



Phenomenal Intentionality and the Problem of Representation

ABSTRACT: *According to the phenomenal intentionality research program, a state's intentional content is fixed by its phenomenal character. Defenders of this view have little to say about just how this grounding is accomplished. I argue that without a robust account of representation, the research program promises too little. Unfortunately, most of the well-developed accounts of representation— asymmetric dependence, teleosemantics, and the like—ground representation in external relations such as causation. Such accounts are inconsistent with the core of the phenomenal intentionality program. I argue that, however counterintuitive it may seem, the best prospect for explaining how phenomenal character represents is an appeal to resemblance.*

KEYWORDS: intentional content, phenomenal intentionality, representation, resemblance

What counts as an appropriate stopping place for explanation can shift dramatically. Consider the claim that intentional and phenomenal properties are identical. When naturalistic theories of representation such as Fodor's asymmetric dependence theory (1994), Dretske's indicator semantics (2002), and Millikan's teleosemantics (1987) were in vogue, it became appealing to explain the phenomenal by means of the intentional (Tye 2000). This first wave of 'intentionalism' offered to reduce the otherwise mysterious 'what it's like' to representational content.

Unfortunately, the program of naturalizing intentionality appears to have stalled, if not failed. Nothing much has changed since 2009, when William Lycan called 'materialist psychosemantics' a 'dismal failure' (2009: 551). The most promising relationships between mental states and their objects in the world seem, in the end, ill-equipped to do the job of tying representations to their *representanda*. Today it seems appealing to run the identification in the other direction. According to second wave intentionalism, phenomenology grounds intentionality (Loar 2003; Horgan and Tienson 2002). Instead of looking at the external relations in which a state stands, one should look to its phenomenal character. A family of theses now flying under the banner of the 'phenomenal intentionality' research program

I would like to thank Galen Barry, Brie Gertler, Ben Jantzen, and Harold Langsam for their comments on earlier versions of this paper and for many fruitful discussions. I would also like to thank two anonymous referees for this journal for their especially helpful criticisms.



(PI) has attracted considerable attention. Put roughly, the core idea is that there is a distinctive kind of aboutness constituted by what it's like to be in a mental state.

The PI research program is well on the way to becoming Kuhnian normal science. There is a fairly obvious problem, however, that has not gotten the attention it deserves: PI has no story to tell about *how* phenomenal character represents properties or objects in the world. The relation that ties states to what they represent is never explained. Nor does this seem to be an accident: the well-developed, if by now suspect, stories about representation all belong to the competing naturalistic model of intentionality. To appeal to any of the familiar naturalistic options would be to sacrifice the distinctive element of PI. But what else is on offer?

Maybe nothing. Perhaps the way in which what-it's-like to see a square represents a square is supposed to be an unanalyzable primitive. I explore the costs of this move in greater detail below, but the chief problem should be obvious. If our project is to zero in on just how 'intentionality is injected into the world', as Uriah Kriegel (2013b: 2) puts it, then we had better have something to say about how a state with such-and-such a phenomenal character represents what it does. In the absence of such an account, critics of PI will justifiably complain that the theory does little to illuminate its target phenomenon. In short, phenomenal character cannot be an appropriate stopping place unless we understand *how* it connects a mental state to something in the extramental world.

My strategy is simple. In the first section, I argue that representation is in fact problematic for PI, despite its tendency to take representation as a primitive. Even if one waives this argument, it remains the case that a theory of representation consistent with PI would, if grafted on to it, make for a more compelling view. The second section offers just such a theory. As strange as it sounds, I shall argue that PI is best served by invoking resemblance to explain how its chosen states get to be about their objects in the world. Those hostile to resemblance as a theory of representation are free to regard the arguments of this paper as a *reductio ad absurdum*: if PI cannot account for aboutness without appealing to resemblance, so much the worse for PI.

As it happens, I think the theory that results from combining PI with a resemblance account is worth pursuing. The third section goes beyond the other two in undertaking a defense of that theory. While I cannot hope to preempt every objection one might generate, I do hope to show that the resemblance-based account is not the obvious nonstarter most philosophers believe it to be. The argument of this third section is logically independent of the other two.

Some caveats are in order before I begin. First, I shall mainly be concerned with phenomenal intentionality in veridical perceptual cases. I want to know what is happening in the simplest possible scenarios, when a sentient subject is in the presence of a world of ordinary objects. Second, I am interested in a very narrow version of PI, one designed to account only for such simple cases. I take no position on whether all intentional content is phenomenal or whether all of its other forms can be derived from the perceptual kind.

I. The Representation Objection

Jane rests her hand on a mossy boulder. She feels the damp moss and the outline of the boulder. A moment later she thinks *the boulder obstructs my view of the waterfall*. Both mental states represent things outside of Jane herself. But they differ in important ways. States of the first kind do not have a propositional structure; as a result, they lack a truth value. States of the second kind have ‘intentional content’ in the sense that they can be true or false. To avoid confusion, I shall reserve ‘representation’ for states of the first, nonpropositional sort. Both representations and intentional contents can be said to ‘be about’ or ‘represent’ something, namely, properties or objects in the first case and states of affairs in the second.

