* ourselves as being. It is conscious in the sense of 'mental appearance'; M is a state that we *appear* to be in [Rosenthal 2011: 431].

Given EXISTENCE and IDENTITY, the non-relationalist rejects, as question-begging, the relational framing of what it is to give existence conditions. The neutral framing is this: what is the difference between (i) having a phenomenally conscious experience of, say, an apple and (ii) having an unconscious perception of an apple?[5] And, for the non-relational HOT theorist, the answer is this: (i) is a matter of S's tokening a suitable HOT M*, which will represent S as being in a FO perceptual state M about the apple (where S need not be in M), while (ii) is a matter of S's being in M alone.

I aim to put the best face on non-relational HO theory. The HOT variant of non-relational HO theory incurs a unique explanatory burden: it needs to explain the character of our perceptual experiences without recourse to FO perceptual states. After all, if S's being conscious, and thus undergoing a conscious experience, is just a matter of S representing herself (via a HOT) as being in some perceptual state M, there is no guarantee that S will be in M. Representation implies the possibility of misrepresentation.[6] So M*, a HOT, must bear the phenomenological burden alone.

But, *prima facie*, it's unclear how a mere HOT can do this. It is standardly assumed that the *format* of HOTs is wholly discursive, or language-like. This is a mistake: if HOTs are wholly discursive, then non-relational HOT theory cannot account for *presentational character*, the distinctive phenomenal character of paradigmatic visual experiences. Non-relationalists should instead insist that M* is *cartographically*

211, 2011], Berger [2014], Brown [2015], and Lau and Brown [2019] reject relational HOT theory. One reason why HOT theorists are read as relational HOT theorists is that relational formulations offer an easy way to gain a sense of the theory, and (non-relational) HOT theorists take advantage of this [Rosenthal 2005: 211]. Confusion over the transitivity principle is a related culprit [Gottlieb 2020]. See note 21.

[4] There are occasional quibbles about the use of 'phenomenal character' in some corners of the HOT landscape, perhaps due to Rosenthal's [2010] claim that mental qualities can occur unconsciously. But, on the present usage, this term just means 'what it is like' for a subject to have an experience. In that sense, HOT theories are interested in explaining phenomenal character.

[5] Many thanks to Jake Berger for discussion here.

[6] Although M will normally be causally responsible for S's being in M*, S's being in M is not constitutive of S's being conscious. So-called 'empty HOT' cases—where M* represents S as being in M when S is not, in fact, in M—are sometimes said to reveal an internal inconsistency in HOT theory [Block 2011]. Yet this is only so on a relational HOT theory, not a non-relational one. See Rosenthal [2011] and Weisberg [2011]. There is empirical evidence for consciousness without M-type states [Lau and Brown 2019], although it is controversial [Block 2019].

formatted, with both discursive and iconic components. I call this *the Map Theory* of consciousness.

Here are a few provisos. First, although I will argue that Map Theory is preferable to discursive HOT theory, the case for Map Theory is partially parasitic on the case for HO theories generally, and for non-relational HO theories specifically. Thus, my overarching thesis is a conditional: if one endorses non-relational HO theory, then one ought to endorse the Map Theory.

Second, I will not settle on a definitive version of Map Theory. Instead, I will focus on setting out its broad contours. I am not so interested in delving into what *really* counts as map; the Map Theory is so-called because, on it, M* has certain features that are arguably essential to, and perhaps sufficient for, being a map—a mix of discursive and iconic elements, the exploitation of spatial isomorphisms, and so on. Another issue that will emerge is whether a HOT theorist can adopt the Map Theory—that is, whether M* can be a *thought* if it is cartographically formatted, or whether it is a higher-order *perceptual* representation (a 'HOP'), as on HOP theory. While a case can be made for M* being a HOT on Map Theory, this depends on how we mark the perception-cognition border. In the end, it might be that the only viable non-relational HO theory is a non-relational HOP theory, albeit one that is different from classic HOP theory (for example, Lycan [1996]).

2. The Problem of Presentational Character

Presentational character is a species of phenomenal character, and so it categorises experiences by what it is like to have them. It is the kind of phenomenal character emblematic of visual experience. When you have a visual experience as of an apple, the apple seems to be directly present to you in a rich, vivid, and detailed way. Experiences have presentational character *with respect to* something, where that something (the apple) is what's phenomenally presented.[7]

The problem of presentational character goes like this:

P1 If non-relational HOT theory is true, HOTs are the sole determinant of phenomenal character.

P2 Presentational character is a kind of phenomenal character.

∴ If non-relational HOT theory is true, HOTs are the sole determinant of presentational character.

P3 If HOTs are the sole determinant of presentational character, then they must have content capable of engendering experiences with presentational character.

P4 HOTs do not have such content.

∴ HOTs are not the sole determinant of presentational character.

∴ Non-Relational HOT theory is false.[8]

[7] Visual hallucinations have presentational character. If we are not aware of, or presented with, anything in hallucination, we can insist that 'being presented' is intensional, and we can cast presentational character in terms of the distinctive way that we *seem* to be aware of things in visual experience.

[8] This is a revamped version of an argument based on presentational character that I have made elsewhere [2015]. The prior version grants the falsity of P4 as a dialectical manoeuvre, and instead argues that, even if P4 were false, the introspective mechanisms posited by HOT theory would (contrary to the phenomenological data) make it seem as if a thought were presenting its object. I now repudiate this variant of the argument.

P1 falls out of IDENTITY. Rosenthal endorses it explicitly [2009: 249, emphasis added]:

> If one has a sensation of red and a distinct HOT that one has a sensation of green, the sensation of red may nonetheless be detectable by various priming effects. But what it will be like for one is that one has a sensation of green. Similarly if one has that HOT *with no relevant sensation at all.*

Here, too, with Weisberg concurring [Rosenthal and Weisberg 2008, emphasis added]:

> What it's like for one [on HOT] is determined by the way the higher-order awareness represents the first-order state … Consciousness is a matter of mental appearance … that mental appearance is due *solely* to the higher-order awareness.

P2 is true by definition. P3 follows from EXISTENCE, IDENTITY, and the fact that HOT theory is a form of representationalism about phenomenal character.

