Phenom Cogn Sci (2007) 6:369–387
DOI 10.1007/s11097-007-9050-6

Representationalism and indeterminate perceptual content

John Dilworth

Published online: 26 January 2007

Abstract Representationalists currently cannot explain counter-examples that involve *indeterminate* perceptual content, but a *double content* (DC) view is more promising. Four related cases of perceptual imprecision are used to outline the DC view, which also applies to imprecise photographic content. Next, inadequacies in the more standard single content (SC) view are demonstrated. The results are then generalized so as to apply to the content of any kinds of non-conventional representation. The paper continues with evidence that a DC account provides a moderate rather than extreme realist account of perception, and it concludes with an initial analysis of the failure of nomic covariance accounts of information in indeterminacy cases.

Key words representationalism · indeterminate perceptual content · levels of content · nomic covariance · direct realism

Informationally speaking, we live in a noisy world. Most, if not all, perception involves informational losses or distortions due to various environmental or contextual factors such as inadequate lighting, foggy conditions, subject or object movement, inadequate eyesight, and so on, with the result that perceptually acquired information about an object X and its qualities is *indeterminate* or *imprecise* in various respects.

But such basic truisms have not yet been adequately theoretically accounted for in informational approaches to perception generally, and in particular in the widely accepted views of *representationalism* and *direct realism* concerning the phenomenal character of perceptual experiences. Among other things, I shall be addressing a

J. Dilworth (✉)
Department of Philosophy, Western Michigan University, Kalamazoo, MI 49008, USA
e-mail: dilworth@wmich.edu

common form of representationalism (thesis R), which holds that phenomenal character is reducible to, or can be completely explained in terms of, representational content.[1]

One of my claims will be that defenders of thesis R currently cannot satisfactorily explain a class of potential counter-examples that involve *degraded* or *imprecise* perceptual content, such as in the case of objects seen blurrily by someone with poor eyesight, or objects seen vaguely in misty conditions. However, my main claim will be, more positively, that these problems for thesis R can be adequately resolved via provision of a more sophisticated view of perceptual content than that which is currently presupposed by supporters of R.

Thus thesis R can be saved from the perils of imprecision, but doing so requires drawing on previously neglected factors in representational theory. Specifically, my claim will be that the representational content of perception must be structured in *two nested levels*, which double content (DC) structure, unlike the standard single level content (SC) account, can adequately explain perceptual vagueness.

To be sure, to avoid charges that it is ad hoc, a DC account must both provide a *general* explanation of any kind of perceptual vagueness, and also it must have at least some plausibility independently of vagueness or imprecision issues. The account to be offered satisfies both requirements.[2]

I shall start by outlining the double content (DC) view via consideration of four closely related cases of perceptual imprecision. Then, after a demonstration that the DC view can also explain imprecise photographic content, inadequacies in the more standard single content (SC) view will be demonstrated. The results will then be generalized so as to apply to the content of any kinds of non-conventional representation. The paper will continue with evidence that a DC account provides a moderate rather than extreme direct realist account of perception, and conclude with an initial analysis of the failure of nomic covariance accounts of information in indeterminacy cases.

Four cases of indistinct representation

To begin, I shall explore the connections between four cases of indistinct or indeterminate perception or representation, each of which results in a similar loss of information about an object, and each of which is phenomenologically closely similar to the others.

1. An object X, having sharp boundaries, is seen through a mist, under which conditions it seemingly has boundaries that are indistinct to degree z.
2. A sharp photograph is taken of the above mist-shrouded object X, which also shows X as seemingly having boundaries that are indistinct to degree z.
3. After the mist clears, a blurred or out of focus picture of X is taken, which photograph also shows X as seemingly having boundaries that are indistinct to degree z.

[1] Michael Tye is a representative author who holds such a position. See, e.g., his two books (1995, 2002a). The view is also sometimes described as 'intentionalism', e.g. by Byrne (2001).

[2] Two recent articles of mine provide a broader introduction to the DC view of content that satisfies the independent plausibility constraint, Dilworth (2005a and 2005b).

4. A viewer of X takes off his glasses, and because of his imperfect unaided eyesight only sees X blurrily, also as seemingly having boundaries that are indistinct to degree z.

Cases 1 and 4 involve perceptual representation, whereas cases 2 and 3 involve photographic representation. Nevertheless, it seems to me that a broadly physicalist, representationalist and information-theoretic view of perception ought to seek a common representational explanation of all four cases, since all of them involve a closely similar loss of visual information, to degree z, about the precise boundaries of object X, while at the same time there are no significant phenomenological differences between the cases that might otherwise introduce perturbing factors.

Thus the primary issue is that of how to explain their common informational loss in some uniform representational way.[3] For after all, blurred vision is generally the result of unsharp focusing of light upon the retina by the lens of the eye, and hence it has a closely similar optical explanation to that of the unsharp photographs resulting from a poorly focused photographic lens.

As a preliminary, an intuitive concept of *transparent representation* will be useful as a standard of comparison. Transparent representation would involve complete preservation, or zero loss, of information regarding its subject. For example, a high-quality color photograph of object X, taken under clear conditions, could approach the ideal of supplying complete, undistorted or transparent information about the visual appearance of X from the perspective of the camera lens, as can normal perception of X from the same position by someone having 20–20 vision. The concept is an intuitive one, because, for instance, looking at a fully transparent photograph of a scene would be phenomenologically exactly like looking at the actual scene itself.[4]

Then various interfering conditions, such as those outlined in cases 1–4 above, serve to *diminish* the transparency of representation in each case, producing equivalent degrees of *opacity* of representation with respect to those represented qualities about which information is lost or distorted in the relevant processes.

