Webster, M.A., D. Kaping, Y. Mizokami and P. Duhamel. 2004. Adaptation to natural facial categories. *Nature* 428: 557–61.
Wexley, K.N., G.A. Yukl, S.Z. Kovacs and R.E. Sanders. 1972. Importance of contrast effects in employment interviews. *Journal of Applied Psychology* 56: 45–48.
Wilson, T.D., D.A. Reinhard, E.C. Westgate, et al. 2014. Just think: the challenges of the disengaged mind. *Science* 345: 75–77.
Wright, L. 1973. Functions. *Philosophical Review* 82: 139–68.
Xu, Y. 2018. A tale of two visual systems: invariant and adaptive visual information representations in the primate brain. *Annual Review of Vision Science* 4: 311–36.

Responses to my critics

Ned Block

1. Adaptation, signal detection and the purposes of perception: reply to Ian Phillips and Chaz Firestone

Ian Phillips and Chaz Firestone have written a wonderful article on the rationale for adaptation as an indicator of perception, and more generally, on the purpose of perception, full of insights and challenges.

1.1 Adaptation

The issue they raise that I find the most interesting and challenging, and that I didn't say enough about in the book, is whether there is any independent justification for adaptation as an indicator of perception or whether my reliance on phenomenology (and also retinotopy) to ground adaptation makes adaptation superfluous.

I will approach the issue by reminding the reader of my three-layer methodology as explained in Chapter 1.

Here are the three layers: (i) use armchair criteria of perception and of cognition to roughly delineate the categories of perception and cognition. (ii) Use those categories to isolate scientific indicators. In particular, I chose perceptual adaptation, rivalry, pop-out, illusory contours and speed of processing, but as I indicated, I could have picked many other indicators. (iii) Use the scientific indicators to isolate the underlying constitutive features of perception and of cognition.

As I also explain, the use of a variety of scientific indicators raises a problem of circularity. The problem is that the justification of any given indicator depends on invocations of other indicators. I argued that the circularity is benign so long as the indicators converge on the same results and those results match up better with the armchair criteria than they would have with alternatives.

My case for benign circularity is threatened by the issue raised by Phillips and Firestone of whether some of the indicators play no real role at all. In particular, do I validate adaptation by appealing to retinotopy and phenom-

Analysis Vol. 83 | Number 3 | July 2023 | 575–588 https://doi.org/10.1093/analys/anac058

enology, where retinotopy and phenomenology just stand on their own? If so, they argue, adaptation has no independent significance as an indicator.

I'll start with a case that fits their criticism in which an indicator has no independent significance: in binocular rivalry, a type of eye movement called optokinetic nystagmus is used as an indicator of which stimulus the subject is consciously aware of. Of course, the eye movements' role as indicators is validated by the subject's phenomenological judgements, but the eye movements are nonetheless incredibly useful in circumstances in which we do not want the subject to make such judgements, because these judgements would register on brain scans and make it more difficult to isolate the neural basis of the experience itself (this methodology is known as the 'no report' paradigm). The general point is that we can have an extremely useful non-phenomenological indicator that is entirely based on phenomenology. Still, optokinetic nystagmus has no independent significance as an indicator. Is adaptation in the same derivative boat as optokinetic nystagmus? I will explain why adaptation is different.

To approach this issue, we need to distinguish between high-level and low-level perceptual phenomenology. Recall that low-level perception is the immediate product of sensory transduction and is the causal basis of high-level perception, that is, perception that is not the immediate product of sensory transduction but rather depends on further processing. Notable low-level perceptual representations represent shape, size, texture, hue, motion, depth, contrast and brightness. High-level perceptions mentioned in the book are perceptions of faces, emotional expressions, causation and numerosity.

Look at Figure 1 in Phillips' and Firestone's article (reproduced here for convenience), specifically at the top halves of parts A, B and C. In A(i), it is perfectly clear from the first person that one is perceiving colours, with hues of red, yellow, blue and green. In B(i), one perceives slightly clockwise tilts on the top and slightly counter-clockwise tilts on the bottom. And in C(i), one perceives red and black vertical stripes and green and black horizontal stripes. We can be sure from the phenomenology that we visually attribute these qualities. If one stares at the (i) parts of these figures and then at the (ii) parts, one experiences the adaptation phenomena in which the perceptual attributions change, but the change in perceptual attributions carries no force that is not already carried by the phenomenology of the pre-adaptation perception. So far, adaptation is playing no real role, just as Phillips and Firestone suggest.