This leaves P4. There is a *prima facie* puzzle here. HOTs are supposed to be just like ordinary thoughts. Yet FO thoughts do not yield the kind of awareness that we associate with presentational character. When I think about the apple, the apple is not phenomenally presented to me.[9] What makes HOTs different? The onus is on the HOT theorist to answer this question and to tell us why P4 is false.

An immediate answer is that only HOTs have higher-order content. HOTs represent their subjects as being in mental states. And when a subject undergoes a visual experience with presentational character, the HOT will represent the subject as *seeing* something. However, the relevance of this difference can only pertain to why FO thoughts are not sufficient for having an experience with *phenomenal* character. This is not enough to explain why HOTs are sufficient for having an experience with *presentational* character—a specific determinate of phenomenal character. In S's experience, the apple is phenomenally presented. But, in order for this to be so, it cannot be sufficient that I have a HOT that simply represents me as seeing the apple. For then the only thing that makes HOTs relevantly different from FO thoughts—states that do not phenomenally present their objects—is the representation of something extra—namely, the seeing-relation. Yet this extra something has nothing to do with what is phenomenally presented—namely, the apple. True, the seeing relation is relevant in so far as one of its relata (the apple) is what is phenomenally presented, but the seeing-relation *itself* is not what's presented. This suggests that the answer resides with *how* the HOT represents the *apple*.

In light of this, notice two things. First, my visual experience of the apple is rich: I am apparently given a lot of information about the apple's features (its size, shape, colour, and location). Second, my experience is fine-grained: the apple is represented in a highly determinate manner. We don't just represent the apple as being coloured, or as being generically red; we represent a minimally discriminable colour, the *precise redness* of the apple [Bourget 2017].

These two conditions (richness and fineness of grain), when combined with HO content, point the way towards our answer:

> PRESENCE For any subject S, S has an experience with presentational character iff (i) S has a HOT M*, and (ii) M* represents S as seeing something, where the content of M* is sufficiently rich and fine-grained with respect to what S is.

[9] On the phenomenal difference between thought and visual experience, see, e.g., Searle [1983], Martin [2002], Bourget [2017], and Gottlieb [2018].

Clause (i) of PRESENCE falls out of EXISTENCE. It tells us what is necessary and sufficient for a subject to undergo an experience. Clause (ii) tells us when a subject has an experience *with presentational character*.

Thus, from a HOT-theoretic perspective, my FO thought about the apple does not phenomenally present the apple, because (i) FO thoughts don't have HO content and (ii) they are not rich and fine-grained. My FO-thought that the apple is red does not represent me as seeing the apple; nor does it represent the apple in detail. Since (i) holds by definition, taking issue with it here is just taking issue with HO theory. What matters is why we think that FO thoughts typically *don't* have rich and fine-grained content, and whether we can say that nonetheless HOTs *do*.

The way forward turns on how the non-relational HOT theorist traditionally understands P4. For her, P4 is just this:

P4* Discursively formatted HOTs do not have such content.

The difference between P4* and P4 is that the former draws an explicit connection between format and content. P4* is about the *actual* world. It might be conceptually possible that P4* is false, and perhaps conceptually possible that a discursive, non-relational, HOT theory is true. Yet when we reflect on the content of *actual* discursive representations, we are provided with little reason to think that HOTs could be wholly discursive, given the role that HOTs are supposed to play. Thus, with PRESENCE in place, the discursive HOT theorist needs to show that P4* is false at the actual world. I claim that this is untenable. P4* is probably true.

To see why, we need a better grip on representational formats. Formats are types of vehicular structures. The sentence 'This is a horse' and a picture of a horse differ in representational format. Sentences, photographs, maps, blueprints, graphs, and hieroglyphs are distinct representational formats.

Compare a picture of Arnold Schwarzenegger sitting in the Governor's chair (see Fig. 1, below) with the sentence:

(1) Arnold Schwarzenegger is sitting in the Governor's chair.[10]

Figure 1 is an *iconic* representation of Schwarzenegger. Sentence (1) is a *discursive* representation of Schwarzenegger. Following Quilty-Dunn [forthcoming-b], icons respect PARTS: parts of icons correspond to parts of what they represent (cf. Fodor [2007]). For example, the middle of Figure 1 corresponds to Schwarzenegger's face and the chair in which he is sitting. By contrast, discursive representations do not obey PARTS. There is no part of (1) that corresponds with any part of Schwarzenegger's chair. Icons also respect HOLISM: parts of icons represent in a holistic manner, representing multiple properties at once. The part of the picture that represents the left side of the collar of Schwarzenegger's shirt also represents its shape, its visual texture, and its colour. Because pictures don't break down into separable elements for each represented item, a picture that represents a scene brings along for free an array of properties represented in that scene. Channelling Fodor [ibid.: 108], icons do not represent by means of a *canonical decomposition* into constituents that stand separately for distinct individuals and distinct features. By contrast again, sentences like (1) do not obey HOLISM. Some parts of (1) represent individual items or properties without being

[10] For this way of making the comparison, see Quilty-Dunn [forthcoming a].

Figure 1. Arnold Schwarzenegger sitting in the Governor's chair.

holistically bound up with other items or properties. For example, 'chair' represents an object without representing any of that chair's various features. Whereas pictures (and icons more generally) are *analogue* in the Dretskian [1981: 136, 137] sense of carrying both nested and non-nested information, the words in (1) are *digital*, and thus do not carry non-nested information. With (1), no more specific information is given about Schwarzenegger's chair other than that he is sitting on it.[11]

Two consequences of this are worth stressing. First, icons are readily representationally rich. The picture of Schwarzenegger represents him as being in his chair, but also represents his sitting position, and the colour and shape of the chair. Yet in (1) we are told that Schwarzenegger is sitting in the Governor's chair without being told anything else.

Relatedly, icons are not subject to *item-effects*. This is because they obey PARTS; icons do not require that *more* items be represented by means of *more* vehicles [Quilty-Dunn forthcoming b]. This enables icons to represent many objects in a compact form. Fodor [2007: 111] puts it well when he remarks how 'a photograph of 60 giraffes takes no more space in your album (or on the screen) than a photograph of 6 giraffes.' On the other hand, discursive representations *are* subject to item-effects,

[11] While common in the literature, 'iconic' and related terms are to some extent misleading since they suggest computer icons (as on a computer home screen) or something of influence or reverence. Here I only employ the technical use as defined above; so, 'icon' covers a family of representations, all of which satisfy PARTS and HOLISM. Later, I introduce what I call *pure icons* which, unlike most maps but like traditional pictures, are isomorphic to what they represent along multiple dimensions.