But how is this kind of information-losing representational opacity to be best explained or understood? Fortunately, broadly information-theoretic approaches to representation already have a central concept, that of *noise*, which can provide a significant initial ingredient in an explanation.[5]

Informational noise arises when an informational channel, such as one carrying visual information about the appearance of an object X, becomes intermixed with information from some *independent* source, such as the misty weather conditions

[3] I shall only consider apparently non-conceptualized perceptual information in this paper, since vagueness problems are most acute for information in this form.

[4] This concept of transparency should be distinguished from that usually employed in debates between phenomenalists and representationalists. E.g., according to Tye, transparency of experience of the surfaces of objects occurs when "None of the qualities of which you are directly aware in seeing the various surfaces look to you to be qualities of your experience." (Tye, 2002a p. 46) – a view according to which perception is 'transparent' just in case it seems to the perceiver to be exclusively object-directed rather than to be, to any degree, about the perceiver's own experience as such.

[5] On information-theoretic approaches to representation, including a brief discussion of noise, see Dretske (1981) Ch. 1, esp. fns 9 p. 238 and 13 p. 239.

involved in cases 1–2 above, or the optical distorting factors involved in the imperfectly focused photograph or blurred vision of cases 2–3 above.

The concept of *noise* employed here – of relevant distorting informational factors independent of the target object X – may be generalized as follows. All of the relevant aspectual or contextual conditions under which an object X is perceived may be regarded as making up a single independent channel of information, the physical basis of which is causally combined with the physical basis of information about object X prior to perception occurring. The result of this intermixing – to be called *aspectual* information – is a single resultant channel of information containing intermixed information both about the independent distorting factors, and, in consequently distorted form, information about the object X.

Now clearly, one central cognitive task for a perceptual system receiving such a mixed aspectual channel of information must be that of employing some technique to *separate out*, as far as possible, the independent informational contributions of the distorting aspectual conditions, and of object X itself. As a further example, in a telephone conversation involving a constantly noisy line, one will be unable to decipher what the caller is saying without having some cognitive techniques available to sort out which factors of the heard sounds are spoken words, as opposed to those other aspectual factors of the same sounds that are non-linguistic sonic factors, independently caused by the noisy line.

Thus, I would claim that, to the extent that perception is a representational process involving cognitive representations that have a perceptual content, operating in an imperfect world in which object-related informational losses are commonplace, the relevant content must have a *two level nested structure*, with the top level being made up the comprehensive aspectual informational content as discussed above, and with the lower level being made up of the purely object-related content with respect to object X – a *double content* view of perception.[6]

The postulation of such a two-level structure of content is unavoidable for several reasons.[7] First, it is undeniable that we do perceive comprehensive aspectual 'packages' such as mist-enshrouded objects or noisy conversations. Second, a perception-based understanding of what is seen or heard in such cases would be impossible unless specifically *object-related* content occurs as an independent level of content in perception, since seeing or hearing aspectual packages alone would not by itself provide any object-related information at all – as in the case of hearing an exceptionally noise-ridden conversation, from which no intelligible linguistic information could be extracted.

Indeed, arguably *all* low level perceptual contents are thus indeterminate, requiring double content processing. Suppose that one attempts to see the *shape* of a round disk. Now the retinal image of such a disk may not be round at all, since it only would be round if the disc were viewed perpendicularly to its surface, since any other viewing angle would produce elliptical retinal data instead. So somehow one must also use that same retinal sensory data to estimate the angle at which one is

[6] For initial developments of this view, independent of noise considerations, see Dilworth (2005a and 2005b).

[7] Also see Dilworth (2005a and 2005b).

viewing the disk, in order to determine the actual shape of the disk. But that needed aspectual angle of view is itself equally indeterminate, in the absence of information about the actual shape of the object. Or in other words, raw sensory information by itself is too indeterminate to specify any definite shape-related perceptual content, since any given shape on the retina could have been caused by many different combinations of actual shapes plus different angles of view – a given retinal ellipse might be the image of a round disk viewed at an oblique angle, or of a similar ellipse viewed perpendicularly, and so on.

Moreover, similar points would apply to any raw *color* retinal data as well. A yellow sensory stimulation might be the result of a white surface viewed in yellow light, or a yellow surface viewed in white light, or many other combinations of surface color plus aspectual lighting conditions. But since arguable all visual sensory data is restricted to data about colors or shapes, *all* retinal data is indeterminate because of interactions between aspectual versus object-related factors.

Applying these points to the current four kinds of cases of apparently blurred shapes, the relevant information processing problem is as to whether one is seeing relatively noise-free, actually blurry shapes – such as in the outlines of some clouds – on the one hand, or on the other hand sharp-edged objects with much aspectual noise, as is actually the case in the examples given, or various intermediate cases of partly blurred objects plus intermediate levels of aspectual noise. But perception would be useless to us unless the perceptual process terminated in some particular *decision* (whether right or wrong) about what the facts are – so the resultant perceptual contents must have two distinctive levels, with each level involving contents embodying a decision about the most likely object-related features on the one hand, and aspectual or noise-related features on the other, that would best jointly explain the initial, low level retinal data.

Thus, to summarize and further explain, in the real world an object X occurs in causal interaction with various aspectual or contextual factors Y. Hence the resultant perceptual data A reaching the senses directly represents, not X itself and its intrinsic features, but only some *contextually modified aspect* Y(X) of X. That aspectual perceptual data A causes the perceiver to be in a concrete perceptual state S, which state directly represents aspect Y(X) rather than X itself. Hence the higher-level representational or informational content of state S will be of form Y'(X'), for aspectual content Y' and X-related content X'.