These two states of Jane—her feeling the mossy boulder and thinking that it obscures her view—are our explananda. We must distinguish two questions: (1) What is it that imbues a state with aboutness? (2) What is it that makes this mental state about this or that object or state of affairs? The two questions need not get the same answer. Even if one finds phenomenal character a plausible answer to (1), that does not commit one to giving the same answer to (2). The representation objection, as I’ll call it, claims that PI so far has no good answer to (2).

Before developing that objection, we need a clearer notion of PI. First, PI claims that at least some intentionality is grounded in phenomenology, where this means, at a minimum, that it is in virtue of having such-and-such a phenomenology that a state represents what it does (Horgan and Tienson 2002: 520). Some (e.g., David Pitt 2004) will go further and explain this grounding relation by means of identity: there is nothing more to the representational content of a state over and above its phenomenology. Moreover, phenomenal intentionality is basic in the sense that it cannot be derived from some other form of intentionality.

Another distinctive feature is the narrowness of phenomenal intentionality: it does not depend for its nature on anything outside the subject. Phenomenal intentionality is supposed to be exhaustively accessible to the subject undergoing it. A brain in a vat and its embodied counterpart will enjoy the same phenomenology and hence the same representational and intentional states (here I am ignoring a variety of complications that surround the internalism/externalism debates; for a careful disentangling of these issues, see Gertler [2013]).

Within the PI camp, there are, of course, important divisions. Horgan and Tienson’s (2002) view encompasses all of cognition. As they see matters, there is a distinctive phenomenology associated with entertaining any proposition whatsoever. ‘What it’s like’ to think $2 + 2 = 4$ explains why one’s thought is about $2 + 2 = 4$. But not all views endorse what we might call cognitive phenomenology (for criticism, see Pautz 2013). One might want to claim instead that the fundamental source of intentionality is perceptual. Perhaps aboutness enters the world when sentient subjects begin experiencing their environments, and more sophisticated cognitive states inherit their aboutness from such experiences. Since both views will maintain that there is at least perceptual phenomenal intentionality, I’ll focus on this overlap region.

For our purposes, the core of PI is the claim that the aboutness of a mental state is not a function of its selectional or causal history but solely of its own nature.

In contrast to the external relations the naturalizing program explored—call them ‘tracking’ relations—what makes a state about its object is, on PI, a nonrelational property. To use an example of Kriegel’s: suppose you ‘undergo an experience with a squarish phenomenal character’ (2013b: 7). Simply in virtue of being in that state, PI tells us, you are thinking about a square object. We can say that aboutness is an external relation for tracking theories but an internal relation for PI.

As I use the terms, an internal relation is one that obtains solely in virtue of the intrinsic properties of the relata while an external one does not. Mass is (*modulo* some problems about the speed of light) an intrinsic property, weight an extrinsic one; *having more mass than* is an internal relation, *weighing more than* an external one. What makes for an intrinsic property is a subject of debate; see, e.g., Langton and Lewis (1998). But the intuitive notion of a property an object has regardless of what is going on outside of that object is clear enough for our purposes. The decisive innovation of PI is its shift from external relations to internal relations. Consider, for example, David Pitt’s claim:

[My theory] maintains that the intentional content of a thought is determined by its intrinsic phenomenal properties, not its relational properties. (2009: 120)

Here are Graham, Horgan, and Tienson:

Physically and apart from phenomenology, there is no ‘one, determinate, right answer’ to the question of what is the content of an intentional state. For... the content of each mental state is not determinately fixed once the physical facts (including perhaps physical facts about the internal-environmental linkages) are fixed. Fortunately, however, for the identity or determinate character of intentional content, content identity or determinacy is fixed phenomenally. (2007: 476)

There are two closely related points here we must pry apart. First is the claim that the subject can ‘just tell’ what she is thinking. Jane does not need to consult a neuroscientist to know that she is thinking *that the boulder is obstructing her view of the waterfall*. That thought wears its character on its face. This first point is largely uncontroversial, at least if we bracket questions about wide content. In this sense, aboutness is not a relation at all; the nature of the state makes it the case that it has the content that it does.

But of course, this is only half the story. Sometimes, there are states of affairs (or properties) that intentional contents (or their constituent representations) stand in some relation to. And that relation, whatever it is, is internal in the sense that it obtains, when and if it does, solely in virtue of the intrinsic properties of the relata.

This second claim is far from uncontroversial. In fact, it should seem *prima facie* implausible, a return to the sort of picture Hilary Putnam (1981: 3) derides

as ‘magical’. For the second point is not a claim about first-person accessibility; it is a substantive claim about representation.