Figure 2. Extreme Richness [RVCJ Media].

precisely because they do *not* obey PARTS. In (1), there is a vehicle that represents Schwarzenegger and a distinct vehicle that represents the chair. If (1) were to represent more things, we would have to increase the number of vehicles. This is not achieved for free. There are computational costs incurred by the implementation and processing of new vehicles. So, when mental states are formatted discursively, and we want to represent increasingly rich scenes, there is an associated burden on processing limits (cf. Quilty-Dunn [forthcoming a]). That's not so with icons.

Second, icons are readily fine-grained. A picture can represent a stop sign as coloured, as red, as crimson, or as crimson$_7$ As Camp puts it, pictures 'deliver information about a continuous spatial array, and the information they deliver about each point in that array is itself typically continuous—for instance, specifying fine-grained color' [2007: 156].

With these points in place, let's return to (4*), focusing on richness. It will help to start with an extreme case—the experience you might have while at Times Square on a busy night (see Fig. 2).

From your vantage point behind the George Cohan statue, your visual field is crowded. You are evidently presented with an immensely rich and detailed scene. The non-relationalist will say that this visual appearance is the result of the tokening of a discursive HOT. This is implausible. The concern is not with our ability to *express* such a HOT: it is not surprising that we can't fully describe, verbally or otherwise, the visual appearances. Nor is the concern that we *can't conceive* of rich discursive HOTs: we can easily envision HOTs as long conjunctive sentences. The problem is that it is not clear that such thoughts are *actually* thinkable, given their complexity. At issue is whether we can have HOTs that can *confer* these appearances. For if HOTs *are* discursive, our ability to think *them* will be contingent on our ability to think an immensely long conjunctive sentence, or many short sentences all at once. The idea behind

this complaint is somewhat familiar [Byrne 1997: 117], but the current iteration renders it precise: worries about the complexities of HOT are a function of (i) their purported discursive format, (ii) the nature of presentational character, and (iii) the fact that HOTs do all of the phenomenological work.

These points about complexity and processing limits can be made without determining precisely what those limits are. Consider Sperling's [1960] partial-report paradigm. The evidence from these experiments suggests that we possess a high-capacity sensory memory store—dubbed 'iconic memory' or 'sensory memory'—that outstrips the 3-4 item limits of visual working memory (VWM). What *explains* why VWM, but not sensory memory, is limited in this way is that the former is discursive and the latter is iconic. That we are limited to 3-4 items in VWM suggests that the represented items are represented by something that is subject to item-effects—namely, a discursively formatted representation.

On its face, these facts, when coupled with the phenomenology of visual experience, make the case for M* being discursive rough going. For, even supposing that the processing limits of M* were not the same as those of VWM, if it is discursive, then, for every extra object or property that M* represents, the number of vehicles comprising it will correspondingly increase. Cases of extreme richness, as in Figure 2, underscore the challenge: whatever the representation underlying such a visual experiences is, its processing limits must be so large that it's not clear what reason we have to call it 'discursive'.

What should the discursive, non-relational HOT theorist say? An initial reaction is to deny the phenomenology. The HOT theorist might insist that visual experiences do not have a rich and fine-grained phenomenology. We don't need to think extremely complex HOTs, because there is no phenomenological data that would be explained by our thinking such HOTs.

Introspective and phenomenological disputes are hard to adjudicate. Yet the richness of visual experience seems to be well-accepted. Not all visual experiences need be apparently *as rich* as the one that you might have when standing behind the George Cohan statue, but the present point does not rely on this. And, at any rate, it would be better, all things considered, to design a theory that takes the phenomenological reports at face-value. Indeed, taking these claims about richness and granularity at face-value is a central plank in many experimental paradigms in cognitive science.

Let's suppose, then, that the phenomenological data are not in dispute. Could the HOT theorist nonetheless insist that we can explain the data with a HOT that is not itself rich? This move is, in effect, to deny clause (ii) of PRESENCE. It is to say that how things *seem* to us is as if we are aware of a richly detailed scene, but that our corresponding HOTs—those states in virtue of which things seem this way—are not themselves representationally rich.[12]

Along these lines, Cohen et al. suggest that conscious experience involves *ensemble representation*, where the visual system represents 'groups of items as a statistic that summarizes different types of information' [2016: 326]. This cuts down the amount of information within a scene that is represented at high fidelity (those parts of the scene that are attended to or foveated), and so might be counted as a view in which (if co-opted by the HOT theorist) HOTs would be relatively 'simple'. Cohen et al. grant the phenomenological point, and explain away the reporting discrepancies in

[12] Thanks to an anonymous referee for pressing me to consider this idea more fully.

change blindness by the fact that the changes across scenes preserve summary statistics [ibid.: 328]. Cohen et al. insist that we experience far more than the 3–4 items to which we are typically restricted by VWM [ibid.: 325, 327, 332], but that everything beyond our attention is represented in summary fashion.[13]

I have two worries here. First, there is evidence that ensemble representations exhibit item-effect size limits [Im and Chong 2014: 664; Quilty-Dunn 2017: 133]. This is especially concerning in cases where one's visual periphery is crowded, like the Times Square case above, and so might tell against M*'s including ensemble representation. Second, richness includes not just the number of items represented (the number of letters on a Sperling grid) or 'object-richness', but also the features of those items, or 'feature-richness'.[14] The Times Square case purports to exemplify both in an extreme manner: you experience very many items and very many features of those items. Here, Cohen et al.'s strategy might seem appropriate. Yet many experiences are merely phenomenologically feature-rich. Suppose that I fixate on a single item—half of a toasted bagel, with the toasted-side facing upwards. This bagel is not a uniform surface, devoid of variation. It is instead highly variable, characterised by many shapes and colours (depending on evenness of toasting). These facts are made manifest phenomenologically. So, limiting actual richness to the few items attended-to does nothing to suggest that M* is discursive. If anything, the bagel case suggests that M* is at least partially iconic. M* can represent me as seeing *one* item, and so be simple in that regard, but nonetheless represent me as seeing one *highly complex* item, and so still be feature-rich.[15]