Although the phenomenology of perception is sufficiently recognizable to tell us that we are perceiving, it isn't much use in distinguishing low-level from high-level perception or either one from the non-ampliative perceptual judgement that is most directly based on perception this is minimal, immediate, direct perceptual judgement; see the critique by Steven Gross in this issue for more discussion of this kind of perceptual judgement. When I say 'perceptual judgement' in what follows, I mean minimal immediate direct perceptual judgement).

Something can look expensive or look like a piano or a pine tree (examples I use in the book). But one cannot tell from the phenomenology whether

when something looks expensive; that stems from perceptually representing it *as expensive* – or whether what one is experiencing is the phenomenology of the colour, shape, texture and other low-level properties, together with the perceptual judgement that it is expensive (or is a piano or a pine tree). (This issue is also discussed in (Begby 2011, Block 2014, Siegel and Byrne 2016)). That is, introspection is useless in deciding between (i) high-level perception and (ii) low-level perception without high-level perception but with high-level perceptual judgement. Of course, there will be high-level perceptual judgement in both (i) and (ii), but in (ii), the high-level perceptual judgement is based on low-level perception.

Without appealing to adaptation, we can't use introspection to distinguish high-level phenomenology from low-level phenomenology, and we can't use introspection to distinguish either one from perceptual judgement.

So how exactly does adaptation help with the problem? Look at the adaptation stimuli in Figure 1D. After staring at a fixation point in the top image

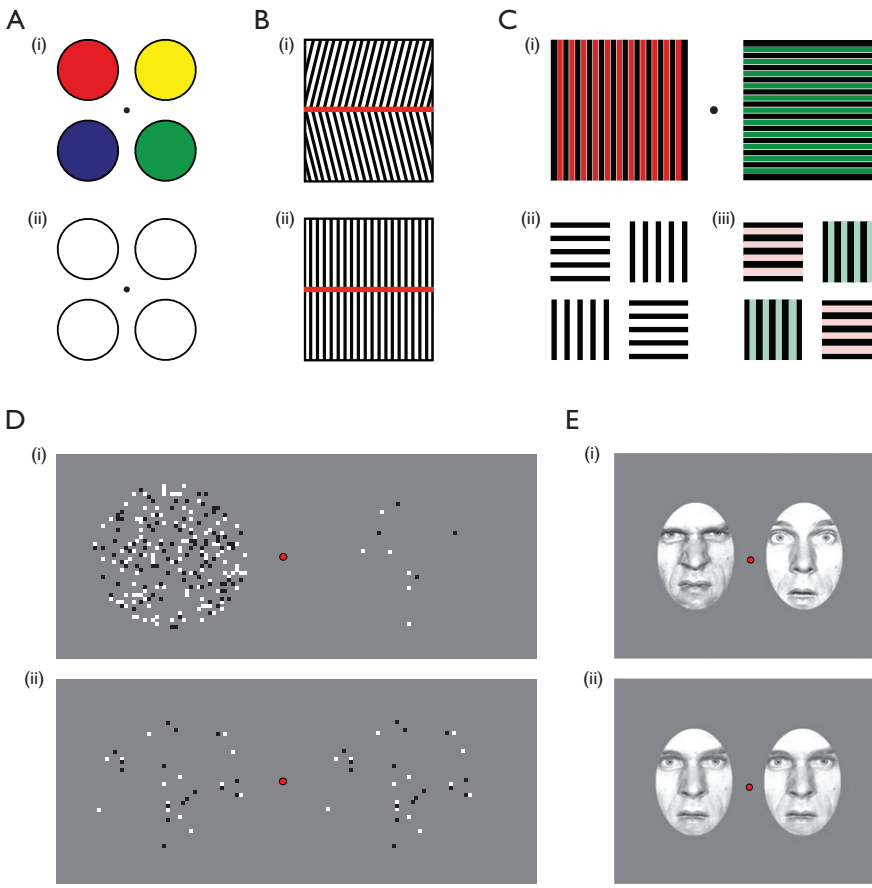

Figure 1 From Phillips and Firestone.

for 30–60 seconds, and then looking at the releaser stimulus in the bottom part of the image, what you experience is that the left cloud of dots on the bottom briefly looks to have fewer dots than the right cloud of dots. Then, after a fraction of a second, they look to have the same number of dots. The explanation is that we have spatiotopic visual channels for high numerosity and low numerosity. Focusing on the top of D decreases the sensitivity for high numerosity on the left and decreases the sensitivity for low numerosity on the right, yielding the repulsive effect when looking at the bottom of D. As I mentioned, students in my classes who view these stimuli often suppose that I have surreptitiously changed the stimulus.

