

Chapter 19

No Problem

David Hilbert and Colin Klein

Pautz refutes tracking intentionalism. We defend it. In what follows we will attempt to undermine the motivation for some of what he says and discuss one of his arguments in more detail. We don't, however, defend tracking intentionalism because we believe it to be true. There are parts of it we do find plausible (intentionalism with external world content). We defend it, though, because we believe some of Pautz's criticisms raise interesting questions about how to think about perception that are independent of the truth or falsity of tracking intentionalism.

19.1 Perception Is Not Magic

One of the themes of the paper is that there is good evidence for internal states that are well correlated with sensory phenomenology and that, at least for the chemical senses, there are no equally well correlated external states.¹ According to Pautz, this

¹Although we will only mention one example here, there are a number of difficulties in interpreting the empirical evidence Pautz brings forward concerning what external properties the internal states he discusses might be tracking. Pautz's discussion of olfaction relies heavily on the very interesting work done on posterior piriform cortex (PPC) by Gottfried and collaborators (Howard et al. 2009; Gottfried 2010; Zelano et al. 2011). The central conclusion of this work is that PPC contains a distributed representation of *odor objects* (Stevenson and Wilson 2007). Odor objects are learned patterns of more basic odors that correspond to olfactory complexes like the smell of chicken. PPC is thus supposed to function in a way resembling visual object and face recognition. Complaining that there is no simple chemical correlate to PPC activity is thus like complaining that there is no simple physical magnitude corresponding to the activity of visual face cells or to activity in

D. Hilbert (✉) • C. Klein

Department of Philosophy, University of Illinois at Chicago, 1420 University Hall,
MC 267 601 S Morgan Street, Chicago, IL 60607, USA

e-mail: hilbert@uic.edu; cvklein@uic.edu

fact raises a *prima facie* problem for tracking intentionalism. A typical version of this claim can be found in the conclusion to Sect. 18.2:

The fact that when it comes to phenomenal character there is “bad external correlation” but “good internal correlation” across the various modalities makes one suspect that there is something very wrong with the radically externalist approach promoted by tracking intentionalists, according to which phenomenal character is fully determined by the external physical properties tracked by our experiences . . . and which accords no serious role to internal factors. (p. 12)

Pautz musters detailed empirical evidence in support of this conclusion but the details are unnecessary here. That sensory phenomenology is better correlated with physiology than any external property is a *consequence* of tracking intentionalism in conjunction with very general (and relatively uncontroversial) empirical considerations.

At the heart of tracking intentionalism are two claims. First, that sensory phenomenology is wholly explained by the content of sensory states (intentionalism). Second, that the content of sensory states is to be explained in terms of how they are connected to the external world (tracking). Exactly which internal states track which external states will depend on both the structure of the world and also the structure of the sensory system (including the brain) of the organism. There won't be internal states that track features like acidity without the presence of sensors that respond to the pH of substances in the mouth and without those sensors being connected to neural circuits that process and deliver the information obtained from the sensors. The causal relationships and correlations that underlie tracking depend crucially on internal features of organisms. That the relationship between internal states and phenomenal experience is systematic and not random is also to be expected. Given very general assumptions about physiology and the evolution of nervous systems what is to be expected is that the internal states that track environmental features will have some systematic structure. If we assume that these states are related to perceptual experience and behavior in systematic ways (an unsurprising feature of actual neurophysiology), then good internal correlation falls out directly. Good internal correlation is not only consistent with tracking intentionalism but to be expected.

19.2 Feeling Curved

Good internal correlation thus can't be a threat to tracking intentionalism. And indeed, Pautz's arguments rely heavily on the notion of a “bad external correlation.” Now, some degree of independence between the property being tracked and the

the ventral stream visual areas involved in object recognition. There are interesting questions for tracking intentionalism here but they are more complicated than the issue of whether there is a simple physical or chemical feature that is correlated with activity in PPC.

internal state doing the tracking is an important part of tracking intentionalism. The particular versions of tracking intentionalism Pautz discusses are very concerned to allow for the possibility of misrepresentation and build their theories accordingly. For example, Tye uses causation under optimal circumstances, rather than plain causation, and Dretske builds representation on indicator functions, rather than plain indication. The inclusion of an optimality requirement and the appeal to function make it possible for sensory states to misrepresent the external world. Since misrepresentation is possible (and actual according to both) external correlation is less than perfect.