Turning to granularity, we might note, with Rosenthal [2005: 186–7], that the concepts that are deployed as part of our HOTs allow us to become conscious of the mental qualities that figure into our first-order perceptual states in increasingly fine-grained ways. A HOT might represent me as seeing an apple that is a very specific shade of red, depending on the concepts that I possess. Rosenthal [ibid.: 187] mentions how the wine novice and connoisseur possess different conceptual repertoires, explaining why the latter can enjoy experiences that are more fine-grained than the former's. The connoisseur might possess the HO concept MENTAL OAKY, and so can appreciate details in her first-order perceptual representation (M) that she otherwise wouldn't be able to appreciate.[16]

Two points now arise. First, the phenomenon of interest is that of visual experiences that are *both* rich and fine-grained (codified in PRESENCE), and being fine-grained does not entail being rich. (See analogue speedometers.) So, the granularity of the HO concepts won't itself solve the problem. Second, even if we grant that concept acquisition can result in changes in phenomenology, at best Rosenthal has only given a condition

[13] See Knotts et al. [2019] for a useful adjunct to this theory, which might explain why the periphery doesn't *feel* summarised.

[14] For a discussion of feature capacity limits and how they relate to object capacity limits, see Green and Quilty-Dunn [forthcoming].

[15] In the context of defending conceptualism about experiential content, Gennaro [2012: 161] offers a 'deflationary' approach to richness. Yet Gennaro himself [ibid.: 163] concedes that at least some experiences are actually rich, only insisting that this does not mean that the 'analog distinction must map onto the conceptual-nonconceptual distinction'. Perhaps so, but this is no matter; the only claim to be defended is that M* must be cartographic, and so have a partially iconic (and thus analogue), and partially discursive, structure. Thus, even Gennaro would grant that, for those experiences that are rich, the current approach is to be favoured.

[16] MENTAL OAKY is the mental counterpart of OAKY: the latter is acquired via perceptual discrimination and applies to worldly entities, while the former applies to my experiences of those entities and is derivative on the latter.

on experience individuation. Given EXISTENCE, what we are trying to establish is a how a HOT—*independent of the first-order state that it describes*—can generate a visual experience with presentational character (cf. Picciuto [2017: 569]). Yet in wine-tasting cases the first-order state exists. Something is being tasted. Saying that the acquisition of HO concepts can make me aware of more fine-grained details of an *existing* and *already* rich and fine-grained first-order perceptual state is not to say that exercising those concepts in a HOT (without that first-order state) is sufficient for me to have a visual experience with fine-grained character.[17,18]

Let me summarise the case for P4*. Visual experience has two perspicuous phenomenological features—richness and high granularity. Regarding richness, we have little reason to think that a discursively formatted HOT could be adequately item-rich or feature-rich (and certainly not both), and little reason to think that a sparse HOT could explain rich phenomenology. Regarding granularity, HOTs in-principle fare better. Yet it's hard to see why simply exercising fine-grained descriptions of existent FO perceptual states in a HOT would be sufficient for presentational character in the absence of actual FO states. Moreover, presentational character is now fettered by the discursive concepts deployable in a subject's mind. By contrast, the representational apparatus exploited by icons faces no such constraint. The cleaner proposal—the one on offer with the Map Theory—is simply to say that experience seems rich and fine-grained *precisely because* it is constituted by a representation that *is* correspondingly so, and in a compact form.

While the problem of presentational character might be surmountable, it is serious enough to warrant our considering alternatives. Yet the challenge deepens when we look elsewhere within the HO-theoretic family.[19]

The culprit is P1. Relational HOT theorists can deny P1, but empty HOTs doom the theory (see Block [2011] and note 6 above). Self-representational variants of HO theory are sometimes cast in this light, but it's misleading: although FO perceptual states are always part of a conscious state for the self-representationalist, the content of these FO states do not play any role in determining phenomenal character [Kriegel 2009: 109–11]. That cuts against denying P1. In addition, self-representational theories face problems with naturalisation [Phillips 2014]. The HOT theorist faces no such problems, as HOTs don't represent themselves.[20] So, if non-relational HOT theory is true, the reduction of consciousness is possible. This is not a benefit to be cast aside.

Thus, there is pressure to retain the basics of a non-relational theory. I propose that we reject P4 by making M* a *mental map*—namely, a HO representation with a cartographic structure. The hybrid nature of maps—their having both iconic and discursive elements—will make us well-placed to offer not only a satisfying

[17] On Rosenthal's [2010] quality space theory, FO perceptual states will instantiate mental qualities (e.g. 'mental redness') that are as fine-grained as the perceptual discriminations that we make in virtue of being in those states. (The quality space of mental qualities is homomorphic to the external perceptible properties that a subject can discriminate.) But Rosenthal needs to show that HO concepts, deployed in the right way, are *sufficient* for presentational character; the appeal to actual mental qualities drops out as soon as we admit that there is presentational character in cases of HO misrepresentation. See also note 6.

[18] Likewise for views where HOTs are demonstrative. As it happens, Rosenthal rejects a demonstrative view [2005: 63n].

[19] One could dispense with the HO-theoretic approach altogether by moving to a FO approach (see note 2). While I set aside this option for narrative reasons, when it comes to existence conditions, scrutiny reveals that the purported difference between many HO and FO views might be merely apparent [Gottlieb 2019].

[20] Reductive accounts of representation typically appeal to causal relations, which are anti-reflexive.

treatment of presentational character, but also other facets of visual experience inherent to the HOT-theoretic approach more generally. Whether this approach will, in the end, result in a HOP non-relational theory, will be broached once the theory is set out.

3. The Map Theory

Like discursive non-relationalism, the Map Theory endorses Existence and Identity. Its distinctive commitment is this: S's undergoing a visual experience with presentational character consists in S's being in an M* with cartographic structure, where M* represents S as being in a visual state (as seeing something).[21]

Officially, the Map Theory leaves open the format of M* in non-visual experiences. But there are options. For experiences that do not have presentational character (like 'free floating' anxiety and the affective components of pain), assuming pluralism about format, the Map theorist can say that M* is wholly discursive. If auditory and olfactory experiences have a form of presentational character, the Map theorist can say that M* will be cartographic here, too. (Clause (ii) of Presence would simply be disjunctive, covering different perceptual relations.) The Map theorist can leverage the wide spectrum of map-types to capture the phenomenological variations between perceptual experiences in different sensory modalities.

