Synthese (2021) 198:6615–6629
https://doi.org/10.1007/s11229-019-02480-7

Transparency, representationalism, and visual noise

Joshua Gert[1]

Received: 11 April 2019 / Accepted: 16 November 2019 / Published online: 28 November 2019

Abstract
Those who endorse the twin theses of transparency and representationalism with regard to visual experience hold that the qualities we are aware of in such experience are, all of them, apparently possessed by external objects. They hold, therefore, that we are not introspectively aware of any qualities of visual experience itself. In this paper I argue that attention to visual noise—also known as 'eigenlicht' or 'eigengrau'—puts pressure on both of these theses, though in different ways. Phenomenally, visual noise does not even *seem* to belong to any external objects, which is a challenge to transparency. Moreover, visual noise is not the normal visual response to any distinctive external property, such as external graininess. Nor is it treated by our visual system as the perception of any such property. Given extant views of visual representation, it is therefore implausible to claim that it is the transparent representation of any such property.

Keywords Transparency · Representationalism · Vision · Perception

1 Introduction

Within contemporary philosophy of perception, the thesis of the transparency of visual experience is close to the received view. Here is one recent and very clear presentation of the thesis, due to Michael Tye:

> in a case of normal perception, if we introspect:
>
> 1. We are not aware of features of our visual experience.
> 2. We are not aware of the visual experience itself.

✉ Joshua Gert
jngert@wm.edu

[1] Department of Philosophy, The College of William and Mary, P.O. Box 8795, Williamsburg, VA 23187-8795, USA

3. We cannot attend to features of the visual experience.
4. The only features of which we are aware and to which we can attend are external features (colors and shapes of surfaces, for example).[1]

Transparency has often been used in order to defend strong representationalist accounts of perceptual experience: ones on which phenomenal character is either identical to, or supervenes on, representational content.[2] But it is not the point of this paper to examine the link between the two theses. Rather, the point is to call both of them into question by noting our capacity to be aware of a certain feature of our visual experience: visual noise.

Challenges to transparency and representationalism based on particular visual phenomena have, of course, already been offered. In particular, the phenomenon of blur has often been appealed to for such a purpose.[3] But there are a number of responses that have been made to blur-based challenges. Some say that in blurry vision we are aware of fuzziness or the absence of a boundary, or, more weakly, that we *fail*, in a salient way, to be aware of a boundary.[4] Others point out that extreme blur obviously impacts representational content—what we are aware of—so that lesser degrees of blur can plausibly be taken to introduce at least some difference in such content as well.[5] Perhaps those who use blur to challenge transparency or representationalism can undermine these responses; I take no stand on this here. But the nature of visual noise, I will argue, makes it much harder to mount parallel defenses. There are reasons—which do not have parallels in the case of blur—that it is not plausible that the experience of visual noise is the transparent representation, or misrepresentation, of anything.

Here is the plan of the paper. In Sect. 2 I briefly explain what visual noise is. Sect. 3 then argues that, while the phenomenology of visual noise may not provide a strict counterexample to transparency, it does a good deal to undermine the positive phenomenological case in its favor. In particular, it does this more effectively than does blur. In fact, it does a sufficiently good job that advocates of transparency will have to look beyond a mere phenomenological appeal for support. One such source of support might be an account of visual representation on which visual noise can plausibly be regarded as the representation of something external. With that in mind, Sect. 4 examines the two primary sorts of account of visual representation—tracking views and phenomenal intentionality views—arguing that neither of them vindicates the claim that visual noise is representational. Rather, both views tend

[1] Tye (2014, p. 40). See also Harman (1990, p. 39).

[2] See, for example, Tye and Harman. Transparency has also been used to defend naïve realist views of perception. See Allen (2016), Kennedy (2009), and Martin (2004). But a discussion of naïve realism is beyond the scope of this paper.

[3] See Pace (2007) and Smith (2008). Kind (2003) and Siewert (2004) also call Tye's strong brand of transparency into question, but not on the basis of particular counterexamples such as blur.

[4] See Dretske (2003), Tye (2003) and Schroer (2002). See also Allen (2013) for discussion of these and other interpretations of blur.

[5] Bourget (2015). This particular argument is a defense only of representationalism, not of the transparency thesis. The consistency of representationalism with the denial of transparency is defended in Siewert (2004) and Pace (2007).

in the opposite direction, though for different reasons. The conclusion is that visual noise provides a powerful case against strong representationalism, and undermines the case for taking transparency as a datum to be appealed to in arguments in philosophy of perception.

2 What is visual noise?

The phenomenon I am concerned with, and which I am calling 'visual noise', has sometimes been called 'eigenlicht' or 'eigengrau'. It is easiest to notice in a room with just enough light to allow one to make out the objects in one's vicinity. One common circumstance of this sort is the scene that presents itself at night in one's bedroom when there is only minimal ambient light coming in through the windows. Under such circumstances one's whole visual field has quite a salient graininess. If you fix your attention on a particular object—say, your hand in front your face—you can appreciate the degree of granularity. Once it is noticed in this way, it is very easy to notice the same grainy phenomenon when one's eyes are closed, whether in darkness or in a well-illuminated space.