To continue, that account of the perceptual content of state S is not yet complete, because so far no X-related information *as such* is available to perception at all. The content Y'(X') of state S provides X-related information only in an *aspectually transformed* or *encoded* form. For object-related perception to be possible, appropriate decoding techniques must be cognitively available to enable the X-related information X' to be extracted.[8]

The resultant decoded information X' must also, as discussed above, be available in perception as a distinct, lower level of content nested within the higher level content Y'(X') – which nesting occurs because the X-related information X' is decoded or extracted from the aspectual information Y'(X'), and not vice-versa.

[8] See Dilworth (2005b) for discussions of how such techniques may work.

Parenthetically, no fundamental advances in representational theory are required in formulating this account, since at least one class of representations having a two-level nested content structure are already familiar, namely pictures *of* pictures, such as a traditional painting A, part of whose subject matter is another painting B, whose own subject is some object C. In such a case, painting A represents both painting B and object C.[9]

Similarly, a perceptual state having a double nested content as described above represents both an aspect Y(X) of X, and object X itself – though in this case there is no further representational relation between Y(X) and X, as there was with the relation of the corresponding objects picture B and subject C in the 'picture of a picture' example.

To be sure, the claim that perception in general has a double content structure may initially seem surprising – entailing, as it does, that all of us have been 'seeing double', in the relevant sense, all of our lives without realizing it. However, as presented here, the primary function of the double content structure of perception is to ensure that we can focus on perceiving objects and their characteristics, free of the irrelevant 'noisy' distractions of aspectual factors – which must be involved in initial perceptual processing, and remain consciously accessible to us, but which we are free to ignore once they have served their purpose of delivering object-related information.

One further consideration should be mentioned. Though the argument was developed around cases involving substantive losses of object-related information, evolutionary considerations practically guarantee that, if double-content perceptual mechanisms are required for any significant percentage of perceptual cases, then they will be used for all cases, even for relatively trivial cases in which their use would seem to be redundant.

As further confirmation of this, the following argument could be invoked. A perceptual system would have no way of knowing ahead of time whether an upcoming stream of perceptual data would categorically require mechanisms for separating out object-related from aspectual information. Indeed, it is only through the use of such mechanisms that the system could separate out relatively trivial cases (with close to zero relevance of aspectual information items) from non-trivial cases. Hence such systems are needed in any case to ensure adequate perceptual sensitivity to such significant informational differences.

The nested content of photographs

Now that perception has been shown to have in general a nested, double content structure, a similar thesis will briefly be argued for photographs, hence completing the announced program in "Four cases of indistinct representation" section of providing a similar explanation for both photographic and perceptual cases of indistinct representation.

To begin, recall that the four examples 1–4 in the introduction were organized so that, as far as possible, there would be no *phenomenological* differences between the

[9] See Dilworth (2005b) for a fuller account of nested representations.

four cases. This means that an experimental situation could be set up in which the subject of the experiment did not know which of the four cases he was observing.[10]

Under those conditions, the subject's perceptual system would presumably operate in the same way in all four cases. So if, as previously argued, the subject's perception must have a nested double content structure in the purely perceptual cases, then it must have a similar structure in the photographic cases 2 and 3 as well.

But what of more normal situations in which viewers of photographs know that they are photographs? Here an evolutionary argument becomes relevant: the perceptual methods developed over millions of years of mammalian development are hardly likely to have given way, in the mere 100 years or so since photography was invented, to quite different methods of perceptual processing, specialized in some completely different way specifically for photographic viewing.

Also, early viewers of the first photographs could not have had such deviant special-purpose mechanisms, and the pressures of social conformity would have ensured that later viewers of photographs saw and interpreted them in much the same way as did the photographic pioneers. In addition, human infants learn to see photographs without any specialized psychological education, of a kind that would be needed to train them to use some special or unusual mode of perception for photographs.

It only remains to point out that the issue of the structure of *photographic* content must be a matter of the structure of *perceptual* content during normal episodes of photographic viewing, because a photograph considered merely as a purely physical, causal product of photographic methods has no intrinsic content of its own. It is only the existence of social practices of viewing photographs, and gaining information thereby, that leads us to say that they have a representational content at all.[11]

Nevertheless, in spite of the derived nature of photographic intentionality, it is important, as noted initially, for a broadly physicalist, representationalist and information-theoretic view of perception to seek a common representational explanation both of ordinary perception of non-representations, and of perception of representations such as photographs. Also, facts about the ways in which concrete physical photographs represent things can provide, at least, fruitful analogies to issues concerning the representational structure of phenomenal states, as will become clear later.

The failings of a single content view

Recently Michael Tye has attempted to defend the representationalist thesis R against potentially troublesome counter-examples involving blurred images.[12] However, I shall show that his account cannot adequately address some common aspects of phenomenal character associated with blurred or imprecise images. As a brief

[10] Case 4), of his removing his glasses, could easily be accommodated by having two pairs of glasses, one normal and the other with lenses having no optical effect, with the latter being substituted to set up a case 4) kind of situation.

[11] Thus to this uncontroversial extent the intentionality or representational content of photographs is derived from that of their viewers, independently of debates about whether all intentionality is derived. For a recent defense of such a moderate view of original versus derived intentionality see Beisecker (2002).

[12] Tye (2002a p.80). See also the related discussion in his (2002b).

overview, my main argument will be that his account is an entirely negative one – namely, that the imprecision of imprecise images of objects, as compared to corresponding precise ones, can be explained entirely in terms of the information that they do *not* provide about those objects. However, my main counter-argument will be that imprecision is caused by information-theoretic *noise*, which noise also has a phenomenal presence in blurry images etc., and which also has a significant representational status that must be addressed in any defense of thesis R in such cases.

