



# Structural and indicator representations: a difference in degree, not kind

Gregory Nirshberg<sup>1</sup> · Lawrence Shapiro<sup>1</sup>

Received: 3 April 2019 / Accepted: 7 January 2020 / Published online: 16 January 2020  
© Springer Nature B.V. 2020

## Abstract

Some philosophers have offered *structural* representations as an alternative to indicator-based representations. Motivating these philosophers is the belief that an indication-based analysis of representation exhibits two fatal inadequacies from which structural representations are spared: such an analysis cannot account for the causal role of representational content and cannot explain how representational content can be made determinate. In fact, we argue, indicator and structural representations are on a par with respect to these two problems. This should not be surprising, we contend, given that the distinction between indicator and structural representations is better conceived as one involving degree rather than kind.

**Keywords** Structural representation · Indication · Disjunction problem · Content determinacy

## 1 Introduction

In recent years, some philosophers have promoted the idea of *structural representations* as a response to challenges that indicator-, or detector-based, theories of representation face (Opie and O'Brien 2004; O'Brien 2016; Gładziejewski and Miłkowski 2017).<sup>1</sup> More specifically, they believe structural representations endow semantic content with a causal role that *indicator representations* cannot, and structural representations enjoy a content determinacy that, according to a standard line of criticism, indicator representations do not. We believe that structural and indicator

---

<sup>1</sup> The precise target of the objections that structural representationalist raise is actually not so clear. They criticize “dyadic” approaches, but also theories like Dretske’s, which (as we will see) are best construed as a “triadic” approaches. Likewise, they raise objections to causal theories of representation, but, again, seem to think that these objections apply to theories like Dretske’s, which is not a causal theory.

---

✉ Lawrence Shapiro  
lshapiro@wisc.edu

<sup>1</sup> Department of Philosophy, University of Wisconsin – Madison, Madison, WI 53706, USA

representations are on a par with respect to these problems. At the core of representation is correspondence, and insofar as both structural and indicator representations exemplify correspondence, they are of a piece. Indeed, we believe that indicator representations are simply a limiting case of structural representations. This means that structural representations and indicator representations share a common fate: one is as causally efficacious as the other; one as content-determinate as the other. We begin our defense of these claims with a presentation of structural and indicator representations. We then explain why both sorts of representations are equally bothered by the causation and determinacy problems. This will lend support to our final argument that both types of representation are a species of the same genus.

## 2 Structural and indicator representations

As structural representationalists frequently note (Opie and O'Brien 2004; O'Brien 2016) representation involves three features.<sup>2</sup> First is the vehicle of representation, i.e. the object that carries representational content. The red heart on the Valentine card, for instance, is a representational vehicle. It is a physical object. Its meaning—its content—is the second feature of a representation. The red heart is the representational vehicle, and *love* is its content. Just how a content is assigned to a vehicle is, as we shall see, a matter of controversy, and no doubt differences exist in how artifacts, like red hearts, acquire their meaning and how natural vehicles, like brain states, acquire theirs, but at the very least we can say that some kind of regular correspondence must exist between the vehicle and its content. Finally, the third feature of representation involves what Peirce (Hardwick 1977), and following him, Von Eckardt (1993) called *interpretation*. Because of the implicit suggestion that interpretation involves mental activity of some sort, thus eliciting charges of circularity when proffering an analysis of *mental* representation, we prefer more neutral terms like *use*, or, to adopt Gładziejewski and Miłkowski's suggestion (2017), *exploitability*. This final component of representation is necessary, for it constrains the set of possible contents. The red heart corresponds to many things. Suppose, for instance, that red hearts became icons in the fifteenth century, in which case they correspond to a time period following the fourteenth century. But the sign means *love*, and not *you are living in a time period following the fourteenth century*, because recipients of Valentine cards think *love* when they see a red heart. They do not think "I am living in a time following the fourteenth century." The red heart is *used* to mean the former and not the latter.

Although the explication above focuses on an artifact, a few tweaks adapt it to natural, or non-convention-based, representations. A brain state too can serve as a representational vehicle—a vehicle, perhaps, for a belief. The content of this vehicle may be something like *The Big Lebowski is playing at the Majestic Theater tonight*,

<sup>2</sup> This claim is not unique to structural representationalists (see, e.g., Bechtel 1998; Dennett 1982); which we take as additional evidence that structural representations do not differ in kind from indicator representations.

and the exploitative act that singles out this content from among the other items with which the brain state might correspond would be the exercise of various behavioral dispositions, e.g. uttering “The *Big Lebowski* is playing at the Majestic Theater tonight,” buying tickets to the show, dressing as one’s favorite Lebowski character, etc.<sup>3</sup>

An indicator representation involves a vehicle, R, that, at some point in its history, due to a nomic regularity, stands in correspondence to its content, C. Dretske, the principal architect of indicator semantics, typically expressed the relation of indication in probabilistic terms, so that the presence of R indicates C if and only if  $\Pr(C|R)=1$  (Dretske 1988).<sup>4</sup> A familiar example of such an indicator is a cell in the frog’s visual system that indicates the presence of flies. The probability of a fly’s presence, given the activation of the cell, is 1. Also true, however, is that on occasions following the “recruitment” of the indicator for purposes of detecting the presence of flies, it may fire in the presence of a some non-fly, like a bee bee. This creates a problem: what does the indicator mean? Does it mean *fly*, or does it mean *fly or bee bee*? We shall return to this question below, but for now we should note that some philosophers, especially Dretske (1988), look to teleological function to answer this question. The activation of the cell represents flies, not flies or bee bees, because natural selection recruited the cell for this purpose.<sup>5</sup> More specifically, natural selection resulted in a mechanism in which the cell was connected to motor neurons that cause the frog to snap its tongue at passing flies. This constitutes the third, exploitative, aspect of representation: activation of the cell disposes the frog to certain behaviors, e.g. fly-aimed tongue snaps. The triadic analysis, introduced above, yields the following: vehicle=cell; content=fly; exploitation=fly-aimed tongue snaps.