In the cases described above, the phenomenon of visual noise pervades one's whole visual field. With one's eyes open and under greater illumination the phenomenon is still present, though more elusive. This elusiveness is unsurprising, since the visual signal is strong—so that there is a higher signal to noise ratio—and our attention is naturally fixed on external objects, not on the comparatively minimal noise in the signal. Still, once one is attuned to it, visual noise can typically be noticed in virtually the whole of one's visual field, except in regions that have quite high informational density already.[6] One can, for example, note it relatively easily in any shaded portion of a uniform expanse, such as a wall or table-top, especially if one avoids having one's visual attention captured by some discrete feature, like a scratch or breadcrumb, and focuses instead on a homogeneous region.

Eric Schwitzgebel has unearthed a number of descriptions of closed-eye eigenlicht from the nineteenth century, when phenomenological investigations were more common.[7] In fact, Schwitzgebel's overall goal seems to be to show that descriptions of what we see when our eyes are closed vary quite a bit, so that we ought perhaps to be somewhat skeptical of our own introspective reports. Still, even in Schwitzgebel's quotations, it is remarkable how often there is reference to dust. Here is a description from Gustav Fechner:

> By paying strict attention, one discovers in the blackness that is seen when the eyes are closed a kind of fine dust composed of light.[8]

[6] Although I do not know of any studies regarding the open-eyed, daylight version of eigenlicht, a surprisingly high proportion of researchers on the Color and Vision Network (CVNet) spontaneously corroborated my own experience.

[7] Schwitzgebel (2011).

[8] Fechner (1860/1966), p. 138.

Alfred Volkmann, also quoted by Schwitzgebel, describes visual experience of a dark field as absolute blackness but with a "light dust".[9] And a subject called 'Mr. H. D.' is described by George Ladd as saying that his experience is one of 'a dancing mass of vari-colored dust'.[10] Finally, Helmholtz mentions 'the "luminous dust" of the dark field'.[11] I myself find the description of my closed-eye experience in terms of dust relatively apt. I suspect that the variation that Schwitzgebel finds—and he does find quite a bit—has its source in a different phenomenon: what different people see *in* the dancing mass of dust: expanding circles, moving lines, and so on.[12] That sort of variation is more like the variation in what different people see in Rorschach inkblots, and it is, as a result, relatively unsurprising.

Of course, when one observes external objects, it is natural to ignore the noise and graininess of our experience. That is, it is natural to "see through" it. As a result, talk of the transparency of experience remains in a certain sense appropriate. But if the arguments of this paper are correct, this sort of transparency is what Amy Kind calls 'weak transparency', since it *is* possible to pay attention to the graininess if one tries, even if it is a mistake to say that one *sees* that graininess in the same sense as one *sees* an apple, or the redness of an apple.[13] As Kind persuasively argues, weak transparency is not the sort of transparency that is of interest to representationalists, since it is perfectly consistent with non-representational qualia. Indeed, it would be better not to call weak transparency 'transparency' at all, at least in the context of debates about perceptual transparency. It is not the phenomenon that Tye describes.

3 Transparency: noise versus blur

In this section I present some of the ways in which visual noise underwrites a stronger case against transparency than does blur. The point is not to show that visual noise provides a strict counterexample to transparency—one that anyone could be brought to acknowledge. That is too much to hope for. Rather, the point is to argue that attention to visual noise does a much better job than blur at weakening the phenomenological case in favor of transparency. In fact, it weakens it sufficiently that it cannot be regarded as a piece of phenomenal data to be used in arguments for representationalism.

[9] Volkmann (1846), p. 311.
[10] Ladd (1903), p. 79.
[11] Helmholtz (1856/1909/1962, p. 7).
[12] Interestingly, Helmholtz's phrase—'the "luminous dust" of the dark field'—occurs in a quotation offered by Freud (1913, p. 24). In that quotation the dust plays precisely this role of the substrate for various interpretations. Freud attributes the quotation to Wundt (1880), but seems to be in error.
[13] Kind (2003).

Here is one possible description of what is going on when one has experiences of open-eyed eigenlicht: one sees external objects by having experiences that themselves have a grain—noticeable primarily because of noise at that level of grain—that is not perceived to belong to those objects, or indeed to *any* objects. One is aware of this grain, and of the associated noise, as features of one's visual experience, and not as features of anything one sees or seems to see. If this is the right description of what is going on, then transparency of the strong form endorsed by Tye and others is false, as is the strong form of representationalism for which some people argue on its basis.