Nor is it optional. If phenomenology is to fix the reference of intentional states, then the relation between the phenomenology and whatever it is about must be an internal one. An external relation, such those in the tracking family, would deprive phenomenology of its alleged explanatory power. After all, nearly *any* state of anything can in principle track any other. By contrast, PI claims that it is the nature of phenomenal character and the states that have it that explains aboutness. Moreover, an external relation would presumably block subjects from using introspection to discern the narrow contents of their own thoughts.

PI is committed, then, to the claim that representation is an internal relation. But what relation, exactly? As far as I can tell, the proponents of PI have not told us. Let me distinguish this ‘representation objection’ from a closely related one. Adam Pautz argues that adherents of PI ‘tend to ignore the whole hard problem of naturalizing the mind’ (2013: 195). Although I agree—and will do the same—that is not the issue here. Instead, the objection is that PI does not so far have a substantive explanation of representation. That explanation need not be naturalistic, whatever one precisely means by that. But without *some* story to tell, PI is significantly less attractive than it at first seems.

Perhaps proponents of PI see no need to offer such a story. Here is one line of thought they might have in mind. Suppose that the only kind of aboutness is propositional and hence that we need only explain how intentional content is tied to states of affairs. If so, there is no mystery about this tie: a state whose phenomenal character can be expressed as *there is a boulder under my hand* represents that state of affairs just in case there is a boulder under my hand. In other words, representation is truth.

The chief problem with this maneuver is that it fails to illuminate intentional content. Even those who agree that ‘what it’s like’ somehow grounds ‘what it’s of’ might want to be told how that grounding is accomplished. But there are a host of other problems as well. None is a knock-down argument against the representation-as-truth maneuver, but each adds to the costs that would have to be paid.

First, suppose my mental state presents the rock as being covered in moss when it isn’t. Is that a false thought about a moss-less rock or a thought about a mossy rock that lacks a truth value because its putative subject—the mossy rock—fails to exist? The best answer is the first, but it is unclear how the present proposal is entitled to it. If the state’s intentional content is false, there seem to be no grounds for saying that it is *about* or represents the state of affairs that makes that content false.

The natural reply is that the thought is about the rock in question. It just falsely predicates *mossy* of the rock. But the proposal under review has no way of tying subpropositional elements like *rock* to their objects in the world. Again, the proposal has it that representing works at the propositional level. The relata then are not simple mental items like *rock* on the one hand and rocks on the other; they must be propositional contents such as *the rock is mossy* and the states of affairs, if such there be, that answer to them.

Second, if thoughts such as *the boulder is mossy* are to play a justificatory role, they cannot be unanalyzable primitives. Instead, they must be composed of simple elements that can recur in other thoughts such as *the boulder is probably damp*. *The boulder* in both thoughts must reach out to the same boulder in the world. This kind of argument is familiar from Fodor (1994). I realize it is controversial. Even so, it is reasonable to ask of a theory of intentionality that it not rule out the possibility of explaining a propositional thought in terms of its constituent parts.

Let me sum up the dialectic so far. PI grounds intentional content in phenomenal character. We then have to wonder whether intentional content—where this is propositional and admits of a truth value—is the basic unit of aboutness. If so, then PI can help itself to an account of representing as truth or satisfaction. I argued that this is an inappropriate stopping place. Moreover, that account brings in its train a host of unattractive consequences.

The most promising version of PI will, then, have to tie subpropositional elements to their objects in the world. Note that not every such element needs to be anchored in the extramental world; syncategoremata such as ‘is’ and ‘as’ will presumably merit a very different treatment. Nevertheless, we have seen that *boulder*, whether it occurs in Jane’s thought in the context of a proposition or on its own, is a representation in its own right.

At this stage, the representation objection becomes more pressing. In virtue of what does what-it’s-like for Jane to touch the boulder represent the boulder and not the waterfall? As always, we have the option of declaring it *sui generis* and moving on. If there is another option, we should take it. Such an alternative would have to start by explaining representation rather than intentional content. And it would have to appeal to an internal relation. Happily or not, there is exactly one possibility left: resemblance.

2. Representation as Resemblance

Few notions have attracted as much contumely as resemblance has. Tyler Burge (1986: 128) claims that it is ‘an intuitively powerful but primitive idea’; Nelson Goodman (1972: 437) calls it ‘a pretender, an impostor, a quack’, always turning up ready to solve problems it can’t (see Isaac 2013 for a catalog of further abuse). And yet for much of philosophy’s history, it was precisely the resource philosophers used to explain representation (see *De Anima* in Aristotle [1984: vol. 1] as well as Lagerlund [2011]; for the use of resemblance in historical figures, see Watson [1995] and Jacobson [2008] and [2013]). We have to be clear about what this ancient, if not primitive, idea is. It is sometimes said that everything resembles everything else, in some respect or other. Gloucester cheese and the square root of 75 resemble each other in that both have just been used in an example. Although quite true, this claim is true in an extremely weak sense of ‘resemblance’. The idea worth attacking (and defending) is what we might call internal resemblance, that is, resemblance in virtue of the intrinsic properties of the relata. In what follows, I shall always use the term in that restricted sense.