Now to proceed. Under certain circumstances, perceptual information acquisition occurs under less than optimal conditions. Tye comments on such a case as follows:

> When one sees a sharp object *as* blurry, one sees it as having indistinct contours and boundaries. This, we can agree, is not what normally happens when one unfocuses one's eyes or takes off one's eye glasses. In these cases, one simply loses information. Likewise, when one sees the world through eyes that are half closed. In seeing blurrily, one undergoes sensory representations that fail to specify just where the boundaries and contours lie. Some information that was present with eyes focused is now missing. In particular, the grouped array [of perceptual information] contains less definite information about surface depth, orientation, contours, and so forth.[13]

Several important points can be extracted from this passage. First, we must isolate as a distinct phenomenon a case of seeing a sharp object *as being blurry*, which would be to incorrectly see the object as having indistinct contours and boundaries. Such a perception, if it were correct, would involve no informational imprecision at all; instead, it would be a case of acquiring accurate and precise information about the actual indistinct contours etc. of the object. In a word, it would be acquiring precise information about an imprecise object, rather than imprecise information about a precise object. (The relevant concept of *precise information* will be used as a reference point in the succeeding discussion).

Second, Tye uses five related concepts to describe the intended case of seeing *blurrily* (rather than of seeing an object as being blurry). He mentions that in such cases a) one *loses information*; b) In seeing blurrily, one undergoes sensory representations that *fail to specify* just where the boundaries and contours lie; c) In such cases, some information that was present with eyes focused is now *missing*; and d) The array of perceptual information contains *less definite information* about surface depth etc.than was present in a fully focused case. Tye also later characterizes cases of less definite experience as being cases of e) *representational indeterminacy* (p. 82). Thus, to summarize so far, Tye contrasts, in five related ways a) through e), those experiences in which *indeterminate* information is gained from blurrily seen things with more optimal cases of seeing, in which instead *precise* information is perceptually obtained.

Another significant element of Tye's view is as follows. In a later passage Tye comments that, with respect to such items seen blurrily as described above, one's visual experience "...makes no comment on where exactly the boundaries lie. Hence there is no inaccuracy" (p. 81).

[13] Tye (2002a p.80).

Thus on his view, we must carefully distinguish a *non-committal* perception, that is accurate as far as it goes without being completely specific, from the distinct case of someone seeing a sharp object *as being blurry*, which would be to *incorrectly* see the object itself as having indistinct contours and boundaries.

Now if this account is to be successful as an explanation of imprecise perceptual information, it must eliminate *all* potential sources of imprecision in the perceptual data. But if such cases of indeterminate information could be analyzed as cases of less than complete, but nevertheless completely precise, information, the problem could be solved.

An underlying model for this approach would be that of a *set of precise items of information*. Optimal or fully precise seeing would involved perceptual reception and experience of *all* relevant items of such a set, while seeing blurrily would instead involve perceptual reception and experience only of some proper subset of those precise informational items: a *precise subset* model of indeterminate information.

Such a view seems to be consistent with at least three of Tye's five descriptions of seeing blurrily. Information is a) lost, c) missing, and b) some failures of specification occur, but all of these are compatible both with the absent or failing items being precise, and with the remaining informational items being precise as well.

As for d) less definite information and e) representational indeterminacy, possible incompatibilities with the 'precise subset' model can be imagined, such as cases involving information that is *intrinsically* indefinite, indeterminate or imprecise, in the sense that at least some of its remaining members are not themselves precise.

However, Tye's own examples of representational indeterminacy remain compatible with the precise subset model. Of a verbal description such as "There's a tall man at the door," he says that it "...leaves open whether he is wearing a hat, the look on his face,..and many other features. It simply does not comment on these matters." While in the case of a picture of a man, he says "...I may well leave unspecified how many stripes are on his shirt, the color of his cheeks, whether he is wearing a belt."[14] Not commenting on features, or leaving them unspecified, are ways of providing *no* information on such features, rather than of providing intrinsically indeterminate information about them.

To all appearances, Tye is assuming that a similar analysis applies to cases of objects seen *blurrily*, so that when he comments, as previously quoted, that one's visual experience of blurrily seen objects "...makes no comment on where exactly the boundaries lie,"[15] he seems to be claiming that such a case provides *no* information on where exactly the boundaries lie, rather than claiming that it provides *intrinsically indeterminate* information on where exactly the boundaries lie.

However, it is undeniable that objects seen blurrily generally provide *some* rough information about the positions of object boundaries. As an initial criticism, Tye seems to have been mislead by the logic of phrases such as 'the exact/precise position of X'. Actual object or boundary positions are concrete matters of fact,

[14] Tye (2002a p. 87).

[15] Tye (2002a p. 81).

whereas talk of the exactness or precision of their position involves conceptualizing such facts relative to some scientific or perceptual scheme of *assessment* or measurement of their position. Or in other words, *ontological* issues about positions etc of objects must be distinguished from *epistemic* issues as to degrees of accuracy of assessment of such characteristics.

Hence to say that a perception provides 'no information about the precise position of X' is simply to say that it provides *no precise information* about the position of X. But this is fully consistent with its providing *imprecise information* about X's position, which surely it must since we can indeed gain *some* information about the position of things via seeing them blurrily. Thus we may already provisionally conclude that Tye's apparent adoption of the precise-subset view of perceptual information in seeing-blurrily cases does not succeed, since postulation of the presence of some *intrinsically indeterminate* information is required in order to explain the rough positional information that we can extract from such cases.

In addition, a more central criticism of Tye's analysis is as follows. He simply does not address at all a central fact in such cases, namely the *phenomenal* presence in perception of blurry or fuzzy elements or factors. In seeing objects enveloped in mist, one sees the mist as well as the objects, and one also sees the *complex blurry results* of the dispersion or scattering of light rays from each object as caused by the mist. Tye's analysis entirely ignores both the mist itself, and the complex blurry results that provide, in intermixed aspectual form, information both about the mist and about the object. Such blurry results would constitute, from Tye's single content (SC) perspective, a case of phenomenal differences without any corresponding representational differences, hence violating the representationalist thesis R.