How might an advocate of transparency respond to the description of noisy visual experiences just offered? In the face of parallel descriptions of the phenomenology of blurry vision, one defensive strategy has been to offer rival descriptions that are consistent with transparency. For example, it has been claimed that when one sees a coffee cup blurrily, it looks as if the cup is made from some fuzzy material. That is, it looks as if an external property—fuzziness—is instantiated in the object of visual perception. Admittedly, based on prior experience, one *knows* that the coffee cup is not fuzzy. So when it is visually presented as fuzzy, one can immediately *infer* that something is wrong with one's vision. But as far as one's experience goes, the property one experiences is fuzziness: a property of an external object. Might not someone who wants to maintain transparency in the face of visual noise make use of the same sort of strategy? That is, couldn't one claim that a visual experience of a blank wall—when we are aware of the visual noise that characterizes it—is one in which we seem to see a wall that is covered in a layer of extremely fine constantly roiling dust, or something like that?

Here is another way of describing the general defensive strategy above. One might suggest that it is always possible to arrange the visual environment in such a way that someone placed in that environment who was accurately perceiving things would naturally—but wrongly—believe that she was having an experience with the purportedly non-representational feature cited by the denier of transparency. For example, if every blurry experience is indistinguishable from an associated veridical experience, it is plausible that blurry experiences transparently—though wrongly—represent certain associated external-world objects and properties. Call this the "doppelganger defense" of transparency from blur-based counterexamples.

Can the advocate of transparency mount a doppelganger defense in the case of visual noise? Constructing the relevant sort of case would require arranging things so that someone would have an experience *without* visual noise, but that *seemed* to be characterized by visual noise. It is far from clear that this is so much as a possibility. As will be discussed in Sect. 4.1, the nature of the visual system suggests that there is always visual noise. It is true that in very brightly-illuminated regions with high informational density visual noise is harder to notice. But, first, that does not mean that it is not there. And, second, even if it disappears under such specific conditions, it is simply not possible to construct noise-free doppelgangers that mimic the noise we experience in dimly-illuminated or less information-dense scenes. Both of these points make doppelganger defense harder to mount in the case of visual noise than in the case of blur.

I do not want to deny that there is some kind of phenomenological aptness to describing visual noise in terms of dust; indeed, I agree with this claim. This provides the advocate of transparency with a *prima facie* case favoring the same externalization strategy that has been used in response to the challenge from blur. After all, being covered in roiling dust is a property of external objects, just as fuzziness is. But if there is to be an appeal to phenomenology, it needs to be even-handed. So it is important that the graininess does not seem to be a property of the blank wall one sees, even if one is aware of the graininess when one looks at the wall.[14] Moreover, the extremely fine grain of the noise makes it phenomenologically unlike any actual texture of any actual object on which we can focus. Further, attempts to focus on it more or less closely have absolutely no effect on the grain or resolution, which means that one of our visual system's methods for determining distance gets no grip. The "dust", consequently, does not seems to be in any particular place. The most one can say is that it seems somehow appropriate to describe it as in front of whatever object one is focusing attention on—whether that be the sky, or the wall of one's office, or, with one's eyes closed, the inside of one's eyelids. But it does not seem to move nearer, or farther, as one focuses alternately on these differently-distanced objects. Nor, as one engages in this variation of focus, does it seem as if one is seeing different *regions* of it. Thinking of it as represented as being at a particular place therefore runs into problems. But anything represented in visual experience must, plausibly, be represented as being located somewhere in the scene. This is true even of a thick fog in which one is immersed; such a fog is represented as being located in all of the space in front of one, up to the point at which one's vision can no longer penetrate. And in the case of fog, one can focus on parts or regions that are closer or farther away, and the phenomenology bears this out; it seems as if one is looking at the near fog or the far fog. None of this is true of visual noise.

As far as phenomenology goes, then, visual noise speaks in a number of ways against transparency. Moreover, the aptness of the language of dust and "being in front of" does not really provide much of a *prima facie* case for transparency, since we can explain our use of such phrases in another way; our *shared public language* admits of no way to describe the phenomenology of experiences *except* in terms of the publicly available objects and properties by reference to which we learn that language. So if we are going to describe the experience at all, we are constrained to use whatever seems *most* apt, even if it isn't quite right. And it is farther from "quite right" in the case of visual noise than it is in the case of blur. The fuzziness that blurry vision might be claimed to represent is a property associated with particular objects we actually do see. In contrast, the dust or fog to which we might refer in trying to capture the phenomenon of visual noise characterizes the whole visual field. If one tries to explain it in terms of, say, an external layer or volume that intervenes between one's eyes and all the objects one sees, the illusion therefore will involve an entirely fictitious object—a cloud or layer of dust or fog. And there is less

[14] In favor of this phenomenological point, consider one's experience as one opens one's eyes in a very dimly lit room, facing a blank wall. The noise that characterizes one's closed-eye experience continues to characterize one's opened-eye experience, but it does not seem to *move* to the surface of the wall.

plausibility in the claim that we *always* have a *pervasive hallucination* of fog or dust than in the claim that we *sometimes* have a *localized illusion* of fuzziness.[15]