Although PI so far has not produced anything more than a few hints about the nature of representation, following those hints leads us to resemblance. Consider, for example, Kriegel's statement of Charles Siewert's (1998) argument for the existence of phenomenal intentionality:

Suppose you undergo an experience with a squarish phenomenal character. If nothing around you is square, your experience is assessable as inaccurate.... Thus phenomenal character can bring in its train accuracy conditions. Since having accuracy conditions is an intentional property, it appears that at least some phenomenal character can guarantee intentional properties. (Kriegel 2013b: 7)

It is entirely possible that 'squarish phenomenal character' is short for 'a state whose intentional content is true just in case there is a square in the environment'. If so, we are back where we started: talk of squarish phenomenal character does nothing to illuminate representation.

Suppose, instead, that we take 'squarish phenomenal character' at face value. What, then, would be the relationship between the square object and the squarish phenomenal character? A natural answer is that both, in some way, instantiate squareness. After all, what is being invoked to explain the intentionality of the phenomenal state is not its simply having *some* phenomenal character or other, but having precisely *this* character, namely, squarishness. And when two things exemplify the same property, they can be said to resemble each other in that respect. (Trope theorists are invited to substitute their own analysis of resemblance. The precise analysis of resemblance makes no difference to my account, so long as it remains an internal relation.) I am not claiming that Kriegel or Siewert would themselves accept this gloss. The possibility that they might be best served by doing so remains.¹

Resemblance also fits neatly with Michelle Montague's (2013) 'matching' view. Montague argues that we achieve perceptual access to a material object when we represent enough of its properties. But how do we represent those properties in the first place? I would argue that the best answer is that we do so by instantiating them ourselves in some way. Admittedly, this view is strange, at least on first hearing, but it is familiar to historians of philosophy (it runs through Aristotle, Aquinas, and Descartes, for example), and has recently been resurrected (see, e.g., Harold Langsam 2011: 63f.) Obviously neither we nor our experiences can instantiate the spatial properties of material objects in the same way those objects do. Instead, the resemblance theory has to appeal to what Descartes called 'objective being'. Much

¹ My thoughts here are very much indebted to the stimulating work of Anne Jaap Jacobson, especially her article 'What Should a Theory of Vision Look Like?' (2008) and her book *Keeping the World in Mind*. (2013). There are significant differences, of course: Jacobson is chiefly concerned with unearthing a theory of representation that squares with current work in cognitive science. Her examples of representational resemblance tend to focus on patterns of neural stimulation and so on. By contrast, I am working entirely at the phenomenal level. Our proposals are in principle compatible.

remains obscure here. But the metaphysics of this kind of instantiation are a further set of issues on which I can remain largely neutral.

There is one further possibility worth mentioning. Someone might argue that there is an isomorphism between the contents of visual experience and the physical objects that experience represents. Very roughly, the idea would be that at least some of the relations among the phenomenal ‘points’ that make up apparent shapes are preserved among the parts of the physical objects that are experienced. Such a view is to be found, in a much more sophisticated form, in the work of Chris Swoyer (1991 and 1995). Despite appearances, this view should still count as appealing to an internal relation to ground representation. What flanks the representation relation will be relations among the constituents of each relatum. If an internal relation is one that obtains solely in virtue of the properties of the relata, then it does not seem to matter whether those properties are themselves relations so long as they hold within, as it were, the two objects. Hence I take the isomorphism account to be a possible option within the resemblance family, rather than a competitor.

Having bracketed the question of naturalism, I take it as no objection that the proposal fails to meet that demand. Like many defenders of PI, I am happily shirking my physicalist duties. And I am helping myself to the notion of properties being instantiated in experience. But note that doing so does not rob the proposal of its explanatory power. In fact, as long as our goal is, with PI, to explain aboutness by means of an internal relation, resemblance is the only game in town.

But how exactly is resemblance supposed to help? Recall that part of the evidence for phenomenal intentionality is the accuracy condition implicit in phenomenology. But merely undergoing an experience with a squarish character is not enough to represent an object in the world as being square. As the moderns liked to put it, simple ideas are neither true nor false.

I think this consideration shows that PI needs an additional mechanism, beyond resemblance, to explain representation. The basic idea is due to Descartes (1984), but has recently been brought up to date by Katalin Farkas (2013). On Descartes’s view, we think of our sensations as ‘of’ an external world only because, at an early stage of development, we gained the ability to move, to twist ‘around aimlessly in all directions in... random attempts to pursue the beneficial and avoid the harmful’ (1984, vol. I: 219). Only at this point does the child realize that the objects it pursues or avoids have an existence outside itself. Similarly, Farkas argues that experiential states on their own do not ‘point’ to anything outside themselves. In this respect, they are like pain, on some accounts. Instead, what explains the object-presenting feature of experience is the structure of such experience. As Farkas puts it, it is ‘the way these [phenomenal] features hang together and respond to movement and inquiry’ that accounts for our taking them to be of objects (2013: 100). The phenomenal features are not piled on top of each other at random. They stand in systematic connections with each other and with the actions of the subject.