So, those are the basic points. Two questions now emerge. Why maps? What kind of maps?

3.1 Why Maps?

There are two motivations for attributing cartographic structure to M*. First, we have independent reason to think that some mental states are map-like. This ranges from retinotopic maps in the visual system (see, for instance, Wandell et al. [2007]) to upstream states including ordinary thoughts [Braddon-Mitchell and Jackson 1996: 177–84; Camp 2007; Rescorla 2009a]. So, the proposal does not commit us to anything fundamentally new.

Second, while we have focused on presentational character, on non-relational HO theories visual experiences also have *non*-presentational character. M* represents more than the apple. It will also represent S herself, and it will represent S as seeing the apple. More generally, the content of M* will always be that *I*, the subject of that very state, am seeing (or smelling, or hearing) such-and-such [Berger 2014: 837]. However, neither S nor the seeing-relation is phenomenally presented. S is not aware of herself, or the seeing relation, in the way that she is aware of the apple.[22]

[21] HO theories are often construed as implementations of the transitivity principle. Quick-and-dirty formulations of that principle deny that M* is a conscious state, instead telling us that the (first-order state) M is conscious because we are aware of M (by being in M*). If M* is itself conscious, then—the worry goes—we are embarked upon a regress. However, at its most neutral, the transitivity principle only points to some interesting connection between phenomenal consciousness and inner awareness. The Map Theory honours this intuition, implementing it at a subject level: a subject S is phenomenally conscious iff S is aware of herself as being in some (first-order, perceptual) mental state M. Here, M* *confers* HO awareness, but in doing so S need not be aware of M* itself. See, e.g., Berger [2014], Brown [2015], and Gottlieb [2020].

[22] That the self is not phenomenally presented is arguably connected to Humeanism regarding introspection and the self. That the seeing-relation is not phenomenally presented is arguably connected to *transparency*. For issues concerning the connexion between transparency and HO theories generally, see Gottlieb [2016].

Maps, having iconic and discursive elements, are well-placed to handle both the presentational and the non-presentational aspects of visual experience espoused by non-relationalists. On the presentational side, maps obey PARTS and HOLISM, making them readily rich and fine-grained (see, for instance, Braddon-Mitchell and Jackson [1996: 181] and Camp [2007: 157–61]). On the non-presentational side, like sentences but unlike pictures, maps need not represent only what is visually perceptible. This is good, since, in whatever sense we can be aware of ourselves and our mental states, this is not literally by seeing them. And if our explanation of presentational characters turns in part on the fact that icons are readily rich and fine-grained in a way in which discursive representations are not, it is natural to treat the *non*-presentational aspects of visual experience discursively. However, that does not mean that we must have two mental states, one iconic and one discursive. The solution is to have just one representational mental state that, by its nature, can *bind together* iconic and discursive elements, in just the way that our awareness of ourselves and the seeing-relation is bound together with our awareness of the world on non-relationalism. This is, of course, precisely what M* would do if it were a map.

3.2 What Kind of Map?

Analogies with ordinary maps have limited explanatory power. Take a seating chart (Fig. 3):

Kelly	Luke	Abraham
Dante	Lucy	Janelle

Figure 3. A Minimal Map [Camp 2007: 158].

Maps are like pictures in so far as they exploit isomorphisms between vehicle and what is represented. But many maps are unlike pictures in so far as these isomorphisms are typically limited to spatial structure and functionally driven: maps typically represent only what is necessary, given their function, abstracting away from unnecessary details [Camp 2007: 158]. And the semantic link between the constituents of maps and the world is rather arbitrary [ibid.: 159]. A seating chart represents the relative locations of each seat, and nothing more. And it does so in an indirect way. 'Kelly' is a *marker*, or what Kulvicki [2015] calls a 'locatable feature'. Like 'Luke' or 'Abraham', it denotes the instantiation of a property at a location corresponding to the marker's location on the map. Because 'Kelly' is adjacent to 'Luke', we say that the map has the content <Kelly's seat is next to Luke's>. But there is nothing special about the use of 'Kelly' here. We could have used a 'K'. Or we could have used icons (see Fig. 4, below).

The maps in Figs. 3 and 4 have the same content. The pig-tailed icon (an iconic marker) and 'Kelly' (a discursive marker) tell us only *that* Kelly's seat is in such-and-such a location. Markers can be icons without functioning iconically. I can mark the location of buried treasure with a treasure chest picture or with an 'x'. Either way, I'm not representing any features of the treasure beyond its location.

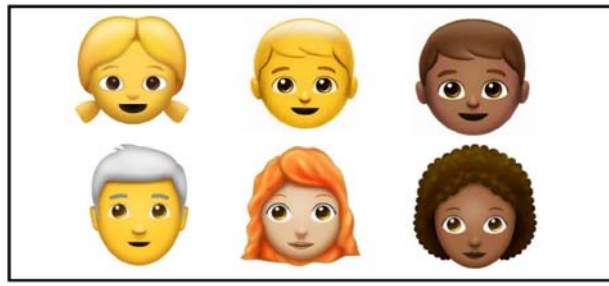

Figure 4. Iconic Markers.

Adopting such an approach is unworkable. Suppose that at *t* S has a visual experience of two apples, one red and one green, with the former appearing to the left of the latter. If our theorising were confined to ordinary maps, we might imagine M* being a map of S's visual field at *t*, with various markers indicating the apparent location of the apples relative to S's vantage point. Letting M* be a map of S's visual field is not the issue. We'll make use of that idea in a moment. The issue is that, even with iconic markers, such a map would say only *that* there is a red apple in the left-hand portion of my visual field, and *that* there is a green apple in the right-hand portion. There would be no rich and fine-grained representational uptake with respect to the items (apples) that are being phenomenally presented. Such a model offers no explanation of presentational character.