Here is another difference between blur and visual noise: one that also makes the latter the basis of a stronger phenomenological objection to transparency than the former. When one focuses one's attention on the blurriness manifested in one's experience, it is at least plausible to suggest that one must focus attention on the external object: the blurrily perceived cup, for example. This is precisely what an advocate of transparency would predict. But to the degree that one tries to focus attention on an external object in an attempt to pay more attention to the phenomenon of visual noise, one will be paying *less* attention to the visual noise. And paying attention to the noise, in parallel fashion, *detracts* from paying attention to the objects one sees *through* the noise.[16] These claims are even more phenomenally compelling when one reflects on visual experience in very low-light conditions, as during a nighttime visit to the bathroom. The visual noise one experiences on such an occasion, looking at a barely-visible sink, does not appear at all to be a characteristic of the surface of the sink. Attempts to focus on the sink make it harder to pay attention to the graininess and noise. And attempts to focus on the noise can cause one to lose visual awareness of the sink altogether.

The advocate of transparency might try to avoid the above point by denying that visual noise is the representation of a property of surfaces. Rather, it might be the representation of something in the space in front of those surfaces. In that case it would be no surprise that attending to visual noise when looking at a sink in very low-light conditions would make it harder to attend to the sink and its properties, and vice versa. But this suggestion is very problematic. If visual noise is the representation of a property of the space in front of visual objects, what property is it? It isn't like ambient illumination, at least if we take the descriptions in terms of dust at all seriously. Is it like dust roiling in a transparent medium? But then we would expect farther objects—including the colors of those objects—to be more obscured by it than nearer ones. But in a moderately well-illuminated room it is typically *visually apparent* that a farther wall is the same color as a nearer wall, and that the intervening space is completely transparent. That is, our awareness of the sameness of color is not the result of an inference, but is the phenomenally immediate deliverance of relatively well-understood mechanisms of color and lightness constancy. Could one perhaps maintain that farther objects are represented as simultaneously more brightly illuminated *and* more obscured by intervening dust, the two effects

[15] One might suggest that we do not in fact always experience visual noise, but only that it is always there to be experienced, if we direct our attention to it. This may be true, but that only increases the relevant difference with blur. A hallucination we can summon at will is much stranger than an illusion caused by whatever particular circumstances cause blurred vision. And the dependence on attention makes good sense if what we are attending to is a feature of the experience, rather than of the represented world, since there would normally be no reason to attend to features of experience that do not represent anything.

[16] This observation casts doubt even on Charles Siewert's more careful formulation of transparency: 'You cannot attend to how it appears to you, by turning your attention away from something that appears to you, and towards your experience' (2004, p. 35).

balanced to the precise degree required to result in actual brightness remaining relatively constant? But that is not the phenomenology at all, and in any case ignores the previous point about color constancy.

I've just argued that visual noise is not the representation of a property of the surfaces of the objects one sees, or of the volume of space in which one sees them. What then is it a representation of? My answer is: it is not a representation of anything. Rather, it is a non-representational feature of visual experience to which we can attend. Of course the possibility remains for the advocate of transparency to say that *some* visual noise is the representation of a property of the space in front of a surface, while *other* visual noise is the representation of a property of the surface itself. But this way out of the dilemma is obviously *ad hoc*. Moreover, it remains a task for the advocate of transparency to argue that this hybrid strategy combines the virtues of each interpretation, without its liabilities, rather than the other way around. And it is far from clear that this can be done. For example, if one attempts to deal with the problems of the volume-representing view by claiming that the noise that characterizes one's experience of a distant green wall is a representation of a property of the surface of that wall, it should be the case that focusing attention on the noise is a matter of focusing attention on the wall. But focusing on the noise actually makes it harder to note genuine features of the wall's texture.

Despite the length of this section, I do not mean to be placing too much weight on phenomenological considerations. The descriptions of experience that one finds apt are liable to be influenced by one's theoretical views, and some readers may find some of my descriptions less than compelling. That is fine; I myself am unpersuaded of the aptness of contrary descriptions in terms of transparent representation, though I don't doubt they are being offered sincerely. What that means is that there is a dispute about the phenomenology. And that is all I really need. The existence of such a dispute is enough to undermine the phenomenological case *for* transparency, especially if transparency is then meant to serve as a bedrock piece of phenomenal data on which arguments for representationalism are to be built.

In light of the foregoing, those who are wedded to transparency might try to bypass phenomenological considerations altogether; they might claim that visual noise *in fact* represents something like roiling dust, independently of whether or not it *seems* to do this. But this strategy runs the risk of putting the represented cart before the transparent horse; transparency is typically supposed to form the *basis* for representational views. Still, might one argue for representationalism on *other* grounds, and then use representationalism to support transparency? Yes, one might; non-phenomenological considerations can sometimes serve to clarify phenomenology.[17] Let me turn, therefore, to the impact of visual noise on representationalist views. As it will turn out, consideration of the ways in which visual experience might get its representational content does nothing to help the strong representationalist, or the advocate of transparency, in their efforts to deal with visual noise. In fact, the situation is quite the reverse; once we start giving more thought to the ways

[17] See Allen (2011).

in which visual experience might get representational content, the idea that visual noise corresponds to some aspect of that content loses even more plausibility.