In purported contrast, the correspondence between vehicle and content at the core of structural representation does not rest in a probabilistic relationship between vehicle and content, but instead depends on resemblance. Sometimes the resemblance between vehicle and content is with respect to a single kind of physical property. A wax figure of Abraham Lincoln represents Abraham Lincoln because properties of the wax figure such as its height, the length of its arms, and the width of its nose,

<sup>3</sup> When offering a *naturalistic* theory of representational content, one must avoid taking the exploitative act to be itself a mental act, such as thinking, because doing so incorporates into the analysis of mental representation just the kind of capacity that one wishes to explain (Opie and O’Brien 2004, pp. 4–5).

<sup>4</sup> Dretske later suggests that he is open to weakening the strength of the connection between the representational vehicle and its content to something less than 1 (Dretske 1994, p. 62). The important point is that indication is a relation between a representational vehicle and its content that observes some sort of nomic regularity. Thus, although a theory of representation like Fodor’s (1990) does not explicitly mention conditional probabilities, its reliance on a nomic regularity between a representational vehicle and its content suffices to make it a species of an indicator account of representation.

<sup>5</sup> Dretske (1988) distinguishes between cases in which natural selection is responsible for the recruitment of an indicator and cases in which the recruitment occurs as a result of learning. Although important, this distinction can be set aside for present purposes.

correspond to these properties of Abraham Lincoln.<sup>6</sup> This is a first-order resemblance, for it is a resemblance between properties of the same kind.

Structural representations may also exhibit second-order resemblance. In such a case, what matters is not a one-to-one correspondence between properties of the same kind, but a correspondence “in which the relations among a system of vehicles mirror the relations among their objects” (Opie and O’Brien 2004, p. 10). A map of Madison, for instance, bears a second-order resemblance to the city of Madison with respect to the relation of distance, insofar as the distance between points on the map corresponds to the distance between points in the city: greater distances in the former correspond to greater distances in the latter. In this example, the correspondence is between relations of the same kind, i.e. distance. But second-order resemblances can be of a more abstract kind, involving correspondences between relations of different sorts. In another familiar example, the curvature of a thermostat’s bi-metallic strip corresponds to ambient temperature. As ambient temperature increases, the curvature of the strip expands. The changing curvature of the strip resembles changes in temperature in virtue of a second-order resemblance—a correspondence between different kinds of relations.

Returning once more to the triadic analysis of representation, we see that the vehicle of structural representation is a set of property instances (if first order), or relations between property instances (if second order), of some object, e.g. a wax figure, or a map, or a bi-metallic strip. The content of these vehicles are the properties, or relations between properties, of objects in the world (the properties of Abraham Lincoln, or the relations between the features of a city, or the relations between temperature and curvature). The exploitation of structural representations involves whatever thoughts or behavioral dispositions that come about as a result of the resemblance that exists between the properties or relations between properties in the representational vehicles and the properties or relations between properties of some object in the world.

### 3 Representation and the content causation problem

Structural representationalists argue that indicator representations fall victim to the *content causation problem* (O’Brien 2016, p. 2). All representationalists agree that contents ought to be causally relevant in the production of behavior. If not—if content does not do anything—then claims like the following would be false: the frog snapped at the fly *because* the content of its visual state was *fly*; the bi-metallic strip turned on the furnace *because* it registered the temperature to be 68°; Greg went to the Majestic Theater *because* the content of his belief was *the Big Lebowski is*

<sup>6</sup> ‘Correspondence’ in this paragraph may appear ambiguous, meaning either something like a probabilistic relationship between a vehicle and its content or something more like an *isomorphism* (or *homomorphism*) between properties. As we will argue later, we deny this ambiguity, for the correspondences in which indicators participate are themselves isomorphisms (or homomorphisms) of a sort.

*playing at the Majestic Theater*. Given the *prima facie* truth of such claims, representational content must be causally relevant.

We shall now argue that in fact structural representations fare no better or worse than indicator representations with respect to causal relevance. This is to be expected if, as we will argue in Sect. 5, indicator representations *are* structural representations.

Why, according to structural representationalists, do indicator representations violate the causal constraint? Consider once more the cell in the frog's visual system. Activation of this cell causes the frog to flick its tongue at passing flies. Notice that this claim differs from the following: the cell's having the content *fly* causes the frog to flick its tongue at passing flies. The activation of the cell, not what the cell means, is causally productive of the frog's behavior. This is obvious when one thinks of causation from the perspective of an interventionist theory (e.g. Woodward 2003): hold fixed the cell's activation but vary the correspondence between this activity and the fly (perhaps by stimulating the cell directly in a fly's absence), and the cell will continue to cause the frog to snap its tongue. The problem the indicator representationalist faces is to explain how a content-fixing correspondence relation can be causally relevant in an explanation of behavior. It is the cell's firing that causes the frog to snap its tongue, not the content-endowing relationship between the cell's firing and the fly. As O'Brien remarks about Dretske's version of indication-based representation, "the relations at the core of his proposal are powerless to explain the required behavioural dispositions" (O'Brien 2016, p. 7).<sup>7</sup>

Structural representations, on the other hand, enter into causal explanations of behavior precisely because of their resemblance to the objects they represent. Consider the bi-metallic strip's causal role in triggering the furnace. What makes it turn the furnace on (or off)? The curvature of the strip resembles (in a second-order way) the temperature of the room. Because the pattern of variation in the metal's curvature mirrors the pattern of variation in temperature, "it is then simply a matter of rigging the innards of the thermostat so that its operation of the furnace is regulated by these internalized surrogates" (O'Brien 2016, p. 9). This appeal to resemblance and its attendant causal relevance is offered in contrast to the more meager resources of indicator representations which, because they fail to resemble their contents, have no story to tell about their source of causal significance.