The phenomenology changes as a result of adaptation even though the stimulus stays the same. But why is this change in phenomenology any more significant than the static phenomenology prior to the change?

We can rule out the most significant low-level explanation of the change due to adaptation: it is clear from introspection that dots do not seem to appear or disappear despite the changes in apparent numerosity. After staring at the top part of D and then looking at the bottom part, it just looks from the get-go that there are more dots on the right and fewer on the left. And then it seems the numerosities are equal on the two sides, without any appearance or disappearance of dots (see for yourself!). The one potential low-level confound is whether the operative change might be in texture density rather than numerosity. This is an empirical issue that has been explored with great ingenuity by creating stimuli in which texture density and numerosity are varied independently. One particularly impressive experiment, summarized in Block 2019, added lines connecting dots, increasing the texture density while decreasing the apparent numerosity. The adaptation effects depended on the numerosity, not the density.

The phenomenology of adaptation is part of the package of evidence that shows that numerosity perception is at least in part a high-level perceptual phenomenon. But the fact that it is the phenomenology of adaptation that I am appealing to does not show adaptation is playing no real role, since static phenomenonology – without adaptation – would not suffice for establishing the result.

A top-down source of the change in perceptual judgement in D is very unlikely since one has no tendency to *believe* that the paper stimulus actually changes. So the change in perceptual judgement is plausibly due to the change in high-level perception. And that suggests that we don't have to appeal to any supposed phenomenology of high-level perceptual judgement to explain the effect. And the reasoning just discussed depends essentially on adaptation, so adaptation is playing a real role.

In sum, in distinguishing low- from high-level perception and distinguishing both from perceptual judgement, adaptation plays a real role.

I haven't said anything yet, though, about whether adaptation plays an independent role in ruling out 'criterion' effects. Decision criteria influence how much perceptual information the subject needs to respond to in a

certain way. Consider a 'detection' task in which the subject is supposed to press one button if there is a stimulus and another if there is no stimulus. If there are very many no-stimulus trials, subjects will tend towards a 'conservative' criterion in which they must be quite sure there is a stimulus to push the yes button. And if you reward the subject for detecting the stimulus without penalizing them for 'false alarms' (incorrect yes button presses), the subjects will respond by liberalizing detection, pressing the yes button without being very sure there is a stimulus.

Can a 'criterion' effect explain away the evidence for numerosity perception? No, because the changes in phenomenology of first seeing the bottom part of D as having more dots on the right and then equal numbers on both sides is clearly a change in perceptual phenomenology and not at all like the phenomenology of a changing criterion of the sort described in the last paragraph. Note that although the appeal here is to phenomenology, it is the phenomenology of adaptation that plays a real role.

1.1.1 *Retinotopy and spatiotopy* Moving to retinotopy and spatiotopy, retinotopic effects depend on where light falls on the retina, and they typically move with the eye. Spatiotopic effects preserve retinal neighbourhood relations but do not move with the eye. The adaptation effects in D and E are spatiotopic because they depend on whether the adaptor and the releaser are on the left or right side of space. The left cloud of dots in D(ii) looks to be fewer than they are after the adaptation because the left cloud of dots in D(i) was numerous.

So why do we need adaptation at all in reasoning about numerosity perception, given that we know which dots are on which side of space and which dots fall on which parts of the retina? The answer is that these facts tell us nothing about numerosity perception, whether there is numerosity perception at all or alternatively, whether vision just attributes low-level properties and we make judgements about numerosity. The reasoning I went through a few paragraphs ago that distinguishes low-level from high-level perception and both from perceptual judgement all depended on adaptation.