4 Visual noise and perceptual content

In this section I address the impact of visual noise on strong representationalist views directly, independently of their dependence on transparency. Broadly speaking, there are two sorts of views to consider: those that take representation to ground phenomenology, and those that take phenomenology to ground representation. These two options do not, of course, exhaust the logical space. But I am unaware of any representationalists who have advocated any third option, so I consider only these two.

4.1 Tracking views

According to *tracking* representationalism about perceptual experience, such experiences are explained in terms of representation, and representation requires some kind of naturalistic connection—causal or informational—between a kind of mental item and a kind of worldly item.[18] Tracking views also require that the explanation of the connection be of the right sort; it must be plausible that the mental item *functions to indicate* the presence of its correlated worldly item. This talk of function can be cashed out in many ways. Often it is explained in terms of evolutionary history.[19] But it can also be explained in terms of a correlation established within the span of a single life by a mechanism that has the (evolutionary) function of establishing such correlations—that is, by learning.[20] Or it can be established by appeal to optimal conditions or a *ceteris paribus* clause.[21]

If we accept a tracking view of representational content, visual noise cannot plausibly be regarded as representational. Such noise does not function to track anything. Rather, it is the unavoidable result of limitations in our visual equipment—primarily in the retina, but in the brain as well.[22] One explanation for such limitations is that (a) neurons function by firing, (b) retinal neurons respond to incoming photons in an essentially stochastic way, and (c) firing is a digital sort of phenomenon. That is, at a given moment a neuron is either firing, or it isn't; it doesn't matter how "strong" each firing is, but what the rate of firing is. And even the rate of firing is not a smoothly-changing value, as a voltage might be, since with each firing, the relevant average rate of firing will also undergo an abrupt change. If the discontinuous nature of neurological processes is what explains visual noise, then—on tracking

[18] For a very nice overview of tracking representationalism more generally, see Bourget and Mendelovici (2014).
[19] Dretske (1988) and Millikan (1989).
[20] Neander (2017), chapter 8.
[21] Tye (2000), pp. 121–122, 136–137.
[22] Pelli (1990).

views—the individual "bits of phenomena dust" cannot be regarded as the transparent perception—or misperception—of worldly bits of dust. Their variation over time is, for all practical and theoretical purposes, completely uncorrelated with any external variation. The strongest tracking claim that can be made will involve the *average value* of lightness and hue within a given non-minimal visual region. But when we pay attention, we can become aware of more than these averages.

It is true that the visual system can indeed represent sparkly dusty things floating about in the air in front of us. But the important question isn't whether it can do this; it is whether it can represent items so small or distant that they occupy the vanishingly small minimal visual regions (in terms of visual angle subtended) that characterize the tiny inhomogeneities of visual noise. In order to support an affirmative answer to this question, an advocate of representationalism who adopts a tracking view of perceptual experience might offer the following suggestion. Suppose that, as I have admitted is plausible, the average lightness or hue within a *non*-minimal visual region is a representation of external lightness or hue within a corresponding *non*-minimal region. In that case, can't we say that each minimal bit of fluctuating visual noise, as a limiting case of a region, is a representation of the lightness or hue of that region? After all, advocates of tracking views typically hold that only certain *basic* states get their content *directly* from a tracking relation; other states then get their contents compositionally. Why not say that the relevant *basic* states involve non-minimal regions, but that individual bits of noise get their contents "decompositionally", by being understood as tiny sub-regions?

The response just described is, to be sure, not incoherent. But it seems to be motivated entirely by the desire to continue to hold on to transparency. That is, it has no independent plausibility of its own. After all, no one is tempted to hold that the tiny cyan, magenta, yellow, and black dots out of which colored newsprint images are composed represent the colors of anything. Nor is anyone tempted to think Seurat's individual spots of paint (mis)represent the colors of tiny portions of his subjects. Rather, in both of these cases, it is *collections* of dots or spots that are required before any representation can be said to be going on. So, unless one simply wants to preserve representationalism, there is no temptation to take the tiniest visual regions, which are constantly fluctuating, to represent fluctuations in the colors of regions of the external world. That is, even if it is true that the visual system's topographic map functions to compose representational wholes out of more finely-grained stimuli in smaller regions of the visual field, there is no reason to take the most finely-grained regions to be representational themselves.