However, thesis R is violated only on the standard single level content (SC) assumption, according to which only purely object-related content has any representational status. If instead the nested double level content (DC) view is adopted, thesis R can be preserved, because on the DC view both the aspectual complex blurry results Y and the object X are represented by the concrete perceptual state, which has perceptual content Y'(X') at the higher level, and X' at the lower level. Hence the phenomenal differences in blurry versus precise cases are completely explained by the DC analysis, since every aspect of the phenomenal changes, whether in Y(X) or X, is explained by corresponding changes in the representational content items Y'(X') and X'. Hence the DC analysis, or at least some view relevantly similar to it, is an indispensable ingredient in any successful defense of thesis R against attempted 'blurry' counter-examples.

Perceptual vagueness as generic representational vagueness

Two tasks remain in order to carry out the program developed in the previous three sections. One is to more explicitly generalize the previous DC analysis of imprecise content so as to cover any kind of representational content. This will be carried out via an investigation of imprecision in photography. The second, concurrent task will be to explain the relevant kinds of perceptual vagueness as being species of more generic kinds of *representational* vagueness, that are ubiquitous in photographic or

other non-conventional representational contexts generally, and which thus have no special relevance to specifically perceptual or other phenomenal contexts.

Hence, it could be argued, such imprecision or vagueness cannot by itself constitute a threat to thesis R – in virtue of providing evidence for cases of phenomenal character that are irreducible to representational content – because the relevant kinds of informational vagueness also occur in *non*-phenomenal cases, for which no alternative phenomenal explanation would be possible.

To be sure, in "The nested content of photographs" section it was argued that photographic content is itself a species of perceptual content, in that it is interpretation of photographs by humans that confers a representational status upon them. Nevertheless, the scientific status of photographic content is much more secure than that for purely internal perceptual phenomena such as after-images or hallucinations. Hence, for example, even those who are skeptical as to whether perception generally is a representational process[16] would presumably concede that specialized perceptual tasks such as the interpretation of photographs are genuine representational tasks.

To begin, the most direct physical analog of cases of seeing blurrily is provided by *photographic* cases of unsharp focus. Seeing blurrily is itself generally the result of inadequate accommodation or focusing in the eyes, which results in light rays not being sharply focused upon the retina. Photographic cases of unsharpness in areas of an image have a similar optical explanation. The relevant light rays from a given point A on the surface of the subject do not converge to a single corresponding point X on the film or digital sensor surface, but instead they are spread across a 'circle of confusion' of non-zero size, having point X as its center. While at the same time, light rays from points B, C,.. adjacent to point A on the subject, also fall on point X, because X lies within their corresponding circles of confusion, with centers Y, Z, ... that are adjacent to X.

The result of these processes, in information-theoretic terms, is that each point X on the photograph has received information from *multiple* points A, B, ... on the surface of the subject, namely all those points such that light from them fell upon point X. But this process also inevitably results in *losses of information*, because the photographic emulsion, or digital sensor array, is unable to separately store information about the precise contribution of each subject point A, B,... Instead it can store only the *summed result* at point X of all of their contributions.[17]

Hence, for example, if point A is much darker than its surrounding points, this information will not be fully preserved at, or recoverable from, point X because information about the relative brightness of the surrounding points will 'swamp' point A's relatively low-level contribution to the summed value at point X. While at the same time, most of the information about a relatively bright point would be lost by its becoming summed with information from less bright points.

[16] E.g., see J. K. O'Regan and A. Noë (2001).

[17] For monochrome or 'black and white' storage of such a summed result, values range from white (full intensity) to black (minimal intensity). If color information is also stored, one standard mode of analysis of colors would distinguish their hue (specific color), intensity (strength of color) and brightness (intensity of light) as independent informational factors.

The preceding discussion was from the perspective of a given point X on the photograph. If instead a given point A on the subject is considered, it would have information about it stored at multiple adjacent points X, Y, ...on the photograph – but as already noted, that information would be subject to losses in each case due to overlapping contributions from A's adjacent points B, C,...

Now it might be hoped that such fundamental informational imprecision factors for particular points could be circumvented by examining, not the light values for such particular points, but instead *average* values for regions. Each region of the photograph, whose size is equal to or greater than the size of an average circle of confusion for the current lens/subject/photograph combination, is such, it might be claimed, that the *average* light value of its points provides precise information about the *average* light values of a corresponding region of the subject.

However, the relevant correspondence is only rough or approximate, because circles of confusion have vague or fuzzy boundaries due to gradual light fall-off at their peripheries, and also variable sizes due to various factors, such as the grain size of the film or sensor, lens aberrations and the varying distance of each distinct subject feature from the lens.

In addition, the precision of any average value is already limited by the imprecision of the available information about each particular point on the subject, from which information the average is computed. Other information-degrading factors, both for averages and particular values, include the fact that even the best available films or digital sensors provide only limited sensitivity to small changes in light values, while at the same time exhibiting grossly non-linear behavior with respect to very large changes in light values – such as in an attempted photograph of a person backlit with the full intensity of the sun, in which light-value distortions are unavoidable. (Related distortions in human perceptual information in such situations are of course equally inevitable).

Thus, as a result, photographic methods, as with any empirical measurement methods, cannot provide *completely* precise information even with respect to *average* light values. Hence any apparent gains in precision by semantic ascent to talk of average rather than particular light-values are entirely illusory. Or, in ontological terms, even if average light reflectance values are genuine qualities or properties of object–regions, such as overall color qualities construed as averages, nevertheless the photographic collection of information about such average values for regions involves substantially the same information-degrading mechanisms as operate also to degrade values for particular points.