Pautz, however, appeals to a more serious mismatch with the external world. In the case of thermal pain, for example, there is an exponential function relating stimulus intensity and judgments of pain intensity. This gives rise to response expansion: a doubling of stimulus intensity gives rise to a more-than-doubling of judged intensity. So according to Pautz, there is a bad correlation between judged intensity and the external stimulus. There is a “perfect correlation” on the other hand, between judgments of pain intensity and internal qualities like S1 firing rates.² Problems for tracking intentionalism should follow. We have noted that good internal correlation is compatible with tracking intentionalism. Are “bad external correlations,” in Pautz’s sense, a threat? You might think so: you might think that bad correlation means poor tracking. This would be a mistake.

First, note that Pautz’s notion of “bad external correlation” has nothing to do with the ordinary scientific use of a “bad correlation.” In ordinary usage, two variables are correlated if and only if there is an association between them such that information about one *reliably* carries information about the other. A perfect correlation means that the value of one quantity is completely informative about the value of the other. Since ‘reliably’ is a graded notion, correlation is a graded notion as well. At the lower ends, however, a poor correlation means that the two quantities don’t have much to do with each other: knowing the value of one doesn’t carry any information about the value of the other.

If internal states and external stimuli were poorly correlated in *this* sense, tracking intentionalism would obviously be in trouble. But that can’t be the claim. For all that’s been said, a subject’s pain intensity judgment lets you predict, perfectly

²See p. 10. We grant this latter claim for the sake of argument, but note that it is problematic in a number of respects. Perfect correlations are mathematically improbable in neuroimaging work, even if there is actually a perfect relationship between a variable and neural response (Vul et al. 2009, p. 275). Pautz cites Coghill et al., but they claim only that the correlations are statistically significant, not that they are perfect (1999). S1 is also a problematic place to locate intensity information. Although there are regions of S1 which have activity that is well-correlated with pain intensity judgments (particularly in BA 3a), these regions are not obviously the substrate of pain experience. Large lesions of S1 do not reliably eliminate pain sensation, and stimulation of S1 does not reliably produce pain sensation (Craig 2003, pp. 18–19). The worry is not just the one that Pautz notes, that pain sensation might be more widely distributed. Rather, it is that a linear correlation between intensity ratings and neural responses need not indicate a necessary part of the substrate of pain experience.

accurately, the intensity of the stimulus to which the subject was subjected. So there is a good correlation between intensity judgments and external states, at least in this straightforward sense.

Pautz's claim, so far as we can tell, is something much weaker: that there is no *linear* relationship between judgments and stimuli. Certainly so. But why should the tracking intentionalist care? So long as the judgments reliably carry information about the external state – and on this account, they do – the tracking intentionalists have all they wanted. Linearity is just one of many possible informative relations that can hold between quantities. Any of those are candidates for representation relations; some, for good engineering reasons, might be preferred to others. So where's the problem?

The tracking intentionalist might rest content here. To his credit, however, Pautz gives an argument for why we should prefer linear relationships (though it is not obviously couched as such). Pautz envisions two fluid-filled columns which track lengths, one with a relationship of n^2 to the length and the other with a relation n^3 .

This provides a schematic illustration of my internal-dependent argument and my structure argument. How could the “tracking intentionalist” or “objectivist” about sensible qualities accommodate the verdict that A and B have sensations of different intensities, and how might he accommodate the truth of their structure judgments? After all, they track the very same objective lengths; and the objective length does not really *double* (or more than double) when it goes from 2 cm to 3 cm (p. 37).

Here's a way to reconstruct this argument: there are indefinitely many information-carrying relationships that might hold between a quality and its internal representation. Different relations will give rise to different judgments about the external quality over some range. Yet each of these representations putatively track the *same* property. As they disagree, they can't all be *accurate*. Further, the linear relation most neatly mirrors the world, and so has the best claim to be the accurate one.

Pautz considers one possible response, which is that the different relations track different properties in the world; he ultimately dismisses this as “Pythagoreanism”. We'll leave the defense of Pythagoreanism to those who find it attractive. Instead, we suggest that there is another, perfectly reasonable, response available to the tracking intentionalist. Different relations track the same property, but provide *different information* about that property. Properly cashed out, this blunts the force of Pautz's argument.

The brain contains regions with a binary response to stimuli.³ So consider two possible neural response functions for thermal pain: a continuous one that increases linearly with stimulus intensity, the other a binary function that changes its output over a certain threshold – say 46 °C. Both of these clearly track the same property of the world: the degree of the thermal stimulus. If we imagine these as instantiated in two different organisms, it's also obvious that they would have

³In the case of painful thermal stimuli, see Bornhöyd et al. (2002).

different experiences: the continuous neural response would give rise to graduated pain sensations that can't be represented by the binary function. So it is possible to have two representations, each of which track the same property, and yet which give rise to different sensations.