2. *Signal detection theory*

I will now move to a disagreement between me and Phillips and Firestone on the proper role of signal detection theory. Signal detection theory is pervasive in their reply though it is only discussed explicitly near the end. Their discussion makes heavy use of the signal detection theory notion of a 'criterion'. I'll start with the issue they raise at the end of their article since it reflects in a simple and straightforward manner why I think the signal detection theory framework is problematic in application to the issues of my book. The issue is my thesis that perception is a winner-takes-all faculty. I argue that there is always a dominant perception and that a likely explanation is that the

purpose of perception encourages the design of a perceptual system that does not wallow in ambiguity.

As I note in Chapter 4, the dominant percept can take a number of forms. In one, the rival percepts merge, and in another, the rival percepts compete. In binocular rivalry, there is an interplay of local and global features in determining winners. Two faces – one projected to each eye – can be sufficiently 'compatible' in a sense described in Chapter 4 for representations of face stimuli to merge when the subject is attending holistically to the face but to compete when the subject is attending to local features. In particular, a masculine and feminine face (one projected to each eye, as shown in Chapter 4) results in a percept of an androgynous face when the subject attends to the face as a whole. But when the subject attends to parts, the two faces alternate. See Figure 2, reproduced from Chapter 4.

As Phillips and Firestone note, it is not easy to see how a winner-takes-all nature of perception is compatible with a signal detection framework in which there is a continuous balance between distributions representing different percepts. This shows an inadequacy in signal detection theory: it concerns the information content of perception, not the percept itself. Let me explain.

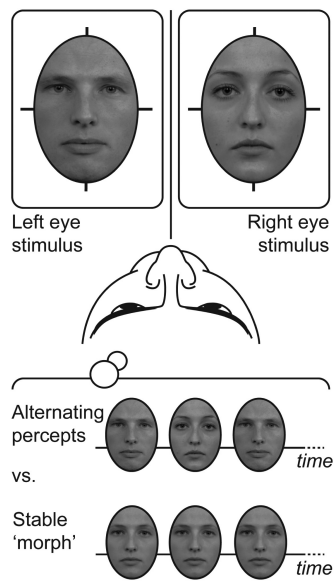

Figure 2 Binocular rivalry stimuli. The masculine face is presented to the left eye and the feminine face to the right eye. If the subject is attending to local features or parts such as the eyes, standard binocular rivalry ensues. This is indicated by the label 'alternating percepts'. If the subject is attending to holistic features such as gender or the identity of the person, the subject sees a morphed androgynous face, as pictured. Thanks to Chris Klink for this figure. See Klink et al. 2017.

As Phillips and Firestone note, in a stimulus situation in which stimuli are degraded, there is substantial information in subjects' second and even third 'guesses'. But the second and subsequent choices reflect the information in a non-dominant percept that is not conscious, or as I will say here with misgivings that I cannot go into, not present to the person.

It is true that the information reflected in the second and subsequent choices are represented in the visual system. When a subject has a percept as of an androgynous (merged) face, the information from the masculine and feminine faces is present in an implicit form. When the subject attends to the parts of the face, and when the feminine face is dominant, there is no hint of the masculine face at the personal level (although there can be patchy versions of both that briefly appear in the transition). The dominant face is the only one present in the person-level percept, although the information about the suppressed face is still represented in the perceptual system. Signal detection theory concerns the information represented in the visual system, but without emendation, it does not concern the percept itself. It is a defect in signal detection theory that it has no notion of what is present at the personal level, and this failure stems from the lack of thresholds in signal detection frameworks.

Matthias Michel and Hakwan Lau suggest, following Michel and Lau 2021, Witt et al. 2015, a modification of signal detection theory in which thresholds are introduced as 'perceptual criteria', as distinct from the report or decision criteria usually understood as part of signal detection theory. Decision criteria influence how much perceptual information the subject needs to respond to in a certain way. To use a variant of the previous example, if in a detection task you reward the subject for correct no-stimulus judgements without penalizing them for incorrect no-stimulus judgements, the effect will be to shift the subject temporarily towards a 'conservative' bias in deciding on and reporting a stimulus, that is, the subject will only say the stimulus was there when they are pretty sure it was.