There is another problem with the averaging approach to information about object qualities, namely that whatever its virtues or lack thereof, it would in any case produce *significantly worse* results for many qualities of objects, such as their outline shape, which involve relatively sharp or discontinuous changes in light received from an object and its background, over the width of a small region of the order of the size of a molecule. Since normal photographs produced without the use of an electron microscope cannot resolve such sharp discontinuities, averaging techniques applied to the details that the lens/photograph combination can resolve will inevitably produce results that are even more informationally degraded than results for relatively continuous qualities such as overall surface color – to the point that the degraded result becomes perceptually easily noticable, as in looking at an

unsharp photograph of the outline shape of an object, or in direct perception of the same object when its boundaries are seen blurrily because of the inadequate optical characteristics of the observer's eyes.

Thus the outcome of the discussion so far is that the kind of information losses involved in seeing blurrily are at least closely related, if not identical, to the information losses involved in photographic cases of unsharp images. These losses cannot be explained by a precise-subset model of informational loss, since *all* relevant informational factors are degraded in such cases, so that the resulting information is *intrinsically indeterminate* in the sense discussed in "The failings of a single content view" section. To be sure, the degrading effect is much grosser, and hence easier to perceptually notice, for object qualities having a relatively precise location, but some degrading is inevitable for any object qualities.

Hence a natural generalization of the discussion is available, to cover *all* cases of photographic information collection, and not only those gross cases in which the result is easily perceivable as being unsharp. A perception of unsharpness occurs when the relevant overlapping circles of confusion for each point on the subject are perceived to be larger than point-size. However, in the real world any subject/lens/photograph combination whatsoever will result in non-zero size circles of confusion, whether or not these are perceiveable as such with the naked eye. Hence all photographic information is intrinsically indeterminate, whether it is viewed as information about points or as average information about regions.

But insofar as normal human vision is a closely similar optical process, all visual information will be similarly indeterminate. And clearly similar arguments could be made for the indeterminateness of information received through other sensory modalities, with the result that *all human perception* is intrinsically indeterminate in the relevant sense. Furthermore, given that photographic information is one typical subset of scientific information, a similar generalization is possible for it too: even the best scientific methods, such as those used in optical cases, will also inevitably provide only intrinsically indeterminate information.

Direct realism is undermined by indeterminate information

The preceding result, that all perceptual, and in general observational, information is intrinsically indeterminate, does serve to solidify the double content (DC) account as a legitimate theory that can provide a general account of perceptual indeterminacy, rather than its merely being an ad hoc explanation of a few 'blurriness' examples. However, that result also has potentially significant implications that go far beyond the defense of the basic representationalist thesis R itself, as I shall now explain.

Many defenders of thesis R, such as Tye and Dretske, may also be regarded as being in addition *direct realists*, in that they would claim that, when one is veridically seeing an object X and its qualities, then the qualities that one then sees it as having are its actual concrete qualities.[18] Indeed, in Tye's case, this probably

[18] E.g., Tye (2002a p.49): "...there are qualities of which the subjects of visual experiences are directly aware via introspection. They are qualities of external surfaces (and volumes and films) if they are qualities of anything."

provides yet another reason for imputing to him belief in the precise subset view of indeterminacy, according to which any missing information makes no difference to the precision, and hence to the realism, of the remaining informational items.

Presumably, however, such realistic views are plausible only to the extent that any veridical information provided by perception about quality Y of object X provides a *sufficient condition* of object X actually having quality Y. Or in other words, if the information about property Y of object X is veridical, then X must actually possess quality Y. However, the current DC indeterminacy view would deny this entailment, on the grounds that information can be veridical in the sense that it is correct as far as it goes, but nevertheless its intrinsic indeterminacy is such that it cannot entail anything about the *precise* qualities of object X. For example, a blurry view of the edge of object X gives information about its approximate position only; but there is no corresponding quality 'having approximate position Y' actually possessed by object X, which instead has some distinct, completely precise positional quality Z.

This criticism can be sharpened as follows. Insofar as writers such as Dretske and Tye are committed to an information-theoretic approach to perception, they must also respect the epistemic limitations imposed by that approach. For example, as Dretske convincingly argues in his classic book *Knowledge and the Flow of Information*, strictly speaking one does not have genuine *information* about some state of affairs, such as object X having quality F, unless the probability of its being so is 1 – not roughly or almost 1, but exactly 1.

Also, the underlying nomic covariance account of information requires that there must be a *lawlike* connection between a representational state and its cause X for the state to provide genuine information about X. But it is simply false that there is a lawlike connection between some particular position of the edge of object X, and some particular blurred perceptual or photographic configuration that represents it. Instead, the blurred configuration is compatible with an *infinite set* of possible exact positions of the edge of object X, whose members are defined by all those positions that would causally produce exactly the same blurred configuration. Thus insofar as the blurred configuration supplies any genuine information at all in the nomic covariation sense, it is only information that object X possesses one of the infinite range of values of the members of that possible set of positions of X – hardly a direct realist view of perceptual information![19] (The next section will provide some initial clarification of the sense in which a broad nomic covariation approach to information can be maintained under these conditions).

Thus, to summarize, my claim is that a direct realist perceptual view is subject to a significant *epistemic* failure, given that all perceptual information is indeterminate. For even if one *seems* to be perceiving the actual qualities of an object – a fact not in dispute-one's perceptually acquired informational evidence for the stronger claim that one actually *is* perceiving those actual qualities could never be strong enough, given the indeterminacy of the perceptual evidence, to justify that claim. Hence the relevant epistemic version of direct realism is false.