At this point, it will help to go back to our discussion of Sperling-style experiments. We noted that what makes so-called sensory memory iconic is its richly high-informational capacity and fine granularity. Yet sensory memory is plausibly 'purely' iconic, and so isomorphic to its content along *multiple* dimensions; subjects don't merely encode spatial information about the stimulus, like the letters' relative positions. So, pure icons are perfect for presentational character. Still, they are inapt for non-presentational elements. Consequently, M* must be a true representational hybrid: its iconic elements must be purely iconic or picture-like, but it must have discursive elements, too.

The preceding can be summarised with four conditions that any version of the Map Theory must meet:

(C1) **M* must be a complex.** M* is a single mental state, comprised of multiple representational vehicles. To ensure that we are not dealing with multiple states, each comprised of a single vehicle, M* must be a *complex*, not a mere sum (cf. Kriegel [2009: 221–2]).[23]

(C2) **M* must have higher-order content.** M* must represent what is phenomenally presented as *being seen* by the subject.

(C3) **M* must have a purely iconic component**. The isomorphisms exploited by M* cannot be limited to spatial structure. It must be closer to a picture.

(C4) **M* must be able to accommodate presentational variation.** By meeting C3, M* will be able to account for the presentational aspects of visual experience. But C3 must be met in a way that allows for variations in the *way* that items are phenomenally presented.

[23] Unlike sums, complexes are wholes that can cease to exist merely by virtue of the relation between their parts failing to obtain, even if the parts themselves persist. Hawai'i *qua* U.S. State is a complex [Kriegel 2009: 221]; a random collection of pebbles is a sum.

Here is one way to satisfy C1–C4. Suppose that S has a visual experience at *t* in which it seems to her as if two apples—one red, one green—are side-by-side on a table. Let M* be a map of the entirety of S's visual field—a map of what S's FO state M would be representing at *t*, if S were in M at *t*. The result is this (Fig. 5):

Figure 5. A Higher-Order Mental Map.

Since this is a visual field map, call it the *VF model* ('VF' for 'visual field') of Map Theory. The VF model satisfies C1–C4.

Regarding C1, on the VF model visual experiences consist of a pure icon (the map proper) and a single discursive vehicle. This discursive vehicle plays a unique role—that of a *title*. The title says (something like) 'Things that I am seeing'. It says what the map in its entirety is *of*, and the relationship between you and what is depicted on the map. (What's on the map is precisely what you are seeing at the moment.) Distinct visual experiences will have the same title, being individuated by *what* you are represented as seeing. And, like any other map, the pure icon and a single discursive vehicle will be bound together as a complex.

Titles are not markers or locatable features, but a kind of *annotation* [Kulvicki 2015: 169–70]. This is evidenced by their location-independence. Changes to a map's locatable features changes the map's semantic identity, and so what it represents: if I move 'Kelly' from the top-left quadrant to the bottom-right, the content of the map changes. Not so with titles: if T is a title, and T were placed on location L* in lieu of L, the map's content would *not* change. For this reason, 'things' in 'Things I am seeing' does not refer to just one apple, or to whichever apple for which the title is closest. This would be bad, as it would imply that I am seeing one apple, when my experience has it that I am seeing two.

As for how this reference relation is secured, we have fairly well-understood cases of reference determination between different strings of vehicles. Suppose that I said 'Frank left the party; he was tired.' This is a case of *anaphora*, where the anaphoric expression 'he' has 'Frank' as its antecedent. Analogously, we want the referent of 'things' in 'Things that I am seeing' (the title of the map) to be fixed by the referent of the iconic representation (the map proper), where (somehow) both vehicles join together as a complex. The point isn't that the friend of the VF model should treat the reference relation between title and icon as anaphoric. Rather, the point is that

what the VF model is proposing is not entirely foreign. Lingering questions here should be seen as a project, not a problem.[24]

What about C2? The title renders M* as a *higher-order* map. M* is a map of the subject's visual field, and so a map of what is represented at a first-order level. But it's not just this. For it also says that *what* is represented at the first-order level is *seen* by the subject. M* is a higher-order representation because (i) the self and seeing relation are part of the title, and (ii) the seeing-relation is a representational relation. Suppose that I had a map with the title 'Cities that I have visited'. I have never been to Des Moines, Iowa. Thus, if the map included a marker on that city's location, it would be inaccurate. If the map is accurate only if I have been to Des Moines, then I am part of the map's content, even though I am not overtly *depicted* on the map. Contents, in one good sense of that term, are the conditions under which a representation is accurate. Likewise, if I did not see an apple, but touched it, M* is inaccurate. So, the self and seeing-relation enter into M*'s content.

Regarding C3 and C4, presentational character is *gradeable*. Items in one's visual field will often be phenomenally presented to a greater or lesser degree. Items in one's periphery will be phenomenally presented but just barely so: relative to items that appear in the centre of one's field, they will be represented at a much coarser grain. But the self and seeing-relation are not presented at all. The VF model captures both points.

M* is partially composed of a pure icon (C3). This allows us to toggle the degree to which various items are phenomenally presented. While the apples at the centre of S's visual field are presented in a deeply vivid manner, the parts of the icon that correspond to the periphery of the subject's visual field can readily encode coarser-grained colour values—precisely as coarse-grained as we need them to be (C4). The result is that, when M* is formatted in this way, the back edges of the table will appear to S to be indistinct and less vivid.

The self and the see-relation, however, are represented discursively, and not represented as being seen. That's why they are not phenomenally presented. So, the VF Model gets the phenomenology right.

Is M* a genuine map, on the VF model? We have little resistance to saying that Google Maps (in 'Street View') are maps, or that aerial photos with various discursive markers (for example, a 'P' for 'picnic area') are maps. The present case is hardly different, merely adding a title. And, like ordinary maps, M* is a structural representation: the spatial relation between the part that represents the red apple and the part that represents the green apple has representational import.