This Cartesian move seems to violate transparency. G. E. Moore, himself a sense datum theorist, might have been the first to claim explicitly that ‘when we try to introspect the sensation of blue, all we can see is the blue: the other element is as it were diaphanous’ (1903: 25). Nothing about such apparent transparency commits us to any particular story about representation (as Tye 2000: 47 in effect points

out). In fact, that our experiences should be transparent makes sense, if they have the kind of structure Descartes and Farkas point to.

One might also worry that I have given away the store by appealing to external relations to explain the object-presenting nature of experience. Note, first of all, that those relations hold among phenomenal states, not between such states and the world beyond the mind. What is more important, however, is that the Cartesian account does not use the structure of experience to answer the chief question we're concerned with. We've been worried about question (2): in virtue of what does this phenomenal character represent this object or property? Instead, the Cartesian account is best construed as an answer to the adjacent question (1): what makes a state a representation at all?

Let me close this section by illustrating the current proposal. Suppose Jane is standing in front of black cube on a white floor. Her visual experience instantiates the property of cube-ishness. It does so, of course, by also instantiating color. Without committing ourselves on the ontology of color, we can say that the experienced color itself has a shape, and it is this shape that resembles the shape of the cube itself. Just as PI requires, it is the phenomenal character of Jane's experience—its cube-ishness—that grounds the aboutness of the mental state that has that character.

3. A Defense of the Resemblance-based Account

Someone might well accept my arguments so far and conclude: so much the worse for PI. If the only way for PI to answer the representation objection is by invoking resemblance, then PI must be rejected. In this section, I argue that this conclusion is too hasty. The typical objections philosophers lodge against resemblance theories are not nearly as powerful as they seem. What is more, resemblance *cum* phenomenal intentionality provides resources for answering those objections that resemblance on its own lacks.

- (i) Ubiquity. For any *a*, there will be a large—though not infinite—number of intrinsic properties it possesses, and, plausibly, there will be at least one *b* such that *b* has at least one of those properties. Resemblance is all over the place; representation isn't. So they cannot be the same thing. (See, e.g., Burge 1986: 128 and Meyering 1989: 3)

The best reply is a concession. One car resembles another; it would be absurd to say that each therefore represents the other. Representation cannot *just* be resemblance. But we have to keep in mind which question we're trying to answer. Throughout we've had to distinguish the question (1) what makes a state a representation from (2) what makes it represent this rather than that? To be taken seriously, the resemblance theory has to be an answer to (2) rather than to (1). A state has to do more than resemble something to become a representation. In

answering (1), the resemblance view has two resources. First, it can hold that only states with phenomenal character are representations. Although not everyone aligned with PI will accept this, some will. More plausibly, the resemblance view can point to the second step in achieving representation: the subject's experiences have to exhibit a certain kind of structure before they become representations. Either of these moves (and there are surely more) would make representation much more scarce than resemblance.

But there is a stronger version of the ubiquity problem lurking: what Satoshi Watanabe (1969: 376) calls the 'Ugly Duckling Theorem'. Suppose that the degree of similarity between *a* and *b* is fixed by the number of predicates that can be applied to them. Watanabe argues that any object will bear precisely the same degree of similarity to any other object. The ugly duckling is just as similar to the fetching cygnet as is any other cygnet. Applied to the present theory, the claim would be that any given mental state will resemble any object to the same degree as any other object. Despite appearances, the experience with the cube-ish phenomenal character does not resemble the cube any more than it does the sphere.

Before meeting the objection directly, let me offer some reasons for being suspicious of the theorem, at least as applied to the present view. It is significant that Watanabe's argument for the theorem appeals to predicates, not properties. And it appeals to predicates that few would agree map genuine properties. Consider Watanabe's (1985: 75) version of the argument: 'To persuade that two objects are similar, it is natural to enumerate the properties that are commonly owned by the objects.... For such an argument it is only fair for the arguers to be allowed to use all possible predicates that are applicable to the objects'. Note that Watanabe begins by talking about properties but shifts to predicates. Allowing *any* applicable predicate at all to count toward similarity makes proving the theorem all too easy, as does permitting external points of resemblance (say, 'being north of the Equator'). But the view I am defending appeals to internal resemblance in respect of properties, taken as real members of the mind-independent world, not in respect of mere predicates. Hence it is not obvious that Watanabe's theorem threatens the resemblance view, properly understood.