[19] Nor would it help to appeal to *ideal* representational conditions, under which the relevant set would have only one member, namely the actual position of object X, for reasons spelled out by Wallis (1994).

Another kind of challenge by the DC theory to the supposed directness of a broadly conceived direct realism is over the issue of the cognitively *indirect* status of the perceptual information it can provide about objects as such. Recall that on the DC account, all initial reception of perceptual data concerning an object X is in aspectual form Y'(X'), in which the relevant information X' about X is present only in an intermixed or encoded form. In order to correctly perceive object X and its qualities, the perceptual system must first decode that encoded X-related information.

But the need for such a decoding procedure shows that information about objects as such is not directly or immediately present in perception – or at least, not in the way in which low level aspectual information is. Instead, such X-related decoding amounts to the application of at least some minimal kind of *inferential procedure*, which presumably disqualifies the resulting X-related information X' from being immediately or directly present in perception, even though X' does become present in perceptual content – as the nested lower-level content in the DC structure – as a result of the inferential decoding procedure.

Thus there are at least two significant respects in which a DC theory of perception is not an extreme or stereotypical direct realist theory. However, neither of these respects involves a claim that there are intermediate mental objects such as sense-data, so the approach is still a broadly or moderately realist one.

Nomic covariation and indeterminate information

This section provides the promised initial clarification of the sense in which a broad nomic covariation approach to information can be maintained under conditions in which the information is intrinsically indeterminate.

To begin, recall that the underlying nomic covariance account of information requires that there must be a lawlike connection between a representational state S and its cause X for the state S to provide genuine information about X. But in the previous 'blurry photograph of sharp edge' example, it is false that there is a lawlike connection between some particular position of the edge of object X, and some particular blurred perceptual or photographic configuration S that represents it. Instead, the blurred configuration is compatible with an infinite set of possible exact positions of the edge of object X, whose members are defined by all those positions that would causally produce exactly the same blurred configuration. Thus insofar as the blurred configuration supplies any genuine information at all in the nomic covariation sense, it is only information that object X possesses one of the infinite range of values of the members of that possible set of positions of X.

Now so far, arguably the nomic covariation approach might be reinstated, via a law governing the relations of the closed region of space around the sharp edge, enclosing all points in the relevant infinite set, and its blurred photographic image caused by any of those points. However, that simple solution cannot work because of the kinds of causal indeterminacy that produce the indeterminate information, based on underlying problems stemming from three well-known factors.

First, any actual observation of the region could produce only an imprecise concept of its extent, such that, for any particular precise set of points chosen, there would always be external additional points, sufficiently close to the chosen

boundaries, which would also cause the same blurry configuration as the members of the set. Secondly, there would also be further external points for which it was *indeterminate* whether or not they caused the photographic blur. And of course these two indeterminacy factors affect the photographic blur itself, in that its boundaries too are vague in related ways.

There is also a third factor due to quantum effects. There may be points in the edge region such that on some occasions they would cause the photographic blur by appropriately refracting photons of light through the lens, whereas on other occasions they would not, hence in that way undermining the desired nomic necessity. Indeed, such effects may hold for all points in the edge region, not just for borderline cases, so that only approximate statistical averages of light level may be replicable at any given point on a photograph, over a series of repeated tests. In addition, arguably all three of these factors are objective causal indeterminacies, rather than mere artifacts of imperfect empirical testing, since even (counterfactually) perfectly precise observations would encounter the same problems.

Hence if laws are understood in nomic or necessary connection terms, strictly speaking there are no laws in the cases being discussed. This also means that the most currently relevant nomic approach, namely the Armstrong–Dretske–Tooley view of laws, that explains their necessity in terms of relations between universals or second-order properties,[20] is also inapplicable to such cases, in that there cannot be any precisely definable property possessed by points in an interval around an edge that is necessarily connected with another property possessed by the corresponding photographic blur.

The relevance of this point to the issue of the indeterminateness of the relevant photographic positional information is as follows. Failing a precise subset view of informational indeterminacy, as discussed in "The failings of a single content view" section, writers such as Dretske and Tye might have assumed, or hoped for, a fallback position to be defensible, in which apparently indeterminate information about an exact position is instead explained as fully determinate information about a *region surrounding* the position, understood as a precise positional interval.

On such a fallback view, the simplicity of the original nomic covariance view of information would have been restored, in that a nomic connection holding in virtue of a property P-namely, presence of the edge in the relevant precise interval-of the relevant object x would have resulted in complete and precise photographic information that object x has property P. However, the above three causal indeterminacy considerations show that there cannot be any such property P, and hence there cannot be a law linking it to a similarly unavailable property Q defined by the relevant blurred photographic interval or area.

The upshot of these points is that the status of the relevant edge-related information as being irreducibly indeterminate or incomplete information about the *exact* position of the edge has been maintained; it cannot instead be analyzed as

[20] E.g., see Armstrong (1983)

determinate information about some property of an interval that includes the edge. Hence it is still true that there is no edge-related *property* of object x for which the photographic blur provides a sufficient condition, and hence the blur cannot provide any genuine information about such a supposed property.

Nevertheless, I would argue that it is still important to maintain a central insight of the original nomic covariance view of information, namely that genuine veridical information must have a probability of 1. It still remains true in the edge case, intuitively speaking, that the photographic blur provides information that there is some *large enough, yet finitely small*, region in which the edge *must* be located – or in which it *always would* be located – even if the necessity or universality of that 'must' or counterfactual 'would' can no longer be explained as strict nomic necessity.