We regard this defense of the superiority of structural representations to be unfounded. To see this, consider the interventionist-based objection to the causal relevance of indicator representations. If activation of the cell in the frog's visual system is held fixed while an intervention breaks the correspondence between this activation and a fly, then the frog will still flick its tongue. We do not deny this. However, the same can be said of the bi-metallic strip. If one holds fixed its

<sup>7</sup> Opie and O'Brien (2004) discuss the *causal constraint* on a theory of mental representation, which requires that content be causally efficacious. However, the reasons they offer for doubting the relevance of content seem to involve suspicions about circularity rather than the inability of a relation between a vehicle and its content to play a causal role in the production of behavior. The circularity involves an analysis of content that looks to the dispositions that contentful states cause as being themselves constitutive of content (Opie and O'Brien 2004, pp. 2–3).

curvature while intervening on the correspondence between curvature and temperature (say, by gluing the strip to a material that does not change its shape with changes in temperature), then it will continue to affect the behavior of the furnace. The curvature of the strip *screens off* variations in temperature from the behavior of the furnace. Hence, when O'Brien complains that the relations on which Dretske focuses are "powerless to explain the required behavioural dispositions" (O'Brien 2016, p. 7), he could as well be criticizing his own view. The curvature of the bi-metallic strip will regulate the furnace's behavior regardless of whatever its relation to temperature.

Gładziejewski and Miłkowski (2017), who, like O'Brien (2016), believe that structural but not indicator representations avoid the content causation problem, seem aware of this point. For this reason, they shift focus from an explanation of a representation's causal relevance to behavior to an explanation of a representation's causal relevance to *successful* behavior. Holding fixed the bi-metallic strip's curvature while varying the temperature in the room will cause the furnace to, say, remain on when the room is very warm, or remain off when the room is very cold. The thermostat succeeds in its goal of maintaining the desired temperature only when the bi-metallic strip's curvature faithfully resembles temperature: "[M]anipulating actions is not the same as manipulating success. Because of this, the effect that the structure of the vehicle has on action does not imply that the same sort of relationship exists between the vehicle's structure and success" (Gładziejewski and Miłkowski 2017, p. 346). In other words, one cannot hold fixed the curvature of the bi-metallic strip and intervene on the correspondence between the curvature of the strip and temperature without also affecting the success of the thermostat in regulating temperature. The strip's curvature does *not* screen off variations in temperature from the *success* of the furnace.

However, the exact same point applies to the indicator representation in the frog's visual system. If we wish to explain the frog's *successful* tongue snaps—those that result in the frog capturing a fly—we must appeal to the correspondence between the activity in the retinal cell and the fly. Holding fixed this activity while breaking the correspondence between it and the fly will also affect the frog's success in capturing prey.

A peculiarity in the structural representationalist's objections to indicator representations is the fact that Dretske—the paradigm of indicator representationalists—develops his view in part on the basis of a discussion of the bi-metallic strip. Is Dretske simply wrong that the relationship between the curvature of the strip and temperature is in fact one of indication? Or, perhaps he is wrong that indication, while present, is in fact causally explanatory? O'Brien seems inclined to the second view. He charges Dretske with missing a distinction between, on the one hand, the *structural correspondence* relation between the curvature of the bi-metallic strip and temperature and, on the other hand, the causally established indication relation, between the two: "it is the fact that the curvature of the bi-metallic strip systematically mirrors the temperature, and not the causal covariation per se, that explains its capacity to operate the furnace in an appropriate manner" (2015, p. 8). Were a correspondence between curvature and temperature maintained while the causal connection between the two were broken,

O'Brien notes, the thermostat would continue to regulate the room's temperature successfully.<sup>8</sup>

As a response to Dretske, we think that this reply is unfair. Above, we saw that Dretske defined indication as a relationship between a representational vehicle *R* and its content *C* such that  $\text{Pr}(C|R) = 1$ .<sup>9</sup> In the counterfactual case that O'Brien imagines, it remains true that the probability is one that the temperature is, e.g. *X*, given that the bi-metallic strip's curvature is, e.g. *Y*. Hence, it remains true that the correspondence he imagines satisfies the definition of indication, and so, from Dretske's point of view, the curvature of the strip would indeed meet a condition for representing temperature, and would do so because it indicates temperature. O'Brien tries to dismiss this line of response, acknowledging in a footnote that Dretske distinguishes indication from causation, i.e. regular correspondence from the causal relations that might ground such a correspondence. However, he says, Dretske's distinction between indication and causation should be ignored, for only by doing so can Dretske be seen as offering a causal theory of content. But, surely, if Dretske had wished to offer a causal theory of content rather than an indicator-based theory, he would have done so. Doubtless, his preference for indication over causation as the grounding relation for representation was motivated in part by just the sorts of concerns that O'Brien's distinction between correspondence and causation raises.

To conclude this section, indicator and structural representations are on a par with respect to the content causation problem. Insofar as holding fixed the activities of either sort of representation while breaking their correspondence to the world results in the same behavioral dispositions, they are causally inert. Insofar as a content-providing correspondence exists between them and the world, their activities will lead to successful behavior—capturing flies in the case of the frog; regulating temperature in the case of the thermostat. But, of course, that either sort of representation can explain successful behavior still does not establish that they do so in virtue of their contents—which is the crux of the content causation problem. Here too we see a parallel in the responses that one might offer to this challenge. The structural representationalist notes that a particular vehicle is “chosen” for its role in a system because of the similarity it bears to its content: the bi-metallic strip is “rigged” to the operation of the furnace because of its correspondence with temperature. But such a story is no less true of the frog's fly-detector: it is “rigged” to the operation of tongue snapping because of its correspondence with flies. We examine considerations having to do with the importance of the exploitation of representational vehicles further in the following section.