The idea of a perceptual criterion is that signals whose strength is less than the perceptual criterion are unconsciously perceived.[1] Second choice accuracy can be accounted for by such a model. A low-threshold model in which the noise distribution is sometimes over the threshold would also explain confident false alarms since confidence judgements could be based in part on the noise distribution (I am indebted to Matthias Michel and Wixted 2020, but see Phillips 2021).

Note that the dominant perception issue is not just about phenomenology and consciousness: it is about the dominant percept, and that is a wider

1 Perhaps there should be another 'perceptual criterion' above which is unconscious perception and below which is no perception at all. A criterion below which there is no perception at all is required to account for binocular rivalry in unconscious perception.

notion than just the notion of a conscious percept. Adaptation effects can occur in unconscious perception of tilt, as noted in Chapter 2. As I also note in Chapter 2, binocular rivalry can occur in unconscious perception. And fruit flies show binocular rivalry though they may not be conscious creatures.

In the book, I made a suggestion concerning why we have dominant percepts – that since the perceiving subject often has to act quickly, 'it won't do for perception to wallow in ambiguity'. As Stephen Jay Gould noted, evolutionary accounts are often 'just so stories' (Gottlieb 2012). They may sound plausible, but other stories that conflict with them can also sound plausible. While mentioning Gould, Phillips and Firestone counter with another evolutionary suggestion:

> But in any event, since in the case of non-reflex actions there must always be a decision to act or not, it is unclear why ambiguity precludes quick action. If one sets a criterion which says: 'Run if there is even a 1% chance of a tiger being present', one does not need a winner-takes-all tiger percept to run quickly.

Of course, complication in perception would presumably come at a cost. For whatever reason, the visual system apparently does not follow the advice of Phillips and Firestone and does not use signal detection parameters in determining perceptual content, even though they play a role in confidence judgements and in guessing. One might speculate (of course, just another just so story) that adding probabilities to perceptual contents would excessively complicate the visual system (there has been a debate about whether perception is indeed probabilistic. See Morrison 2016, Munton 2016, Denison 2017, Block 2018).

2.1 'The' purpose of perception

Turning now to the issue of the title of Phillips' and Firestone's critique, I don't think there is much of a disagreement.

As I mentioned in Chapter 2, perceptual adaptation operates via many different mechanisms. It is a network phenomenon and not merely a cellular phenomenon. It encompasses quite different kinds of perception, notably both opponent process perception and multi-channel. It operates in both high- and low-level perception. In its multiple mechanisms, it reminds me of depth perception. Multiple mechanisms do suggest evolutionary selection but exactly what the relevant purposes are is unclear.

As they note, I say, 'The evolutionary purpose of perception is acquiring information about what is happening here and now. Call that "news"'. I'm slightly embarrassed by my use of the word 'the' here, since everyone knows that postulating evolutionary purposes is very often a matter of 'just so' stories, and postulating single purposes seems especially speculative.

If the evolutionary pressure is general enough, perhaps it can be justified, and indeed the purpose of taking in information about the here and

now does seem general enough that, as they say, it 'will be hard to disagree' with. But that sets me up with a problem which they express in the form of a dilemma:

> In other words, Block faces a dilemma: either understand 'news' broadly as 'information about what is happening here and now' and give up on the idea that this purpose explains perception as characterized by adaptation; or understand 'news' more narrowly as referring to changes along a specific feature dimension, but then lose the plausibility of the claim that detecting news is a general purpose of perception.

Now we all learned in Dilemma 101 that the horns of a dilemma have to be incompatible – or at least we have to be forced to choose between them. One horn is the one just discussed, acquiring information about the here and now, and we seem to agree that that is a very general purpose of perception. Information about specific changes in the environment is a special case, so perception has both a general and a number of specific purposes that are special cases of the general purpose.

Still, Phillips and Firestone are right that it is not obvious how either the general or specific purposes just mentioned really show why adaptation is so important in perception.

3. Word representations and ampliative transitions: reply to Steven Gross

In a wonderfully insightful article, Steven Gross points to a serious problem with the views in my book. Luckily for me, he suggests a number of ways that I can modify my views while keeping the main lines of them intact.