That is why even blurred photographs can still provide genuine positional information that must be correct as far as it goes, even though the information is indeterminate. It is genuine information insofar as a) the region appealed to is large enough so that in *all* possible cases the edge must be in that region, hence ensuring a probability of 1, and b) the region is some finite proper subset of all of space, whose specificity or lack of indeterminateness increases as the size of that region decreases.

Also, such genuine though indeterminate information must still somehow be distinguished, as previously in discussions of informational nomic covariance, from pseudo-information provided by mere accidental but universally holding correlations that provide no genuine information whatsoever.[21]

To be sure, some independent basis is now needed to defend that contrast, given that a simple nomic versus non-nomic distinction is no longer available for the purpose. But considerations of causal determination or relevance versus irrelevance are still available, at least in cases involving issues of information. In effect my "Four cases of indistinct representation" section discussion assumed some such account of causally determined but not strictly lawlike behavior, in treating cases of indeterminate but genuine information as arising via causal interactions of objects with other aspectual factors, that served to degrade but not completely destroy the resultant object-related information. Thus an object X is able to supply information about itself by its causing changes in aspectual factors Y, which changes in turn cause a perceptual state S to arise in some cognitive system.

Hence there are two possible kinds of deviance from a nomic relation between X and S. First, a gross deviance, not previously considered, in that the precise state S caused by X might vary widely, depending on which intervening aspectual factors happen to obtain. And second, even if such aspectual factors are held fixed, as in the main discussions here, there will still be more subtle deviations from lawfulness in the X to S relation, arising from the causal interactions of X, Y and the photograph, which interactions both explain why the resultant information about some property of X is not complete, and also explain why the information is intrinsically indeterminate (as a result of the three kinds of causal vagueness discussed).

[21] E.g., as discussed by Dretske (1981).

Summary

The results of the paper will now very briefly be summarized in a wider context. First, the discussion of informational indeterminacy will be related to broader considerations about vagueness. It is generally agreed that issues about e.g. the vagueness of a distinction between heaps and non-heaps involve primarily conceptual or epistemic rather than ontological factors. So it is prima facie implausible to suppose that cases of vagueness or indeterminateness are to be explained in terms of objectively real vague objects or properties – such a simplistic form of direct realism about vagueness is hard to defend, and certainly it is not a reasonable default or initial assumption.

But this point lends added force to my discussion of *informational* vagueness or indeterminacy. The assumption by representationalists such as Tye that cases of perceptual vagueness are compatible with a perceptual direct realism can, in this perspective, be seen to be a highly questionable initial view, ripe for the kinds of refutation to which it has been subjected in this paper.

To be sure, some people, including Tye himself, do hold that there are cases of genuine ontological vagueness, rather than all vagueness being merely an epistemic or conceptual matter.[22] However, even if all objects and properties were ontologically vague, this would be irrelevant to *informational* indeterminacy or vagueness, because information about such properties as the supposedly objectively vague position of an edge would still be subject to the additional layer of informational indeterminacy discussed in this paper, so that perceptual direct realism about such properties would be just as false as it is for completely precise objects or properties.[23] Thus, as a result, the prima facie implausibility of direct realist views of vagueness as applied to informational indeterminateness persists, whether or not ordinary non-informational kinds of vagueness can be ontologically explained.

In addition to undermining the supposed direct realist implications of representationalist views in vagueness cases, this paper has also argued that representationalism itself cannot be defended against likely phenomenalist or qualia-based attacks – such as those based on phenomenal vague qualities, to which no vague properties of an object correspond – without a comprehensive theoretical restructuring of the kind provided by the current general analysis of indeterminate information, plus the accompanying double content (DC) view of perceptual contents.

This paper has also argued that the underlying nomic covariance view of information, which provides the basis for appeals to informational concepts by some major representationalists such as Dretske and Tye, is also in need of a comprehensive overhaul, in order to adequately account for the pervasive *noisiness* of perceptual information – an ironic result indeed, in that purely mathematical theories of information flow, such as that of Claude Shannon, are of course fully adequate in describing either noisy or non-noisy information channels.

[22] E.g., see Tye (1990).

[23] A similar point about the inapplicability of direct realism would also apply to 'average properties,' if any, as discussed in "Perceptual vagueness as generic representational vagueness" section.

But the attempts by Dretske and others to build a broadly semantic informational theory on those mathematical foundations have yet to come to terms with the fact that most, if not all, perceptual information is genuinely indeterminate, in spite of its semantic correctness or incorrectness. It will be work for future papers to provide a more constructive theory of semantic information that is compatible with these indeterminacy results.

References

Armstrong, D. M. (1983). *What is a law of nature?* Cambridge: Cambridge University Press.
Beisecker, D. (2002). Dennett and the quest for real meaning: in defense of a 'myth'1. *Philosophy in the Contemporary World, 9*, 12–18.
Byrne, A. (2001). Intentionalism defended. *Philosophical Review, 110*, 199–240.
Dilworth, J. (2005a). The twofold orientational structure of perception. *Philosophical Psychology, 18*, 187–203.
Dilworth, J. (2005b). The double content of perception. *Synthese, 146*, 225–243.
Dretske, F. (1981). *Knowledge and the flow of information*. Cambridge, MA: MIT.
O'Regan, J. K., & Noë, A. (2001). A sensorimotor account of vision and visual consciousness. *Behavioral and Brain Sciences, 24*, 939–973.
Tye, M. (1990). Vague objects. *Mind, 99*, 535–557.
Tye, M. (1995). *Ten problems of consciousness*. Cambridge, MA: MIT.
Tye, M. (2002a). *Consciousness, color and content*. Cambridge, MA: MIT.
Tye, M. (2002b). Representationalism and the transparency of experience. *Nous, 36*, 137–151.
Wallis, C. (1994). Representation and the imperfect ideal. *Philosophy of Science, 61*, 407–428.