<sup>8</sup> Exactly how the correspondence could persist while curvature was no longer under control of temperature is unclear.

<sup>9</sup> Or, at any rate, something close to  $\text{Prob} = 1$ . Nothing that follows depends on how close to perfect correlation the indication relation approaches.

## 4 Representation and content determinacy

The disjunction problem, or the problem of content indeterminacy (Fodor 1984, 1987) challenges the representation-theorist to explain the possibility of misrepresentation. The problem arises because the content of representational states can always be re-described so that any purported case of representational error becomes an instance of a correct representation with a disjunctive content. For instance, an account of representation should be able to say why the neural state in the frog has the content *fly* as opposed to *fly or bee bee* even though the same neural state is tokened in the presence of both flies and bee bees. Here, the indicator theorist leans on exploitation to explain content determinacy. For instance, Dretske appeals to the way the fly detector is “used” or “harnessed” or “recruited” by the sensory system in which it is embedded to produce relevant tongue flicking behavior (Dretske 1988). It is this exploitation that fixes the representational content of the neural state—that makes it *fly* rather than *fly or bee bee*.

We will show that the most plausible solution to this problem for a structural representationalist requires a significant concession. More precisely, it requires the structural representationalist to lean as heavily on the exploitation feature of representation as does the indicator representationalist. In consequence, the features of structural resemblance that its advocates see as distinctively valuable, namely the first- or second-order resemblance between vehicles and their contents, turn out no longer to be playing a starring role in the account. Instead, structural representation-ists end up appealing to the same “use-criterion” on which indicator representationalists depend to avoid the disjunction problem.<sup>10</sup>

First, begin with any potential structural representation, whether it be a map of the city of Madison or a bi-metallic strip, the shape of which varies with temperature. According to the structural representationalist, these structures all represent their relevant objects, if they represent anything at all, in virtue of bearing a relation of structural resemblance to their represented objects. Consider, though, that any system or structure, A, that bears a relation of structural resemblance to another system or structure, B,<sup>11</sup> will also bear a relation of structural resemblance, in principle, to a large set of other systems or structures.

For instance, take a map of Madison and now suppose that there is another city, Waterbury, to which the map bears a similar relation of structural resemblance. Further, consider that as a map of Madison (or any city) becomes more coarse-grained, it will successfully, though accidentally, bear a relation of structural resemblance to an increasingly larger set of other cities and towns, e.g. a map of a small town that leaves out only but the most basic details may bear a relation of structural resemblance to any number of other small towns. In fact, such a coarse-grained map may bear a relation of structural resemblance to an amusement park, an office floor plan, and so on. Similarly, the changes in the curvature of a bi-metallic strip might mirror

<sup>10</sup> Isaac (2013) stands as a possible exception to structural representationalists who neglect the significance of use in addressing the problem of misrepresentation.

<sup>11</sup> Whether in a first- or second-order way.



not only changes in temperature, but also changes in pressure, and so bear a second-order relation of structural resemblance to both. Clearly, then, it is possible for A to bear a relation of structural resemblance to both B or C, and yet, intuitively, represent only B. The structural representationalist must have a principled way to account for this.

We claim that this now leaves the structural representationalist facing a dilemma. On the one horn, the structural representationalist could accept that a system or structure represents everything to which it bears a relation of structural resemblance, in which case the content of the representation is indeterminate. Though some structural representationalists have embraced this horn (notably Cummins (1996)), its concession to content indeterminacy presents a strong *pro tanto* reason to reject a structural representation account. The problem is that if content is indeterminate, it would seem to follow that content can play no explanatory role in cognition or behavior. For example, because the curvature of a bi-metallic strip bears a relation of structural resemblance to both temperature and pressure, the structural representationalist needs to offer a compelling reason to think that it is because it represents *temperature* that the bi-metallic strip's varying curvature explains the behavior of the furnace. However, the mere fact that temperature happens to be *one* of the things to which the bi-metallic strip bears a relation of structural resemblance is insufficient to do this.<sup>12</sup>

Pretty clearly, structural representationalists wish to avoid the above conclusion. O'Brien, for instance, discusses just this problem of indeterminacy with respect to the representational content of the bi-metallic strip: "The most we can say about the thermostat's bi-metallic strip is that its curvature represents that potentially large and motley collection of objects with which it systematically corresponds" (2016, p. 11). However, this would be a concern only if structural resemblance on its own was intended to fix a particular content to a representational vehicle. As we have seen, structural representationalists suggest that, in addition, structural resemblance must be exploitable by the system.<sup>13</sup> For the structural representationalist, this means that the behavioral dispositions of the system towards the vehicle's represented object are modified in virtue of the relation of structural resemblance between the vehicle of the representation and the representational object (Opie and O'Brien 2004, p. 5).

<sup>12</sup> We should note that the present comments on this first horn of the dilemma do not adequately address Cummins' full defense of structural representations. Though isomorphism is sufficient for representation on Cummins' view, he is aware of the need to limit content ascription. His strategy involves invoking the function of tokening a particular representation on an occasion of use. Though it's beyond the scope of this paper to address Cummins' view in any detail, we believe that the above strategy would result in Cummins' account impaling itself on the second horn of the dilemma, which we outline below.

<sup>13</sup> Philosophers of mind often distinguish two questions about mental representations (Ramsey 2007, 2016). The first is a question about representational status: why think some state or structure is genuinely representational? The second is a question about representational content: what makes it the case that the representation has the content that it does? It is unclear whether the acknowledgment that structural resemblance is insufficient for representation is intended to address the former question or the latter (or both). Regardless of intent, we take it that a theory of representation is not complete until both questions have been answered, and, further, that exploitation is the best candidate available to the structural representationalist for fixing representational content.