3.1 Abstract word representations

The problem Gross raises involves my thesis that perceptual representations are constitutively iconic in format. What is iconic format? I discuss a number of somewhat different characterizations of iconicity used in cognitive science, but the one that I take to be most important is a 'mirroring' characterization based on Roger Shepard's notion of second-order (Shepard 1978, Shepard and Chipman 1970). The main idea of the 'mirroring' characterization is that relations among perceptual representations are analogues of relations among the properties in the world that are perceptually represented. But there are other notions of iconicity used in cognitive science, and Gross's critique does not depend on any one particular kind or notion of iconicity.

The problem raised by Gross is that I hold that one of the constitutive features of perception is iconic format, while at the same time leaning on adaptation as a strong (though imperfect) indicator of perception. He notes a difficulty with that combination of views based on an experiment by Hanif et al. (2013).

The way the experiment works is diagrammed in the figure in Gross's paper and with different details in Figure 3. An adaptor, in this case an upper or lower case word, is presented for 5 seconds. The word pairs 'area/name' and 'crane/nerve' were used because the words are matched for familiarity, imageability and concreteness, and it is possible to make a stimulus that is ambiguous between them both in upper and lower case. What is illustrated in Figure 3 is the version in which the adaptor is either 'area' or 'AREA'. A mask is then presented briefly. Experiments like this often use masks to prevent the perceptual representation from continuing after the stimulus is ended.

Then, after 300 ms, an ambiguous stimulus is presented (this time with no mask). In the case illustrated in the figure, the ambiguous stimulus is ambiguous between an upper case 'AREA' and upper case 'NAME'. Then the subject has a choice between lower case 'area' and 'name'. The 'congruent' version is when the ambiguous stimulus is the same case as the adaptor; the 'incongruent' version is when the ambiguous stimulus is in the opposite case from the adaptor. The choice is always between items of the opposite case and the adaptor (so what is pictured assumes the adaptor is 'AREA' and that the case is congruent). The signature of adaptation is that the subject will be biased against seeing the ambiguous word as the word initially presented, that is, the adaptor. And that is what happens. Seeing 'AREA' for 5 seconds makes one much less likely to choose 'area', that is, much more likely to choose 'name'. This is the classic 'repulsive' effect of adaptation.

Hanif et al. showed a 'cross-case' adaptation effect for visually presented words in which an adaptor word had a 'repulsive' effect on seeing another word, irrespective of whether the adaptor words were written in a different case from the adapting words (they showed similar results for varying font and handwriting styles, but I will discuss only the case result). The problem for me is how there can be an iconic representation of a word that abstracts from case, given that there doesn't seem to be an iconic format shared between the lower case and upper case words.

The key finding from my point of view is that the effect works in the incongruent version. That is, when the adaptor is a lower case 'area', that makes one more likely to respond to the upper case ambiguous stimulus as 'NAME', and when the adaptor is an upper case 'AREA', that makes one more likely to respond to the lower case ambiguous stimulus as 'name'.

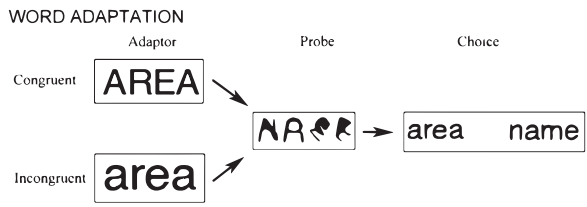

Figure 3. From Hanif et al.: 69. Thanks to Jason Barton for this figure.

There is some indication in the data that the congruent effect is bigger than the incongruent effect. See Figure 4.

The difference shown in Figure 4 is large but not significant. Philosophers may wonder how that can be? If I measure the height of the students in my small seminar and find that the students whose name begins with 'A' are on average two feet taller than the students whose name begins with 'B', we have a large difference that can be explained by a chance effect in a small sample size. The case experiment had only 14 subjects and not many trials. However, Hanif et al. also showed in a different experiment from the one just reported that there is an adaptational effect of case, so it may well be that the effect of Figure 4 is real.

However, the problem for me is that there is any effect at all in the incongruent version, since that suggests a perceptual word representation that is neutral between upper and lower case, and that does seem in tension with an iconic theory of perceptual representation.