Yet surely, there is no *mystery* here about how that can be the case: the difference in sensation is due to the fact that the two functions carry different information about the world. The amount of information carried by a representation is a function of how many potential states the representational vehicle might be in. A binary representation which can only be in two states can only carry one bit of information: in our example, whether or not the stimulus is painful. The continuous linear stimulus can carry more information: not just whether or not the stimulus is painful, but the intensity of the stimulus. But these are just differences in the amount of information conveyed by the representation, not what it is information *about*: both representations track the same feature of the world.⁴

Indeed, the same point can be made with two *linear* functions. Actual neural response functions are not continuous: there are some discriminations that are too fine for neurons to make. So consider two hypothetical entities *Coarse* and *Fine*. Unlike us, both have linear response functions for thermal pain: a doubling of the stimulus exactly doubles the judged intensity of pain. So there is “good external correlation” in Pautz’s sense. Yet *Coarse*’s neural mechanisms can only discriminate with an accuracy of 1 °C, while *Fine* can discriminate with an accuracy of 0.1 °C. So *Coarse* will lump together as similar many states that *Fine* will distinguish. Intuitively, *Coarse* and *Fine* will also have different pain experiences, and their different pain experiences will also be well-correlated with the state of their putative neural mechanisms.

Now, we have a curious case. Both *Coarse* and *Fine*’s sensations should have both good internal and external correlation in Pautz’s sense. Yet the same stimulus can give rise to different sensations in each. Further, if you’re a tracking intentionalist, you have a perfectly good story about how this works: both *Coarse* and *Fine* track thermal stimuli, but the mechanisms by which they track it carry different information. That difference in information is a difference in representation, which gives rise to a difference in phenomenology – exactly as the tracking intentionalist predicts.

What this shows, we think, is that “good correlation” in Pautz’s sense is a red herring. What matters for tracking intentionalism is ‘correlation’ in the old fashioned sense: that is, in the sense of carrying information. There are many ways of carrying information, however, and each gives rise to a different way of tracking the world. Different ways of tracking the world may lump stimuli together as more or less similar: which of these lumpings is preferable is not an *a priori* matter. It is an engineering one: a well-designed system should treat as similar states which need to be treated similarly in output.

⁴In the philosophy literature these points were given prominence by Dretske (1981, 1995).

If this is right, then it is an easy step to dispense with Pautz's case of *Mild* and *Severe*. *Mild* and *Severe*, remember, have response functions each with a different steepness. Intuitively, there should be a difference in their response for the very same stimulus. We agree: they are tracking the very same stimulus, but carrying different information about it. Because of this, they respond differently to the same stimulus. Given the granularity of any neural response function, the two functions will carry different information about the very same facet of the world. So it is unsurprising that their phenomenal experience should also differ. But this is a fact that can be fully accounted for by the tracking intentionalist.

19.3 Tracking Systematically

Tracking intentionalism is not threatened either by good internal correlation nor poor external correlation (in Pautz's sense). We believe that Pautz, although he often puts his point in terms of poor external correlation, is actually making an additional, distinct argument. One of the morals that Pautz wishes to extract from his survey of sensory physiology and psychophysics is that, for many senses, there is no external property to be tracked at all. It's not just that the correlation between the property supposedly being tracked and the internal state supposedly doing the tracking is imperfect; it is rather that there are no plausible candidates at all for the properties being tracked. This argument does not apply to all modalities (thermal pain, for example, surely tracks some straightforward external property), and is most plausible in the case of the chemical senses. The form of the argument is familiar from the color literature but Pautz very usefully extends this argument form to a much broader array of sensory properties. We conclude by considering some of these cases.

Pautz offers four examples of cases in which there is supposedly sameness in content combined with difference in experience. These are then generalized to form the internal-dependence argument which, in effect, asserts that some such example is (probably) possible (p. 277). Each case follows the same pattern. There are two individuals who are stipulated to have differing activity in a neural area correlated with perceptual phenomenology while tracking the very same property. Pautz then argues that because of the difference in neural activity it's plausible to suppose that the two are undergoing phenomenally different perceptual experiences. These cases are thus putative counterexamples to tracking intentionalism.

This form of argument immediately runs into a problem. For concreteness take the case of *Yuck* and *Yum*. *Yuck* and *Yum* have very different neural responses to a specific variety of berry. If this is all we are told about the case then we don't really have an argument, just a case that might elicit conflicting intuitions. If *Yuck* and *Yum* were very similar in their neural circuits that process taste information then the empirical evidence Pautz brings forward might support the claim that differing neural activity implies phenomenally different experiences. But one thing we do

know about *Yuck* and *Yum* is that they differ in the neural circuits that process taste information since they have significantly different neural activity in response to the same stimulus.