As noted, if it's insisted that M* is not a map, then that is fair enough; above all, we want a complex mental state that binds together discursive and iconic vehicles in the

[24] The psychological mechanism that binds these vehicles together into a complex will probably involve *functional integration*. According to Kriegel [2009: 235], this is a 'process of unifying disparate bits of information into a single representation in a functionally significant way'. This is well-studied in intramodal cases [Treisman 1996], where *neural synchrony* correlates with feature-binding. Kriegel appeals to neural synchrony in service of his self-representationalism, but there is nothing about neural synchrony itself that yields self-representation. The idea is simply that, in being a complex, M* will cease to exist whenever the firing rates in the relevant neural areas for each vehicle fail to be synchronised. The binding of vehicles here is not akin to writing 'I want to go to |===|' which, without a metalinguistic interpretation, is unintelligible. For present purposes, it will suffice to note that M* is closer to ['I am tired' + '私は寝たい'], which means 'I am tired, and I want to go to bed.' The '+' picks out the annotation relation between the icon and the title, allowing us to string together syntactically divergent forms without violating rules of language. (Thanks to Ali Rezaei for discussion here.)

right way. Maps do that. But so does the VF model, and without adverting to anything that we do not already need for independent considerations. What we call M* is to some extent irrelevant.

Earlier, we mentioned that non-visual perceptual experiences might have a form of presentational character. Suppose that auditory experiences do so. How should we treat such experiences? The answer is 'Much like visual experiences.' All auditory experiences will consist of a pure icon and a discursively formatted title, bound as a complex. The differences with visual experiences will just be the title ('Things that I am hearing') and the kind of information encoded in the icon. Whereas visual icons have parts that simultaneously 'coordinate' fine-grained values along hue, brightness, saturation, and spatial dimensions, auditory icons can be thought of as having parts that simultaneously 'coordinate' pitch, timbre, and volume, with location [Quilty-Dunn 2017: 160].

The Map Theory does not end with the VF model. Given the range of map types, there is room to grow. Instead of exploring other options, however, I'll close on this question: can a HOT theorist endorse the Map Theory?

4. HOT or HOP?

Is M* a HOT on Map Theory? My initial reaction is mild ambivalence: if non-relational HO theories are only tenable as non-relational HOP theories, so be it. However, others might not share this reaction, given hesitance towards HOP representations more generally [Roche 2013]. Few doubt that we *have* HOTs; what's up for debate is whether HOTs constitute consciousness. The good news is that, while the issue is delicate, there is a path for the HOT theorist.

M* is not familiar from folk psychology. Nor will M* subjectively seem like a thought. However, as evidence against HOT-compatibility, both measures are too austere. Our conception of thoughts shouldn't be beholden to folk psychology. Moreover, HOT theorists are clear that HOTs are *posits* that earn their keep by their explanatory power [Rosenthal 2005: 9]. We are seldom aware of HOTs. And, when we engage in introspection, they need not be revealed to us *as thoughts*. When S introspects, she will have a distinct HOT M** directed at M*. The content of the former need not be <I am *thinking* that I am seeing an apple>; it could just be <I am *aware* that I am seeing an apple>. So, it won't seem to S that her *HOT* is rich and fine-grained. What seems rich and fine-grained to S are the appearances conveyed by this HOT. These points can be adopted by the Map theorist. That M* does not seem like a thought is irrelevant.

Whether M* is a HOT or a HOP depends on how we mark the cognition-perception border. One view privileges format: perception is iconic, while cognition is discursive (for instance, Carey [2009: 8]). However, this is controversial. For example, Beck [2012] argues that we can have cognitive states that have analogue or iconic components in addition to discursive components. And Quilty-Dunn [forthcoming a] argues that we can have perceptual states that likewise have both iconic and discursive components. Thus, the proponent of the Map Theory might insist that the fact that M* has both discursive and iconic components tells us nothing about whether M* is cognitive or perceptual. Instead, we should look to non-representational aspects of mental architecture to delineate the border, like stimulus-dependence [Beck 2012, 2018]. While M* can be part of a causal chain that includes your visual system being in

causal contact with the object seen, it need not be. Nor is M* exhausted by elements that have the function of being stimulus-dependent: the representation of the self has no such function (cf. Beck [2018: 11–12]). By these measures, M* is cognitive.[25,26]

Disclosure Statement

No potential conflict of interest was reported by the author.

ORCID

Joseph Gottlieb http://orcid.org/0000-0001-9014-8487

References

Beck, J. 2012. The Generality Constraint and the Structure of Thought, *Mind* 121/483:563–600.
Beck, J. 2018. Marking the Perception–Cognition Boundary: The Criterion of Stimulus-Dependence, *Australasian Journal of Philosophy* 96/2 :319–34.
Berger, J. 2014. Consciousness Is Not a Property of States: A Reply to Wilberg, *Philosophical Psychology* 27/6: 829–42.
Block, N. 2011. The Higher-Order Approach to Consciousness Is Defunct, *Analysis* 71/3: 419–31.
Block, N. 2019. Empirical Science Meets Higher-Order Views of Consciousness: Reply to Lau and Brown, in *Blockheads! Essays on Ned Block's Philosophy of Mind and Consciousness*, ed. Adam Pautz and Daniel Stoljar, Cambridge, MA: The MIT Press: 199–214.
Bourget, D. 2017. Why Are Some Phenomenal Experiences 'Vivid' and Others 'Faint'? Representationalism, Imagery, and Cognitive Phenomenology, *Australasian Journal of Philosophy* 95/4: 673–87.
Braddon-Mitchell, D. and F. Jackson 1996. *Philosophy of Mind and Cognition: An Introduction*, Oxford: Blackwell.
Bronner, B. 2015. Maps and Absent Symbols, *Australasian Journal of Philosophy* 93/1: 43–59.
Brown, R. 2015. The HOROR Theory of Consciousness, *Philosophical Studies* 172/7: 1783–94.
Burge, T. 2010. *Origins of Objectivity*, Oxford: Oxford University Press.
Byrne, A. 1997. Some Like It HOT: Consciousness and Higher-Order Thoughts, *Philosophical Studies* 86/2:103–29.
Camp, E. 2007. Thinking with Maps, *Philosophical Perspectives* 21:145–82.
Carey, S. 2009. *The Origin of Concepts*, New York: Oxford University Press.
Cohen, M.A., D.C. Dennett, and N. Kanwisher. 2016. What Is the Bandwidth of Perceptual Experience? *Trends in Cognitive Sciences* 20/5: 324–35.
Dretske, F. 1981. *Knowledge and the Flow of Information*, Cambridge, MA: The MIT Press.
Dretske, F. 1995. *Naturalizing the Mind*, Cambridge, MA: The MIT Press.
Edwards, J. and D. Platchias 2018. Epistemic Warrants and Higher-Order Theories of Conscious Perception, *Pacific Philosophical Quarterly* 99/2: 343–64.