I agree that there is a level of word form representation that is more abstract than either the lower case or the upper case word representation. The problem is how that representation can be iconic. Focus on the word 'AREA' or the word 'area' in Figure 3. If the word that you focused on were slowly rotated, the mirroring 'rotation' of the perceptual representation would depend on whether what is rotated is 'AREA' or 'area'. One indicator of iconicity is whether two tokens of the same iconic type can be superimposed, preserving format, and upper and lower case formats do not pass this test.

It is possible to make a picture of a word that is ambiguous between lower and upper case – and the experimenters did that for both upper and lower case 'area', 'name', 'crane' and 'nerve'. But just eyeballing the one example

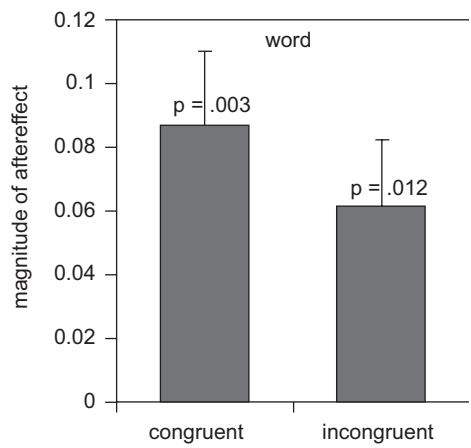

Figure 4. Slightly modified from Hanif et al.: 64, Figure 3. Thanks to Jason Barton for this figure and permission to modify it slightly.

they give in Figure 3, the techniques are to alternate upper and lower case letters and use ink blobs, and that technique does not seem to lend itself very well either to the mirroring or superimposition criteria.

Gross offers a number of possible ways out. One way, he suggests, is to treat this as an 'edge-case' that no more challenges the joint than twilight challenges the distinction between day and night. I discuss a number of 'borderline' examples similar to the twilight example. For example, I note that the existence of glasses that are hard as with solids but have the amorphous structure of liquids does not impugn the explanatory significance of the division of matter into liquid, solid and gas. There is more than one way to make a liquid into a solid. In the case of water to ice, the way involves changing an amorphous molecular structure into a crystalline lattice. That is explanatorily significant, even if super-cooled liquids exemplify a more gradual way of solidifying. A case of perception that was discursive, however, would threaten the explanatory unity of iconicity, non-conceptuality and non-propositionality.

One promising option suggested by Gross is that there may be associative links between upper and lower case representations so that activation of one activates the other. The thought here is that 'area' has an adaptive effect on 'AREA' because 'area' triggers an instantiation of 'AREA', and that has an adaptive effect on 'AREA'. This theory would also explain why the incongruent effect is smaller than the congruent effect – because the incongruent effect is mediated by an association and so is not direct.

As Gross notes, there might also be an associational effect mediated by overlearned associations between orthography and phonology. Look at Figure 3. You see (say) 'area' for 5 seconds. We know from many experimental results (e.g. the Stroop effect) that for most adults, reading is automatic. Anyone looking at 'area' for 5 seconds is going to read the word, activating the representation of 'area' at all levels, including at the phonological level. Then one sees the ambiguous stimulus for 300 ms. Since there is no mask after the ambiguous stimulus, one might continue to process it for another 150 ms before the choice between the two words is given. It is well known that both meanings of ambiguous words are automatically processed in parallel (Lucas, 1987), and something analogous might apply to both readings of the morph. Then, the story goes (as applied to when seeing 'AREA' biases one against 'area' and towards 'name'): seeing 'AREA' activates the sound of the word, biasing one against the sound of 'area', thereby negatively priming the orthographic 'area'. If we write the sound of 'AREA' and 'area' as /area/, we could put it like this: seeing 'AREA' activates /area/, seeing the ambiguous stimulus activates both /area/ and /name/, but the activation of the first /area/ suppresses the second, leading to favouring /name/.

3.2 Ampliative transitions

Gross considers the question of whether the transition between perception and basic perceptual belief is 'ampliative' in the sense that content is changed.

As Gross notes, I suggest precisifying Burge's notion of a basic perceptual belief as a minimal immediate direct perceptual judgement. I won't repeat the definitions of these terms except for 'minimal', which I defined in terms of conceptualizing each representational aspect of a perception and no more. On my account, perceptual beliefs are conceptualized versions of perception. I say that conceptualization is the product of global broadcasting in which long-range axons connect perceptual areas with prefrontal cortex, allowing for an active neural coalition that yields inferential promiscuity (to use Stich's term). It is that inferential promiscuity that makes the representation a concept.