Pautz adds detail to his cases which serves to make plausible the claim that *Yuck* and *Yum* are having phenomenally different experiences. *Yuck* finds the berries similar in taste to poison dart frogs (he shows an amazing willingness to stick things in his mouth) while *Yum* finds the berries similar in taste to bananas and there are corresponding similarities in neural activity for each of them. Their behavioral responses are also very different. Nevertheless when they are tasting the berries they are both representing them to have the very same property. Thus, according to Pautz's intentionalist, there is property that *Yuck* represents and that he represents in a way that is similar to the way he represents a property possessed by poison frogs. That very same property is also represented by *Yum* and his representation of it is similar to his representation of a property of bananas. The added detail, although crucial to establishing that *Yuck* and *Yum* are experientially different, is in some tension with the stipulation that *Yuck* and *Yum* are in states with the same intentional content.

Any version of tracking intentionalism that is able to offer a substantive account of perceptual similarity will be able to offer an account of these types of cases. If berries and poison frogs have similar tastes for *Yuck*, then there must be some similarity in the taste properties they are represented as possessing. It's not enough for similarity that *Yuck* represents the one as having taste A and the other as having taste B. Without some structure to the representations of taste, there will be no basis for judging the tastes to be more or less similar.

So, *Yuck* could represent the tastes by representing (some aspect of) the chemical structure of the objects in which case similarity in taste would track similarity in those aspects of chemical structure. Or *Yuck* could be representing the taste (at least partly) in terms of the effects of the substance on his digestion, in which case similarity in taste would track digestive effects. Given that *Yum* judges different tastes as similar, the one thing we can be sure of is that his representation of taste tracks different aspects of the world from the ones tracked by *Yuck*.

It's only if we think of the content of each taste perception as independent of the contents of the others that it could look at all plausible to make the assignments of content that Pautz stipulates. Perception represents objects in complex and systematic ways that allows comparison across different representations in order to judge of different aspects of perceptual similarity and difference. Note that this follows more or less directly from the considerations about information we advanced in the previous section. Carrying information is not done, as it were, state by state. It works only against the background of an ensemble of different potential states, each of which represents an equivalence class of possible specific states. There may be trouble for some versions of tracking intentionalism here but, without further argument, these kinds of cases pose no problem for tracking intentionalism itself.

19.4 Conclusion

There is much more in the paper than the small set of issues we've commented on, including much more on the issues we raise in our comments. Pautz's paper is rich in philosophical and empirical detail; both are worth engaging with. Nevertheless, we don't think that Pautz has succeeded in refuting tracking intentionalism (although some of the arguments may work against some versions of the theory). In particular, we don't think that Pautz has succeeded in casting doubt on the thesis that the phenomenology of perception can be explained in terms of contents involving external properties.

References

- Bornhöyd, K., M. Quante, et al. 2002. Painful stimuli evoke different stimulus–response functions in the amygdala, prefrontal, insula and somatosensory cortex: A single-trial fMRI study. *Brain* 125(6): 1326–1336.
- Coghill, R., C. Sang, J. Maisog, and M. Iadarola. 1999. Pain intensity processing within the human brain: A bilateral, distributed mechanism. *Journal of Neurophysiology* 82(4): 1934–1943.
- Craig, A. 2003. Pain mechanisms: Labeled lines versus convergence in central processing. *Annual Review of Neuroscience* 26(1): 1–30.
- Dretske, F.I. 1981. *Knowledge and the flow of information*. Cambridge, MA: MIT Press.
- Dretske, F. 1995. *Naturalizing the mind*. Cambridge: MIT Press.
- Gottfried, J.A. 2010. Central mechanisms of odour object perception. *Nature Reviews Neuroscience* 11(9): 628–641.
- Howard, J.D., J. Plailly, M. Grueschow, J.D. Haynes, and J.A. Gottfried. 2009. Odor quality coding and categorization in human posterior piriform cortex. *Nature Neuroscience* 12: 932–939.
- Stevenson, R.J., and Donald A. Wilson. 2007. Odour perception: An object-recognition approach. *Perception* 36: 1821–1833.
- Vul, E., C. Harris, et al. 2009. Puzzlingly high correlations in fMRI studies of emotion, personality, and social cognition. *Perspectives on Psychological Science* 4(3): 274–290.
- Zelano, C., A. Mohanty, et al. 2011. Olfactory predictive codes and stimulus templates in piriform cortex. *Neuron* 72(1): 178–187.