Some advocates of tracking representational views might worry about my claim that variation over time, in the tiniest regions of the visual field, is uncorrelated with any external variation. After all, even if it is admitted that some—perhaps much—of this variation is noise, isn't at least some of it due to variations in the environment? If this is so, can't the tracking representationalist claim that these fluctuations do indeed function to track changes in the external world, even if the presence of noise prevents them from doing so with much accuracy? Alas for the tracking representationalist, the most plausible answer to this question is 'No'. The reason for this is that mere correlation is not sufficient to vindicate claims about representation, even on tracking views. Rather, the correlation must have other features: features that are

missing when we consider visual noise. Let me make this clearer by offering a couple of examples.

First, let us consider Ruth Millikan's account of representation. Millikan holds that a purportedly representational state must be made use of by other systems in a certain way if it is to count as actually representational. One advantage of paying attention to the "consumer" of the state in this way is that it allows Millikan to acknowledge that representational states need not correlate very robustly with what they represent—even under normal or optimal conditions. The representational state need only be correlated well enough that the system that "consumes" the state benefited from that correlation in a way that explains its evolutionary persistence. And "correlated well enough" is consistent with the existence of a lot of false positives.[23] For example, it might be that sudden noises produce a kind of mental state in mice that causes a fleeing response. Even if that mental state is only very loosely correlated with the presence of a predator, it might still be taken to represent that, because it is the correlation with predators—even though it is weak—that explains the evolutionary persistence of the consumer system. But, to return to visual representation, the weak correlation of, say, the colors of the bits of "phenomenal dust" with the colors of external objects does not explain the evolutionary persistence of any consumer system. The only correlations that do that will involve larger regions. This point can be reinforced when one considers that the colors of bits of phenomenal dust do not persist long enough for any consumer system to make practical use of them.

A second sort of tracking account is defended by Tye. He holds that, to support a representational claim, the relevant correlation must be explained in such a way that, "if optimal conditions were to obtain, [the representing state] *S* would be tokened in [creature] *c* if and only if *P* were the case; moreover, in these circumstances, *S* would be tokened in *c* because *P* is the case".[24] Since optimal conditions may not often obtain, this account, like Millikan's, makes representation consistent with weak correlation. Nevertheless, Tye's view cannot provide any defense of the claim that bits of phenomenal dust represent bits of the external world. Given the stochastic way the visual system works, there are no conditions—neither normal nor optimal—that support the counterfactual claim that under those conditions the color of a bit of phenomenal dust would match the color of a bit of the external world. Under any circumstances in which the visual system continues to be the actual human visual system with which we are concerned, the colors of the phenomenal dust will be fluctuating, even while the visual experience of which they form a part is representing an unchanging scene as well as is humanly possible.

It is worth noting that neither of the arguments I've just offered regarding visual noise work well in the case of blur. Fuzziness is something that is certainly perceptible, and there is no special reason to think it presents a problem for tracking accounts of representation. As a result, it remains possible for the advocate of a tracking account of representation to argue that seeing something blurrily is merely

[23] See Millikan (1989), pp. 283 and 288.

[24] Tye (2000), p. 136.

a matter of misrepresenting it as fuzzy. I do not say how plausible this strategy is, but it is at least available. But, as I have been arguing, the parallel strategy in the case of fine-grained visual noise is not available. When objects are so small, or so distant, that they subtend visual angles as small as bits of phenomenal dust, their colors are *not* perceptible—at least on tracking accounts. If we accept a tracking account, we cannot explain visual noise in terms of the perceptual representations of such tiny or distant bits of colored dust.

4.2 Phenomenal intentionality

Distinct from tracking views, a second view of the relation of visual experience to its content goes by the name *phenomenal intentionality*. Advocates of this view deny that the phenomenal nature of experience is grounded in its intentional content—a content that is in turn explained in terms of some naturalistic sort of tracking relation. Rather, those who endorse phenomenal intentionality hold that intentionality is grounded, somehow, in the very nature of phenomenal consciousness.[25] The "directedness" of phenomenal mental states (a directness *on objects* in the case of perceptual states) is an unexplained explainer, and is the ultimate source of all intentionality, including the intentionality of language, pictures, and so on. If such a view is correct, couldn't it be used to support transparency by allowing us to argue that it is just phenomenally apparent that visual noise is a matter of tiny variations in the hue and brightness that external objects are represented as having—or of bits of fine roiling dust in front of them? The fact that such phenomenal variations do not serve to track external variations would do nothing to undermine this claim.