In other words, the relation of structural resemblance must dispose the subject to behave appropriately towards the representational content. Thus, we find O'Brien describing exploitation as an effective means by which to address indeterminacy<sup>14</sup>: "...interpretation plays an important content-limiting role. Specifically, a system's behavioural dispositions will anchor its representing vehicles to particular represented domains" (2016, p. 11).

However, we claim that this response to the disjunction problem threatens the very distinction that structural representationalists hope to draw between their view and the indicator representationalist's. This danger was already implicit in our critique in the previous section, in which we claimed that indicators and structural representations are on a par with respect to the content-causation problem. Current considerations make more explicit this aspect of our critique. Our claim is that the only way for the structural resemblance between a vehicle and its object to *in fact* dispose the system to behave appropriately requires that the system already be set up to use the vehicle as a representation of the object. Rather than the vehicle's structural features all on their own enabling the system to behave appropriately, it is the way a system is designed to use the vehicle that results in its appropriate behavior (a point to which we shall return in the next section). For example, the structural resemblance between a map of Madison and Madison disposes the user of the map to behave appropriately with respect to navigating Madison only if she antecedently knows that the map is meant to be a map of Madison. Similarly, changes in curvature of a bi-metallic strip that bear a relation of structural resemblance to changes in temperature dispose the thermostat to behave appropriately with respect to temperature only if the bi-metallic strip is embedded in a system designed to use the changes in the curvature to control temperature.

Recall once more O'Brien's observation that, given the structural resemblance between the curvature of the bi-metallic strip and temperature, "it is then simply a matter of rigging the innards of the thermostat so that its operation of the furnace is regulated by these internalized surrogates" (2016, p. 9). However, once the points above about the importance of exploitation are acknowledged, the relation of structural resemblance loses its centrality in an account of representation, leaving one to wonder why it should be preferred to an indicator-based account.<sup>15</sup> On either

---

<sup>14</sup> Though Gładziejewski and Miłkowski (2017), Shea (2014) and O'Brien (2016) incorporate exploitability to, in part, address concerns of "panrepresentationalism," they are not explicit about whether exploitation is also the means by which they intend to account for misrepresentation. However, we take the problem of panrepresentationalism and the disjunction problem to be two sides of the same coin.

<sup>15</sup> Ramsey (2016) argues that indicator and structural representation theories need not be seen as competing theories of representation, but, rather, as complementary answers to distinct questions about representation. Specifically, he suggests structural representations are more naturally construed as playing a genuinely representational role, while indicator theories can provide a better account of content determination. A detailed discussion of this possibility is outside the bounds of this section (which is concerned exclusively with the question of content determination), however, we direct the reader to Rupert (2018) for an argument that indicators can meet what Ramsey sometimes refers to as the job description challenge (2007). In addition, in our next section we will argue that indicators *are* structural representations. Though our paper isn't intended as a defense of indicators as such, if one is already convinced that structural representations can meet the job description challenge, it ought not be surprising that indicators can do so too.

approach to representation, we contend, as much of the heavy lifting with respect to explaining successful behavior rests on the design of the system in which the representing vehicles are embedded as it does the properties of the representing vehicles themselves.


On reflection, what it means for a representing vehicle to possess the capacity to dispose a cognitive subject to behave appropriately towards the object of the representation is unclear. Place a bi-metallic strip into a system not designed to take advantage of the correspondence between the curvature of the bi-metallic strip and temperature, and no such appropriate behavior with respect to temperature will result (a bi-metallic strip installed into a washing machine will not insure that a room's temperature remains constant). Of course, we do not deny that something that bears a relation of structural resemblance can be recruited to cause appropriate behavior in the system in which it is embedded. Our claim is that once the idea of exploitation becomes central in the account of structural representation, one must wonder whether structural resemblance, in the sense that its proponents typically characterize it, is any longer necessary to cause the relevant appropriate behavior. A bi-metallic strip whose shape covaries with temperature is a useful structure to recruit, but an indicator, or set of indicators, would work just as well. This is unsurprising, given that, as we will see in the next section, indicators are themselves a limiting form of structural representations.

To be clear, our claim is not that the solution to the disjunction problem should forsake the content-determining capacity of exploitation. Rather, our claim is that in embracing exploitation as a means to a solution, structural representations lose what proponents have lauded as their primary advantage over other approaches to representation. As structural representationalists characterize their view (Opie and O'Brien 2004, p. 17), the intrinsic properties of the vehicles are supposed to run the show, in purported contrast to indicator approaches, which must appeal to use conditions to ground content. However, if the previous considerations are correct, on either approach to representation, use is playing an integral content fixing role.

## 5 Indication is structural resemblance

We have argued that structural and indicator representations are comparable in at least two important respects. Both kinds of representation stand in a correspondence relation to their contents and, because of this, are recruited for some purpose—regulating a furnace, or capturing flies. The relation that a vehicle bears to its representational content explains its presence in a system—causes it to play the role in the system that it does—even though this relation itself does not cause the system to exhibit the behavior that it does. Secondly, because both structural and indicator vehicles correspond to myriad properties or relations, assignments of particular contents risk indeterminacy. To the extent that this difficulty might be resolved, appeals to the use to which a system puts a representation grow in significance. We submit that the parity of structural and indicator representations in the above respects is no coincidence, but instead owes to the fact that they do not, contrary to the claims of structural representationalists, differ in kind.