But let us waive these considerations. What is needed to answer the objection directly is some way of weighting the properties in respect of which the similarity obtains. And here is precisely where the resemblance account's marriage to PI pays dividends. The salient features are precisely those intrinsic properties that capture our attention. One notices the shared shapes of the comely cygnets and notes their difference from that of the duckling. There are no doubt plenty of similarities that we simply do not attend to even if they somehow feature in our phenomenal experience. In fact, something very like this reply is common currency in cognitive science. As Stevan Harnad puts it, in his reply to Watanabe's argument: 'But of course our sensorimotor systems do not give equal weight to all features: they do not even detect all features. And among the features they do detect, some (such as shape and color) are more salient than others (such as spatial position and number of feathers)' (2005: 32). Replace 'sensorimotor' with 'phenomenally perceptual', and you have a fair approximation of this point. In short, the Ugly Duckling theorem tells against the present account only if it assumes that there is no way to

privilege some of the features shared by experiences and objects over others. And that assumption is false.

- (ii) Symmetry. Representation is asymmetrical: if *a* represents *b*, then, at least typically, *b* doesn't represent *a*. Resemblance, by contrast, is symmetrical: each Winklevoss twin resembles the other.

Here again the difference between questions (1) and (2) makes all the difference. True, resemblance is symmetrical and representation is not. So while my experience of a cube resembles the cube and vice versa, it would be absurd to say that the cube *represents* my experience. *Why* it is absurd depends on just how we answer question (1). Suppose that only mental states can be representations. Then the reply is clear: the representation relation is asymmetrical not because resemblance goes in only one direction—it doesn't—but because the cube is not a representation at all.

- (iii) Singularity. Suppose I am looking at a cube; my experience and the cube share a property, cube-ishness. But cube-ishness is shared by every cube in the actual world and out of it. How does my experience now get tied to this cube and not all those others, actual and possible?

It's important not to exaggerate the threat posed by objection (iii). If we limit ourselves to features of experience that are accessible from the first-person perspective, then any view will have to tolerate some degree of unspecificity. And one of PI's core commitments is the subjectivity of content. Such specificity as there is must be tied to introspectible features such as the distribution of color, size, shape, and position. In principle, a single experience can always be an experience of more than just the particular environment in which it happens to occur. If this is a problem, it is one any adherent of PI will have to live with.

But remember that the resemblance account under consideration can use the resources of PI. It is true that, taken singly, a phenomenal experience that instantiates, say, black-cube-ishness will likely be indistinguishable from another such experience. Fortunately, we hardly ever take our experiences singly. Jane's phenomenal world nearly always includes more than a single phenomenal property, such as the feel of the boulder; she also experiences the color of the moss and the variegated shading of the trees. It is not as if it is up to just one of these aspects of her experience to tie that experience to the world. Instead, all of them can work together, at least to the extent that they are objects of her attention. If we keep this fact in mind, the worry that Jane's experience is insufficiently singular should at least be greatly reduced or even vanish. For the number of subjectively identical 'experiential scenes' is far fewer than the number of subjectively identical squareness-es. The final step is achieved by looking not just at a single slice in time, but at Jane's experience over a period of seconds or even minutes. For part of what knits Jane's experiences over time together with her environment is surely the fact

that the properties instantiated in the objects around her, both synchronically and diachronically, are also instantiated in her experiences.

Finally, note that objection (iii) should not be mistaken for the claim that singular thoughts are such that they depend for their existence on the existence of their objects (see McDowell 1998). Whether or not there are such object-dependent thoughts, anyone committed to the narrow intentionality of phenomenal states will deny that perceptual states are among them.

- (iv) Any story about representation has to account for *misrepresentation*. But if *a* resembles *b*, *a* cannot misrepresent *b* on pain of not being a representation of *b* at all.

Suppose we take resemblance to be (literal) property-sharing. If *a* resembles *b* in virtue of the fact that *Fa* and *Fb*, *a* represents *b* just as *b* in fact is, at least in respect *F*. It seems that the only way for *a* to misrepresent *b* is for it to lack *F*, but since representation is resemblance, then it is no longer a representation of *b*.

The present view has a ready response. Suppose the felt shape of Jane's experience fails to resemble the shape of the boulder underneath her hand, due to some sort of haptic illusion. There will still be *other* qualities of the boulder that her experience gets right, such as its distance from her, its contact with her hand, and so on. The degrees of misrepresentation correspond to the number of respects in which Jane's experience resembles the world.

But now suppose Jane's experience is as of a pink elephant (to borrow Montague's example [2013: 39]). Now none of the qualities instantiated in her experience resembles the qualities of the boulder. In such a case, there's no object in the world that she's misrepresenting. That in fact sounds like the right result, not a counterexample.

One thing any story about experience wants to preserve is this: when Jane hallucinates, she can undergo an experience that is the same, from her perspective, as one she would undergo were the world as she represents it. The present view is entitled to this counterfactual, which is made true by the intrinsic nature of her experience. Now, Jane might hallucinate in a way that just happens to match the way the world is. On the view we're considering, she is not misrepresenting the world in that case; nevertheless, we might want to deny that her experience is veridical since it is not caused by the world itself. Not all epistemic failures are failures to represent the world as it is.