Alas for defenders of transparency, I do not think an appeal to phenomenal intentionality will do much good. First of all, arguments in favor of phenomenal intentionality do not depend on transparency in any way, so there is no support provided to transparency in the form of an inference to the best explanation. Rather, those who endorse phenomenal intentionality often appeal to other features of phenomenal states that suggest that they must be intentional. Charles Siewert, for example, claims that phenomenal states are always assessable for representational accuracy, that this is apparent simply from their phenomenal nature, and that this entails that they possess intentional content in virtue of that nature.[26] But Siewert's argument is consistent with these states having other noticeable features that are independent of the content that makes them assessable for accuracy. Another common claim offered in support of phenomenal intentionality is that it is only phenomenal states that can yield determinate content for intentional states. To put the point somewhat sloppily, the idea is that an appeal to tracking relations unavoidably leaves it unclear whether we mean 'plus' or 'quus' by our use of the symbol '+', or whether 'gavagai' means 'rabbit' or 'undetached rabbit parts'. What then does determine what we mean? Advocates of phenomenal intentionality claim that it is the phenomenal

[25] See Horgan and Tienson (2002), Loar (2003) and Pautz (2013).
[26] Siewert (2004), p. 26.

nature of our thoughts and perceptions, which have an "original intentionality" that—they claim—avoids worries about indeterminacy.[27] But again, this argument for phenomenal intentionality does nothing to support transparency or strong representationalism.

Not only do arguments in favor of phenomenal intentionality not depend on transparency, advocates of phenomenal intentionality typically *deny* transparency—at least when when they mention it at all.[28] This is because they tend to hold that the intentionality of conscious experience involves two things: (1) a content and (2) a *way*, or *mode*, or *aspect*, under which that content is presented. Awareness of such a way, or mode, or aspect, is not awareness of the properties of the external objects we seem to perceive. But those who hold to transparency claim that we are only ever aware, in perceptual experiences, of such objects and properties.

On reflection, the denial of transparency by proponents of phenomenal intentionality is not particularly surprising. Advocates of phenomenal intentionality are not attempting to *explain* phenomenal character in terms of representational content—as those who offer naturalistic tracking views of representation are. As a result, they are under no pressure to find representational contents to match any particular phenomenal feature. All they have to hold is that phenomenal states are intrinsically representational. But that is perfectly consistent with the claim that many of their features are non-representational: these might be the modes or aspects under which the intentional objects are presented. What the endorsement of phenomenal intentionality rules out is only the claim that those non-representational features are capable, on their own, of fully characterizing a "stand-alone" phenomenal experience: an experience that would, then, lack any representational content. Because those who advocate phenomenal intentionality are not compelled to explain all phenomenal features in terms of representational content, when such a theorist notices visual noise, she can simply say, "Well, that's part of the way in which visual experience presents the wall". Such a theorist has no reason to say that the noise is a systematic misrepresentation.

5 Conclusion

Visual noise casts considerable doubt on three of the four claims internal to the thesis of transparency. In particular, attention to visual noise makes it plausible that

(1′) We *are* aware of a feature of our visual experience: that it is characterized by noise with a certain degree of graininess.[29]

[27] I should register that I do not find this argument at all persuasive. It seems to take the form of an argument from elimination: nothing else can determine content, so phenomenal character must do so. But the same arguments that eliminate other candidates seem to me to eliminate phenomenal character as well, leaving us with the original puzzle about determinate content.

[28] Loar (2003), Siewert (2004), and Crane (2000).

[29] I do not mean that we *see* it, as we see external objects. Rather, in noting visual noise, we are aware of a feature of the *way* in which we see external objects.

(2′) We *can* attend to this feature of visual experience.
(3′) It is not true that the only features of which we are aware and to which we can attend are external features (colors and shapes of surfaces, for example).

It must be admitted that visual noise does not provide a strict counterexample to transparency, or its related form of representationalism. But that isn't really a weakness of the argument. As Fred Dretske points out, putative counterexamples to transparency or representationalism can *always* be accommodated by claiming—with varying degrees of plausibility—that the phenomenal feature at issue represents *something*.[30] In the case of visual noise, that *something* might be the state of one's retina, or a transparent but "noisy" membrane that constantly interposes itself between one's eyes and whatever one is looking at. As long as one refuses to commit to any substantive account of what it is that gives particular visual experiences their particular contents, such suggestions will always be available to advocates of transparency who are willing to dig in their heels.

Where does the phenomenological stalemate as regards transparency leave us? Well, the case for representationalism on the basis of transparency is an instance of inference to the best explanation. Such arguments are peculiar in that the conclusion—the "best explanation"—can bolster the case for the premise, as long as that conclusion has its own independent plausibility. So a lot will depend on the relevant representational account visual content—the one that provides the "best explanation". If it has something plausible to say about what visual noise represents, this will bolster the case for transparency; if it lacks anything plausible to say, this will undermine the case for transparency. This was the point of the discussions, in Sect. 4, of phenomenal intentionality and of tracking views. With a tracking view in mind, it becomes important that visual noise is not the product of selection pressures; it is a by-product of the stochastic processes that underlie visual perception. And with phenomenal intentionality in mind, it becomes important that there is no need—nor indeed any motivation—to postulate a property that corresponds to the noise that characterizes our perception of external objects. Is there some other sort of account of visual representation that would support the idea that visual noise represents, say, the fluctuating state of the individual rods and cones in one's retina, or tiny bits of a transparent-yet-colored aether that pervades all of visual space? Perhaps. But until someone produces such a surprising account, visual noise remains a serious challenge to those who are drawn to representationalism. And as long as it remains such a challenge, transparency should not be regarded as firm ground on which other theories can be built.