**Fig. 1** A spring scale making use of *spring1*, a spring that extends 1.0 cm in length for each 1.0 g weight attached to it



Perhaps the easiest way to see this point is to imagine a spring scale, of the sort pictured in Fig. 1, that makes use of a metal spring with a slightly unusual property. Attaching a weight of 1.0 g to the end of the spring will cause its length to extend by some fixed proportion, e.g. 1.0 cm. The peculiarity of the spring, which we shall call *spring1*, is this: it will not extend any further in length until another 1.0 g of weight is attached to it. That is, *spring1* will remain fixed in length for all weights between 1.0 and 2.0 g, but will stretch another 1.0 cm in length when the weight reaches the 2.0 g threshold. *Spring1*, like the thermostat's bi-metallic strip, stands in a structural resemblance relation to weight. However, unlike the bi-metallic strip, whose changes in curvature are isomorphic to changes in temperature (at least for as long as it remains a solid), *spring1* bears only a homomorphic relation to weight. This means, roughly, that for every change in the length of *spring1* there exists a change in weight, but not vice versa, because, as just noted, changes in weight of less than 1.0 g do not correspond to changes in the length of *spring1*. In effect, *spring1* changes in length for every gram added (or subtracted) to its end, but it does so discretely.

So, *spring1* provides a discrete, rather than continuous, measure of weight. This difference between *spring1* and the bi-metallic strip should not, we think, preclude it from bearing a structural resemblance to weight and, if used to modify the behavioral dispositions of a larger system, like a grocer's scale, from representing weight. As noted, changes in *spring1*'s length do, after all, maintain a *homomorphism* with changes in weight: for every change in the length of the spring, there is a change in weight, but some changes in weight (those less than 1.0 g) are not matched by changes in *spring1*'s length. Just as the discreteness of changes in a digital watch

do not prevent it from representing the relation between units of time, the discreteness of changes of length in spring1 should not prevent it from representing weight. The system that exploits spring1 will not be as precise as one that uses a continuous measurement of weight, for it cannot be set to detect weights between 1.0 and 2.0 g, or between 5.0 and 6.0 g, but, depending on the purpose for which it is used (e.g. weighing potatoes), it may be precise enough.

But if the homomorphism between spring1's length and the weight of the objects it measures suffices to establish a structural resemblance between the two, then the same should be true of spring2, which is like spring1 except that it is sensitive to changes in weight of 2.0 g increments rather than 1.0 g increments. Whereas spring1 changes in length by 1 cm for every change in weight of 1.0 g, spring2 changes in length by 1 cm for every change in weight of 2.0 g. The spring2-equipped scale is even less precise than the spring1-equipped scale, because it cannot be used to identify the weight of things that are, e.g. exactly 3.0 g or exactly 5.0 g (because these weights fall between 2.0 and 4.0 g, and between 4.0 and 6.0 g), but spring2 nevertheless continues to bear a homomorphism, and hence structural resemblance, to weight.

And, of course, if changes in the length of spring2 can be said to structurally resemble and, hence (*modulo* the appropriate exploitation), represent weight, why wouldn't the same be true of spring4, and spring8, and spring16, and so on. Let spring $N$  be a spring that changes 1 cm in length for some very large change in weight. Suppose that the crucial weight that marks the boundary point at which spring $N$  "jumps" in length is 100 kg. The scale that makes use of spring $N$  is now, in a sense, maximally insensitive. It will "detect" something hanging from its end when and only when the object's weight is 100 kg or greater. Changes in the length of spring $N$ , we contend, continue to bear a homomorphism to changes in weight, for the same reason that changes in the lengths of spring16, spring8, spring4, spring2 and spring1 did. To be sure, the structural resemblance between changes in the length of spring $N$  and changes in weight is very coarse-grained, but this no more disqualifies it from representing weight than a map's being coarse-grained prevents it from representing a city. A map that depicts merely the lakes that surround Madison still represents Madison despite not including symbols for streets, buildings, and landmarks.

On our view, spring $N$  is like a switch that indicates when an object's weight is equal or greater to 100 kg. This is because the probability that the object weighs less than 100 kg given that the length of spring $N$  is  $X$  cm is equal to 1; and the probability that the object's weight is equal or greater to 100 kg given that the length of spring $N$  is  $X$  cm + 1 cm is equal to 1. However, given spring $N$ 's continuity with the other increasingly sensitive springs, we find no reason to deny that its behavior is homomorphic to, and thus resembles, changes in weight. We conclude that indicators are simply limiting cases of structural representations.

We are not the first to claim that structural representations are of a kind with indicator representations. Morgan draws the same conclusion (although by different means):

mechanisms like oil lights on car dashboards, smoke alarms, and pregnancy tests might participate in very simple homomorphisms with the systems they measure, but they nevertheless *do* participate in homomorphisms with those systems; that’s how they’re capable of measuring them. An oil light is no less a measuring instrument than a fuel gauge, simply because it can occupy fewer states; it’s just a less *discriminating* measuring instrument (2014, pp. 232, his italics).

We agree with Morgan. Indicators, such as the one the frog uses to represent the presence of flies, may not be very discriminating—the cell’s firing rate, e.g. may not vary in proportion to the tastiness of the fly—but its fire/don’t fire states nevertheless stand in a homomorphic relation to states of the world and thus qualify as structural representations.