[25] If stimulus-dependence is the way to mark the cognition-perception border, then it is less concerning that maps and icons have significant limits in expressive power [Camp 2007] or perhaps don't involve predication [Rescorla 2009b]. Moreover, neither point has much independent force in the present context anyway. Suppose that maps cannot accomplish negation. That's fine. Experiences never *deny* anything. My experience doesn't tell me that the apple is *not* in a bowl, even if this is a face-value *implication* of my experience. Regarding predication, Rescorla [ibid.] uses the 'absence intuition' to claim that maps don't involve semantic predication. While this might be a pragmatic phenomenon [Camp 2007: 163; Bronner 2015], even if icons don't involve semantic predication, we might, with Burge [2010], still insist that they attribute properties to objects. It's unclear why the Map theorist would need anything more than this.

[26] For comments and discussion, I wish to thank audiences at Coastal Carolina University and Texas Tech University, Jacob Berger, Sam Coleman, Jessica Gottlieb, Saja Parvizian, and two anonymous reviewers and the editors of *AJP*. I owe special thanks to Ali Rezaei for countless marathon discussions about maps and HOT theory.

Farrell, J. 2018. Higher-Order Theories of Consciousness and What-It-Is-Like-Ness, *Philosophical Studies* 175/11: 2743–61.
Fodor, J. 2007. The Revenge of the Given, in *Contemporary Debates in Philosophy of Mind*, ed. Brian McLaughlin and Jonathan Cohen, Oxford: Basil Blackwell: 105–16.
Gennaro, R. 2012. *The Consciousness Paradox: Consciousness, Concepts, and Higher-Order Thoughts*, Cambridge, MA: The MIT Press.
Gottlieb, J. 2015. Presentational Character and Higher-Order Thoughts, *Journal of Consciousness Studies* 22/7–8: 103–23.
Gottlieb, J. 2016. Transitivity and Transparency, *Analytic Philosophy* 57/4: 353–79.
Gottlieb, J. 2018. Consciousness and the Limits of Memory, *Synthese* 195/12: 5217–43.
Gottlieb, J. 2019. The Collapse Argument, *Philosophical Studies* 176/1: 1–20.
Gottlieb, J. 2020. On Ambitious Higher-Order Theories of Consciousness, *Philosophical Psychology* 33/3: 421–41.
Green, E.J. and J. Quilty-Dunn forthcoming. What Is an Object File? *The British Journal for the Philosophy of Science*.
Im, H.Y. and S.C. Chong 2014. Mean Size as a Unit of Visual Working Memory, *Perception* 43/7: 663–76.
Knotts J.D., B. Odegaard, H. Lau, and D. Rosenthal 2019. Subjective Inflation: Phenomenology's Get-Rich-Quick Scheme, *Current Opinion in Psychology* 29/49–55.
Kriegel, U. 2009. *Subjective Consciousness: A Self-Representational Theory*, Oxford: Oxford University Press.
Kulvicki, J. 2015. Maps, Pictures, and Predication, *Ergo* 2/7.
Lau, H. and R. Brown 2019. The Emperor's New Phenomenology? The Empirical Case for Conscious Experience without First-Order Representations, in *Blockheads! Essays on Ned Block's Philosophy of Mind and Consciousness*, ed. Adam Pautz and Daniel Stoljar, Cambridge, MA: The MIT Press: 171–98.
Martin, M.G.F. 2002. The Transparency of Experience, *Mind and Language* 17/4: 376–425.
Picciuto, V. 2017. Keeping it Real: Intentional Inexistents, Fineness-of-Grain, and the Dilemma for Extrinsic Higher-Order Representational Theories, *Pacific Philosophical Quarterly* 98/4: 555–75.
Quilty-Dunn, J.K. 2017. Syntax and Semantics of Perceptual Representation, Ph.D. dissertation, City University of New York.
Quilty-Dunn, J. forthcoming a. Perceptual Pluralism, *Noûs*.
Quilty-Dunn, J. forthcoming b. Is Iconic Memory Iconic? *Philosophy and Phenomenological Research*.
Rescorla, M. 2009a. Cognitive Maps and the Language of Thought, *The British Journal for the Philosophy of Science* 60/2: 377–407.
Rescorla, M. 2009b. Predication and Cartographic Representation, *Synthese* 169/1: 175–200.
Roche, M. 2013. A Difficulty for Testing the Inner Sense Theory of Introspection, *Philosophy of Science* 80/5:1019–30.
Rosenthal, D. 2005. *Consciousness and Mind*, Oxford: Clarendon Press.
Rosenthal, D. 2009. Higher Order Theories of Consciousness, in *Oxford Handbook of the Philosophy of Mind*, ed. Brian McLaughlin, Axel Beckermann, and Sven Walter, Oxford: Clarendon Press: 239–52.
Rosenthal, D. 2010. How to Think about Mental Qualities, *Philosophical Issues* 20/1: 368–93.
Rosenthal, D. 2011. Exaggerated Reports: A Reply to Block, *Analysis* 71/3: 431–7.
Rosenthal, D. and J. Weisberg 2008. Higher-Order Theories of Consciousness, *Scholarpedia* 3/5: 4407.
Searle, J.R. 1983. *Intentionality: An Essay in the Philosophy of Mind*, Cambridge: Cambridge University Press.
Sperling, G. 1960. The Information Available in Brief Visual Presentations, *Psychological Monographs: General and Applied* 74/11: 1–29.
Treisman, A.M. 1996. The Binding Problem, *Current Opinion in Neurobiology* 6/2: 171–78.
Tye, M. 2000. *Consciousness, Color, and Content*, Cambridge, MA: The MIT Press.
Wandell B.A., S.O. Dumoulin, and A.A. Brewer 2007. Visual Field Maps in Human Cortex, *Neuron* 56/2: 366–83.
Weisberg, J. 2011. Abusing the Notion of What-It's-Like-ness: A Response to Block, *Analysis* 71/3: 443–48.
Wilberg, J. 2010. Consciousness and False HOTs, *Philosophical Psychology* 23/5: 617–38.