I argued that the process of conceptualization changes the format of representation at least somewhat and also encloses the perceptual representation in a discursive envelope, so, contrary to Burge, the formation of basic perceptual belief is ampliative.

One of the kinds of changes that I mention is that basic perceptual belief is coarser grained than perception. Gross mentions also that there are cognitive 'biases' that, for example, favour the categorical centre. These are just two of the ways in which the transition is ampliative. As I mentioned in Chapter 5, perceptual beliefs depend on the task. Yuna Kwak and Clay Curtis (2022) used two kinds of stimuli on different trials, oriented gratings (Gabor patches) and clouds of moving dots. The subjects' task was to indicate the orientation of the grating or the direction of the moving dots after a delay period. They scanned the subjects using fMRI during the delay period prior to doing the tasks. One result was that decoding trained on the grating task also worked on the dot task and vice versa. This fact shows that the working memory representation was sufficiently abstract as to be common between the two perceptions. This working memory result suggests that two different perceptions may yield the same basic perceptual belief, and if so, that transition would certainly be ampliative.[2]

Departments of Philosophy and Psychology, Center for Neural Science
New York University
USA
ned.block@nyu.edu

References

Begby, E. 2011. Review of origins of objectivity. *Notre Dame Philosophical Reviews*. <http://ndpr.nd.edu/news/24627-origins-of-objectivity>

Block, N. 2014. Seeing-as in the light of vision science. *Philosophy and Phenomenological Research* 89: 560–72.

Block, N. 2018. If perception is probabilistic, why does it not seem probabilistic? *Philosophical Transactions of the Royal Society B: Biological Sciences* 373: 20170341.

2 Thanks to Susan Carey, Steven Gross, Matthias Michel and to Ian Phillips and Chaz Firestone for comments on an earlier version.

Block, N. 2019. Tyler Burge on perceptual adaptation. In *Blockheads! Essays on Ned Block's Philosophy of Mind and Consciousness*, eds. A. Pautz and D. Stoljar, 71–78. Cambridge, MA: MIT Press.
Denison, R. N. 2017. Precision, not confidence, describes the uncertainty of perceptual experience: comment on John Morrison's "perceptual confidence". *Analytic Philosophy* 58: 58–70.
Gottlieb, A. 2012. It ain't necessarily so. *New Yorker*.
Hanif, H.M., B.L. Perler and J.S. Barton. 2013. The visual representations of words and style in text: an adaptation study. *Brain Research* 1518: 61–70.
Klink, P.C., D. Boucherie, D. Denys, P.R. Roelfsema and M.W. Self. 2017. Interocularly merged face percepts eliminate binocular rivalry. *Scientific Reports* 7: 7585.
Kwak, Y. and C.E. Curtis. 2022. Unveiling the abstract format of mnemonic representations. *Neuron*. 110: 1822–28.e5.
Lucas, M.M. 1987. Frequency effects on the processing of ambiguous words in sentence contexts. *Language and Speech* 30: 25–46.
Michel, M. and H. Lau. 2021. Is blindsight possible under signal detection theory? Comment on Phillips (2021). *Psychological Review* 128: 585–91.
Morrison, J. 2016. Perceptual confidence. *Analytic Philosophy* 57: 15–48.
Munton, J. 2016. Visual confidences and direct perceptual justification. *Philosophical Topics* 44: 301–26.
Phillips, I. 2021. Bias and blindsight: a reply to Michel and Lau (2021). *Psychological Review* 128: 592–5.
Shepard, R. 1978. The mental image. *American Psychologist* 33: 125–37.
Shepard, R. and S. Chipman. 1970. Second-order isomorphism of internal representations: shapes of states. *Cognitive Psychology* 1: 1–17.
Siegel, S. and A. Byrne. 2016. Rich or thin? In *Current Controversies in Philosophy of Perception*, ed. B. Nanay, 59–80. New York and London: Routledge.
Witt, J.K., J.E.T. Taylor, M. Sugovic and J.T. Wixted. 2015. Signal detection measures cannot distinguish perceptual biases from response biases. *Perception* 44: 289–300.
Wixted, J.T. 2020. The forgotten history of signal detection theory. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 46: 201–33.