Aware of Morgan’s point, Gładziejewski and Miłkowski (2017) seek to defuse it in an effort to re-establish the distinctness between structural and indicator representations. To show this, they imagine a thermostat with a detector strip that enters different states, each of which corresponds to a distinct temperature, while, at the same time, *not* bearing relations to each other that mirror the relations between temperatures. For example, in a normal bi-metallic strip, the curvature of the strip when indicating 33° is closer to the shape of the strip when indicating 34° than it is when indicating 17°. However, in their imagined example, the shape of the bi-metallic strip when indicating 33° is now more similar to the shape of the strip when indicating 17° than it is when indicating 34°. Gładziejewski and Miłkowski maintain that despite the fact that the relations between the indicator states no longer structurally resemble ambient temperature, it is possible for the functioning of the thermostat to remain the same, as long as it is appropriately rigged to switch the furnace on or off in response to the relevant indicator state. The example shows, they say, that “it is not *necessary* or *essential* for the relational structure of possible indicator states to replicate the relations between different variants of ambient temperature in any particular way (2017, p. 348).” In short, structural representations succeed in their jobs *because* of their resemblance to their objects; the success of indicators doesn’t so depend.<sup>16</sup>

However, we believe that Gładziejewski and Miłkowski once more fail to appreciate the importance of the role of exploitation in determining relations of structural resemblance. One might suppose that if a complex of indicators correctly regulates the behavior of a furnace, then the indicators must stand in the same relation to each

<sup>16</sup> Gładziejewski and Miłkowski’s full defense of the distinction between structural and indicator representations involves establishing two crucial differences between them. One difference is captured by the argument above having to do with the role that resemblance plays in structural, but not indicator, representations. The other purported difference is that structural representations involve an endogenous source of cognitive and behavioral control, whereas indicators are purely reactive (2017). The argument we develop in the rest of the paper focuses on disputing the plausibility of the former difference. However, see Rupert (2018) for a response to the tenability of the latter distinction. Rupert argues that indicator-based theories (in particular, Dretske’s) give a more robust role to indicator representations than Gładziejewski and Miłkowski recognize, thus elevating them from being merely “reactive”.

**Table 1** A two-dimensional array of cells that represents a square-shaped figure

4	Off	Off	Off	Off
3	Off	On	On	Off
2	Off	On	On	Off
1	Off	Off	Off	Off
	1	2	3	4

**Table 2** The same square-shaped figure represented in a single row of cells

State	Off	Off	Off	Off	Off	On	On	Off	Off	On	On	Off	Off	Off	Off	Off
x	1	1	1	1	2	2	2	2	3	3	3	3	4	4	4	4
y	1	2	3	4	1	2	3	4	1	2	3	4	1	2	3	4

other than degrees of temperature do, even if the nature of this resemblance becomes conspicuous only when in possession of the proper “key”.

We find support for this idea in a simplified version of Kosslyn’s concept of a “quasi-pictorial” representation (Kosslyn 1983). Consider an area of visual cortex that consists in a sheet of cells (indicators) that can be either “on” or “off” (Table 1). Using the format (x, y) to identify the cells in the cortical sheet, we can see that the cells structurally resemble a figure of a square because (2, 3) is to the left of (3, 3) and above (2, 2), which is to the left of (3, 2). Gładziejewski and Miłkowski, presumably, would assert that these cortical cells represent a square because of the structural resemblance that they bear to it.

But now suppose that the region of cortex that represents a square figure is arranged in a single row (Table 2). Although these cells are arranged in a single row rather than four rows, they needn’t be. They may not be contiguous at all, instead spread about in what appears to be a haphazard organization. On Kosslyn’s view, this *apparent* lack of structural resemblance disguises a genuine resemblance that comes into focus in virtue of the processes that operate on the cells. When “interpreted” *so*, the 1-row array of cells contains all of the information about the shape of the square object that the 4-row array of cells contains.<sup>17</sup> Likewise, the imagined haphazard arrangement of cells could also, when viewed through the appropriately tuned lens, resemble a square. Indeed, we should not be fooled by the obviousness of the resemblance between the “on-cells” in the four-column array and a square. The cells’ “squareish” appearance does not entail that they represent a square to the system that exploits them. If the system used the “aboveness” relation as indicating, e.g. diagonality, then the 4-row array, despite looking squarish, may actually be depicting a diamond shape.

<sup>17</sup> Some structural representationalists acknowledge that something like the above example would count as a second-order structural resemblance (see: Isaac 2013, p. 700). We of course welcome this conclusion, and view it as further support for our claim that indicator and structural representations do not differ in kind.

In sum, not only do indicators structurally resemble their objects, as our argument involving spring $N$ , showed, but arrays of indicators, unsurprisingly, do as well. Moreover, they appear to do so in virtue of a structural resemblance that they bear to their object even if, as the above discussion of quasi-pictorial representation illustrates, the visibility of the resemblance requires knowledge of how the relations between representing states are exploited.

Though our response to Gładziejewski and Miłkowski uses an example in which exploitation reveals a structural resemblance between a set of indicators and their contents, our argument should not be taken as support for the claim that any collection of indicators can or should always be treated as being systematically related to each other, or that doing so is necessary for an indicator to be considered a structural representation. Our main claim continues to be that a single indicator, in virtue of resembling its object, structurally represents its object, albeit in a very coarse-grained way. Consider the following example from Shea (2018, p. 119) involving alarm calls made by vervet monkeys in the presence of predators. Vervets make three distinct kinds of alarm calls to warn conspecifics of the presence of eagles, leopards and snakes. Shea explains that while vervets exploit the correlation between each alarm call and what that call corresponds to, the relation between the kinds of calls plays no role in how the calls represent the predators that they do. Thus, even if we suppose some systematic relation to exist between the kinds of calls, and surely there will be arbitrarily many such relations, no appeal to this relation is necessary to explain how the vervets use the alarm calls to warn each other of predators. For instance, Shea imagines that the alarm calls can be systematically related to each other in order to capture the *higher than* relation between predators (eagles are usually higher up than leopards, which are usually higher up than snakes). But, because, in fact, vervets make no use of a correspondence between the relations between kinds of alarm calls and the heights of predators, Shea concludes that the array of alarm calls is therefore not a structural representation.<sup>18</sup>

We agree with Shea that, *taken as an array*, the alarm calls are not a structural representation. However, we hold that, taken individually, the alarm calls *are* structural representations. Indeed, the example of the alarm calls reinforces the point we have been emphasizing throughout this section. The alarm calls, as prototypical examples of indicators, can, were they exploited in the appropriate way, be taken to stand in some sort of systematic relation to each other, just like the states of Gładziejewski and Miłkowski's aberrant bi-metallic strip. But, unlike in the case of the bi-metallic strip, no benefit accrues from this particular exploitation for the vervets, as the *higher than* relation does not track anything of use to them. This marks a contrast to our example involving the array of indicators. In this case, an appropriate exploitation of the relations that the cells bear to each other does confer a benefit., viz. the accurate representation of a square.