- (v) If resemblance is property-sharing, then typical visual experiences will not resemble, and hence not represent, material objects. As the visual arts demonstrate, the apparent spatial properties of objects very rarely coincide with their real ones. The properties instantiated in Jane's experience are viewer-dependent properties that cannot be identified with the viewer-independent properties of real objects.

Imagine Jane watching a hawk flying overhead. As it comes closer, it takes up more and more of Jane's visual field. And yet Jane's experience typically doesn't represent the hawk as growing in size. The 'typically' is important here, for with attention we can notice the shift in experienced size.

To use Gibson's (1950) terms, my claim is that Jane can experience both a visual world and a visual field. It's very unlikely that she experiences both at the same time, all the time. In fact, as with figure/ground illusions, it might be impossible for her to instantiate both sets of properties at once. In typical conditions, Jane's experience will amount to experience of a visual world. Only in special circumstances (such as when reflecting on optical illusions or trying to draw something) will Jane experience her own visual field.

How does any of this help with our present problem, namely, the seeming mismatch between viewer-dependent and viewer-independent properties? The objection assumes that the *only* features instantiated in experience are dependent on the viewer. What we've just seen, however, is that there is another layer to experience. There are also features instantiated in experience that do resemble the mind-independent world. What Jane attends to, in the usual case, is the visual world, which co-instantiates many of the properties of the actual world.

I have no doubt other objections will spring to the reader's mind. Answering all of them would require fully developing the resemblance account, which I do not have space to do here. And like other participants in PI, I have sailed blithely past the whole issue of naturalism. My goal, however, is to open a spot, however small, for the resemblance account in the logical space of the debate. Readers who think there is no such spot available are invited to regard the argument of the paper as a *reductio* of PI. For whether it succeeds or fails, the resemblance account is the best answer PI has to the representation objection.

4. Conclusion

If I am right, the phenomenal intentionality program as it stands now does not provide us with an appropriate stopping place in our search for an explanation of aboutness. So far, those pursuing the program have either rejected the need for a robust theory of representing or simply ignored the question altogether. I have been arguing that this is the wrong approach. Without such a theory, the program will remain little more than a promissory note.

Unfortunately, the most prominent theories of representation all lie in enemy territory. A core tenet of PI is the narrowness of representation—if it is, or is fixed by, phenomenal character, matters could hardly be otherwise. So none of the familiar candidates, from asymmetric dependence to teleosemantics, can be grafted onto PI. Any of those tracking theories makes representation a function of a state's external relations, whether historical or modal.

The narrowness of content means that it has to be fixed by the internal relations in which phenomenal character stands. As I argued above, anything less knocks phenomenal character from its throne. For any external relation could be satisfied

by nearly any state. PI claims that a state has precisely the intentional content it does in virtue of its phenomenal character. Once this is accepted, the representing relation has to be understood as internal.

As far as I can tell, the only internal relation that can explain phenomenal intentionality is resemblance. Only resemblance can respect the need of PI to appeal to ‘what it’s like’ to ground aboutness. Thus, the chief result of sections one and two is a conditional: if PI is true, representation must involve resemblance. At this stage, even those who agree with my arguments thus far are liable to deploy *modus tollens*: a view that requires a resemblance theory of representation is not likely to win many converts. In [section 3](#), I went further and tried to establish that resemblance’s ill-repute is undeserved. If the resemblance view fails, it will not be for any of the usual reasons.

In closing, let me note that there are sound motivations, entirely independent of those given above, to try to resurrect a resemblance theory. For if tracking theories are in fact doomed, anyone wishing to preserve sensory representation would be advised not to dismiss resemblance too quickly. That is true whether one locates representation in phenomenal character or not.

WALTER OTT
UNIVERSITY OF VIRGINIA
wrottjr@gmail.com

References

- Aristotle. (1984) *The Complete Works of Aristotle*. Edited by J. Barnes. 2 vols. Princeton, NJ: Princeton University Press.
- Burge, Tyler. (1986) ‘Cartesian Error and the Objectivity of Perception’. In P. Pettit and J. McDowell (eds.), *Subject, Thought, and Context* (Oxford: Clarendon Press), 117–36.
- Descartes, René. (1984) *The Philosophical Writings of Descartes*. Edited by John Cottingham, Robert Stoothoff, and Dugald Murdoch. 3 vols. New York: Cambridge University Press.
- Dretske, Fred. (2002) ‘A Recipe for Thought’. In D. Chalmers (ed.), *Philosophy of Mind*. (Oxford: Oxford University Press), 491–99.
- Farkas, Katalin. (2013) ‘Constructing a World for the Senses’. In Uriah Kriegel (ed.), *Phenomenal Intentionality* (Oxford: Oxford University Press), 99–115.
- Fodor, Jerry. (1994) ‘A Theory of Content, II: The Theory’. In S. Stich and T. Warfield (eds.), *Mental Representation: A Reader* (London: Blackwell), 180–220.
- Gertler, Brie. (2013) ‘Understanding the Internalism-Externalism Debate: What is the Boundary of the Thinker?’ *Philosophical Perspectives*, 26, 51–75.
- Gibson, James J. (1950) *The Perception of the Visual World*. Boston, MA: Houghton Mifflin.
- Goodman, Nelson. (1972) ‘Seven Strictures on Similarity’. In Nelson Goodman, *Problems and Projects* (Indianapolis, IN: Hackett), 437–46.
- Graham, George, Terence Horgan, and John Tienson. (2007) ‘Consciousness and Intentionality’. In S. Schneider and M. Velmans (eds.), *The Blackwell Companion to Consciousness* (London: Blackwell), 468–84.
- Harnad, Stevan. (2005) ‘To Cognize is to Categorize’. In Henri Cohen and Claire Lefebvre (eds.), *The Handbook of Categorization in Cognitive Science* (New York: Elsevier), 20–45.
- Horgan, T., and J. Tienson. (2002) ‘The Intentionality of Phenomenology and the Phenomenology of Intentionality’. In D. Chalmers (ed.), *Philosophy of Mind: Classical and Contemporary Readings* (Oxford: Oxford University Press), 520–33.