References

Allen, K. (2011). Revelation and the nature of colour. *Dialectica, 65*(2), 153–176.
Allen, K. (2013). Blur. *Philosophical Studies, 162*(2), 257–273.
Allen, K. (2016). *A naïve realist theory of colour*. Oxford: Oxford University Press.

[30] Dretske (2003), p. 80. In fact, Dretske's point concerns representationalism, but it applies equally well to transparency.

Bourget, D. (2015). Representationalism, perceptual distortion and the limits of phenomenal concepts. *Canadian Journal of Philosophy, 45*(1), 16–36.
Bourget, D., & Mendelovici, A. (2014). Tracking representationalism: Lycan, Dretske, and Tye. In A. Bailey (Ed.), *Philosophy of mind: The key thinkers* (pp. 209–238). New York: Bloomsbury Academic.
Crane, T. (2000). Introspection, intentionality, and the transparency of experience. *Philosophical Topics, 28*(2), 49–67.
Dretske, F. (1988). *Explaining behavior*. Cambridge, MA: MIT Press.
Dretske, F. (2003). Experience as representation. *Philosophical Issues, 13,* 67–82.
Fechner, G. T. (1860/1966). In D. Howes & E. Boring (Eds.), *Elements of psychophysics* (Vol. 1) (H. Adler, Trans.). New York: Holt, Rinehart and Winston.
Freud, S. (1913). *The interpretation of dreams* (A. Brill, Trans.). New York: MacMillan Co.
Harman, G. (1990). The intrinsic quality of experience. In J. Tomberlin (Ed.), *Philosophical perspectives 4: Action theory and philosophy of mind* (pp. 31–52). Atascadero, CA: Ridgeview.
Helmholtz, H. (1856/1909/1962). In J. Southhall (Ed.), *Helmholtz's treatise on physiological optics*. Mineola, NY: Dover.
Horgan, T., & Tienson, J. (2002). The intentionality of phenomenology and the phenomenology of intentionality. In D. Chalmers (Ed.), *Philosophy of mind: Classical and contemporary readings* (pp. 520–533). Oxford: Oxford University Press.
Kennedy, M. (2009). Heirs of nothing: The implications of transparency. *Philosophy and Phenomenological Research, 79*(3), 574–604.
Kind, A. (2003). What's so transparent about transparency? *Philosophical Studies*, *115*(3), 225–244.
Ladd, G. T. (1903). Direct control of the 'retinal field': Report on three cases. *Psychological Review, 10*(2), 139–149.
Loar, B. (2003). Phenomenal intentionality as the basis for mental content. In M. Hahn & B. Ramberg (Eds.), *Reflections and replies: Essays on the philosophy of Tyler burge* (pp. 229–258). Cambridge, MA: MIT Press.
Martin, M. (2004). The limits of self-awareness. *Philosophical Studies, 120*(1), 37–89.
Millikan, R. (1989). Biosemantics. *Journal of Philosophy, 86*(6), 281–297.
Neander, K. (2017). *A mark of the mental: In defense of informational teleosemantics*. Cambridge, MA: MIT Press.
Pace, M. (2007). Blurred vision and the transparency of experience. *Pacific Philosophical Quarterly, 88*(3), 328–354.
Pautz, A. (2013). Does phenomenology ground mental content? In U. Kriegel (Ed.), *Phenomenal intentionality* (pp. 194–234). Oxford: Oxford University Press.
Pelli, D. (1990). The quantum efficiency of vision. In C. Blakemore (Ed.), *Vision: Coding and efficiency* (pp. 3–24). Cambridge: Cambridge University Press.
Schroer, R. (2002). Seeing it all clearly: The real story on blurry vision. *American Philosophical Quarterly, 39*(3), 297–301.
Schwitzgebel, E. (2011). *Perplexities of consciousness*. Cambridge, MA: MIT Press.
Siewert, C. (2004). Is experience transparent? *Philosophical Studies, 117*(1), 15–41.
Smith, A. (2008). Translucent experiences. *Philosophical Studies, 140*(2), 197–212.
Tye, M. (2000). *Consciousness, color, and content*. Cambridge, MA: MIT Press.
Tye, M. (2003). Blurry images, double vision, and other oddities: New problems for representationalism?'. In Q. Smith & A. Josic (Eds.), *Consciousness: New philosophical perspectives* (pp. 7–32). Oxford: Oxford University Press.
Tye, M. (2014). Transparency, qualia realism and representationalism. *Philosophical Studies, 170*(1), 39–57.
Volkmann, A. (1846). Sehen. In R. Wagner (Ed.), *Handworterbuch der physiologie* (Vol. 3, pp. 265–341). Braunschweig: Vieweg.
Wundt, W. (1880). *Grundzüge der physiologischen Psychologie* (Vol. 2). Leipzig: Engelmann.