---

<sup>18</sup> Shea, in using this example, isn't interested in the difference between structural and indicator representations. His primary aim is to establish the need for an exploitation condition for structural representations. A relation of structural resemblance between a set of vehicles and a set of objects is only a structural representation if it is exploited in an appropriate way by the system.



However, one must keep in mind that, as we illustrated with our example of the spring scale, indicators need not bear systematic relations to other indicators to bear a homomorphic relation to their contents. Spring1 bears a homomorphic relationship to weight because the different states of spring1 structurally resemble (or correspond to) *the presence of* weights of a particular range hanging from the bottom of the spring. The states of spring $N$ , too, correspond to the presence of certain weights. It just so happens that spring $N$  has only two possible states (corresponding to weights less than 100 kg or greater than or equal to 100 kg). When exploited in virtue of this resemblance, the spring represents weight. The calls of the vervet monkey are no different. Each call structurally resembles the presence of a different predator, and, when exploited by the vervets in virtue of that resemblance, represents the presence of that predator.

On reflection, one must wonder how, on the structural representationalist's view, collections of indicators could possibly represent if each individual indicator did *not* resemble its object. All natural forms of representation must depend on resemblance of some sort, even if the resemblance is no more specific than the on–off behavior of indicators.<sup>19</sup> Without a correspondence between a representing state and its content of at least this minimal kind, a system's ability to use the state to tell it something about its content would be completely occult—akin to using tea leaves to tell the future.

## 6 Conclusion

Structural representationalists believe that resemblance is a crucial feature of representation. In virtue of resemblance, structural representations avoid the content causation problem. Similarly, specific content assignments to structural representations are possible because of their exploitability. We have argued that if these claims are true of structural representations, they are true as well for indicator-based representations. The reason for this is simple: indicator representations structurally resemble their objects.

**Acknowledgements** Thanks to two very thorough referees for comments that greatly improved this paper. Thanks also go to Gerard O'Brien for useful discussion, and to Rob Rupert for comments on an earlier draft.

## References

- Bechtel, W. (1998). Representations and cognitive explanations: Assessing the dynamicist challenge in cognitive science. *Cognitive Science*, 22(3), 295–317.
- Cummins, R. (1996). *Representations, targets, and attitudes*. Cambridge: MIT Press.

<sup>19</sup> We specify “natural” forms of representation, for conventional representations, e.g. “Let this egg represent Earth,” *create* correspondences rather than being recruited for correspondences that occur in nature. For more on conventional representation, see Dretske (1988: 52 ff.).

- Dennett, D. (1982). Styles of mental representation. *Proceedings of the Aristotelian Society*, 83(213–226), 213–226.
- Dretske, F. (1988). *Explaining Behavior*. Cambridge: MIT Press.
- Dretske, F. (1994). The explanatory role of information. *Philosophical Transactions: Physical Sciences and Engineering*, 349, 59–70.
- Fodor, J. (1984). Semantics, wisconsin style. *Synthese*, 59(3), 231–250.
- Fodor, J. (1987). *Psychosemantics*. Cambridge: MIT Press.
- Fodor, J. (1990). *A theory of content and other essays*. Cambridge, MA: MIT Press.
- Gładziejewski, P., & Miłkowski, M. (2017). Structural representations: Causally relevant and different from detectors. *Biology and Philosophy*, 32(3), 337–355.
- Hardwick, C. (Ed.). (1977). *Semiotic and signification: The correspondence between Charles S. Peirce and Victoria Lady Welby*. Bloomington: Indiana University Press.
- Isaac, A. M. C. (2013). Objective similarity and mental representation. *Australasian Journal of Philosophy*, 91(4), 683–704.
- Kosslyn, S. (1983). *Ghosts in the mind's machine*. New York, NY: W.W. Norton.
- Morgan, A. (2014). Representations gone mental. *Synthese*, 191(2), 213–244.
- O'Brien, G. (2016). How does mind matter? Solving the content causation problem. In T. Metzinger (Ed.), *Open MIND philosophy and the mind sciences in the 21st century* (Vol. 2, pp. 1137–1150). Cambridge: MIT Press.
- Opie, J., & O'Brien, G. (2004). Notes toward a structuralist theory of mental representation. In H. Clapin, P. Staines, & P. Slezak (Eds.), *Representation in mind: New approaches to mental representation*. Amsterdam: Elsevier.
- Ramsey, W. (2007). *Representation reconsidered*. Cambridge: Cambridge University Press.
- Ramsey, W. (2016). Untangling two questions about mental representation. *New Ideas in Psychology*, 40(A), 3–12.
- Rupert, R. (2018). Representation and mental representation. *Philosophical Explorations*, 21(2), 204–225.
- Shea, N. (2014). Exploitable isomorphism and structural representation. *Proceedings of the Aristotelian Society*, 114(2), 123–144.
- Shea, N. (2018). *Representation in cognitive science*. Oxford: Oxford University Press.
- Von Eckardt, B. (1993). *What is cognitive science?*. Cambridge: MIT Press.
- Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford: Oxford University Press.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.