- Isaac, Alistair. (2013) 'Objective Similarity and Mental Representation'. *Australasian Journal of Philosophy*, 91, 1–22.
- Jacobson, Anne Jaap. (2008) 'What Should a Theory of Vision Look Like?' *Philosophical Psychology*, 21, 641–55.
- Jacobson, Anne Jaap. (2013) *Keeping the World in Mind*. New York: Palgrave Macmillan.
- Kriegel, Uriah. (2011) *Sources of Intentionality*. Oxford: Oxford University Press.
- Kriegel, Uriah, ed. (2013a) *Phenomenal Intentionality*. Oxford: Oxford University Press.
- Kriegel, Uriah. (2013b) 'The Phenomenal Intentionality Research Program.' In U. Kriegel (ed.), *Phenomenal Intentionality* (Oxford: Oxford University Press), 1–26.
- Lagerlund, Henrik. (2011) 'Mental Representation in Medieval Philosophy.' In E. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. Stanford, CA: Stanford University. Available at: <http://plato.stanford.edu/entries/representation-medieval/>.
- Langsam, Harold. (2011) *The Wonder of Consciousness*. Boston, MA: MIT Press.
- Langton, Rae, and David Lewis. (1998) 'Defining "Intrinsic"'. *Philosophy and Phenomenological Research*, 58, 333–45.
- Loar, Brian. (2003) 'Phenomenal Intentionality as the Basis for Mental Content'. In M. Hahn and B. Ramberg (eds.), *Reflections and Replies: Essays on the Philosophy of Tyler Burge* (Cambridge, UK: Cambridge University Press), 229–58.
- Lycan, William. (2009) 'Giving Dualism its Due'. *Australasian Journal of Philosophy*, 87, 551–63.
- McDowell, John. (1998) *Meaning, Knowledge, and Reality*. Cambridge, MA: Harvard University Press.
- Meyering, Theo C. (1989) *Historical Roots of Cognitive Science*. New York: Springer.
- Millikan, Ruth G. (1987) *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.
- Montague, Michelle. (2013) 'The Access Problem'. In U. Kriegel (ed.), *Phenomenal Intentionality* (Oxford: Oxford University Press), 27–48.
- Moore, G. E. (1903) 'The Refutation of Idealism'. In G. E. Moore, *Philosophical Studies* (London: Routledge), 1–30.
- Pautz, Adam. (2013) 'Does Phenomenology Ground Mental Content?' In U. Kriegel (ed.), *Phenomenal Intentionality* (Oxford: Oxford University Press), 194–234.
- Pitt, David. (2004) 'The Phenomenology of Cognition, or, What is It Like to Think That P?' *Philosophy and Phenomenological Research*, 69, 1–36.
- Pitt, David. (2009) 'Intentional Psychologism'. *Philosophical Studies*, 146, 117–38.
- Putnam, Hilary. (1981) *Reason, Truth, and History*. Cambridge, UK: Cambridge University Press.
- Siewert, Charles. (1998) *The Significance of Consciousness*. Princeton, NJ: Princeton University Press.
- Swoyer, Chris. (1991) 'Structural Representation and Surrogative Reasoning'. *Synthese*, 87, 449–508.
- Swoyer, Chris. (1995) 'Leibnizian Expression'. *Journal of the History of Philosophy*, 33, 65–99.
- Tye, Michael. (2000) *Consciousness, Color, and Content*. Cambridge, MA: MIT Press.
- Watanabe, Satosi. (1969) *Knowing and Guessing: A Quantitative Study of Inference and Information*. New York: Wiley.
- Watanabe, Satosi. (1985) *Pattern Recognition: Human and Mechanical*. New York: Wiley.
- Watson, Richard. (1995) *Representational Ideas from Plato to Patricia Churchland*. New York: Springer.