Analytic Philosophy Vol. 58 No. 1 March 2017 pp. 58–70

DISCUSSION

PRECISION, NOT CONFIDENCE, DESCRIBES THE UNCERTAINTY OF PERCEPTUAL EXPERIENCE: COMMENT ON JOHN MORRISON'S "PERCEPTUAL CONFIDENCE"

RACHEL N. DENISON
New York University

1. Introduction

Morrison puts forth a view he calls "perceptual confidence," defined as "the view that perceptual experiences assign degrees of confidence" (Morrison 2016). He contrasts this view with "postperceptual confidence," which holds that degrees of confidence are assigned subsequent to perceptual experience. The question of the relation between perceptual experience and confidence is an important one, but the perceptual confidence view has a problem. Namely, degrees of confidence cannot be assigned to an experience; they can only be assigned to a decision outcome, and so cannot be a general attribute of perception. Perceptual uncertainty is not a new idea, but it is separate from and precedes determinations of confidence. Morrison's approach, however, conflates uncertainty and confidence. Here, I raise concerns about the perceptual confidence view. I ask a series of questions about the scope of perceptual confidence, which I suggest is limited at best; discuss how Morrison's construction of perceptual confidence conflicts with the scientific literature; and describe how sensory representations can underlie both uncertain perceptual experience and, only at a later stage, confidence. I propose that perceptual precision, not confidence, is a general approach for describing the uncertainty of perceptual experience.

2. Confidence Reflects Decision Outcomes, Not Perceptual Experience

Given a perceptual experience, we can assign degrees of confidence to certain categories for that experience, but not until we know what the

Analytic Philosophy Vol. 58 No. 1

possible categories are – that is, not until we are faced with a perceptual decision. The perceptual confidence view implies that our perceptual experience is made up of a privileged set of categories; this is a major weakness of this view. Let's consider Morrison's example of perceiving the color of a tablecloth in a candlelit room as more and more candles are lit. At a medium light level, if someone asks, "How confident are you about the color of the tablecloth?" you might answer "Very confident", because you are sure it's red and not blue. Or you might answer "Not very confident", because you are not sure if it's scarlet red or crimson red. The same perceptual experience, then, can be associated with different levels of confidence depending on the perceptual decision undertaken. If someone asks you, "How confident are you about how the tablecloth looks?" you might be unsure how to answer. What visual feature are you being asked to report on?

Confidence, which refers to the subjective probabilities of decision outcomes, requires at least four processing steps: (i) representing perceptual information; (ii) specifying a question about one's perception; (iii) specifying the possible answers (decision outcomes); and (iv) assigning probabilities to those outcomes (see Figure 1). Because confidence cannot be assigned until perceptual decision outcomes are specified, it requires computations beyond those required for the perceptual experience itself. Perceptual experience exists prior to and independent of any particular choice of decision outcomes. It is the same regardless of the specific question you ask about it, and whether you ask any question at all.

To further illustrate how the same perceptual experience can support multiple assignments of confidence, let's take another example: the blurry eye chart. Without your glasses on, you can describe with high confidence what you see: fuzzy, black blobs on a white background. It is only when the optometrist asks you to decide whether the fuzzy blob on top is an E or an F that you report lower confidence for that particular perceptual decision. The same applies to a foggy day, now with your glasses on. You can report quite confidently what you see—white mist, diffuse light, a dark shape in the distance—and you could accurately reproduce what you see, in a painting, for instance (as the Impressionist oeuvre attests). It is only when asked to identify the dark shape in the distance that you report lower confidence for that decision.

3. What Is the Scope of Perceptual Confidence?

These considerations lead to three questions about the scope of the perceptual confidence view:

1: How complete is the assignment of perceptual confidences? As we have seen, for a given perceptual experience, category

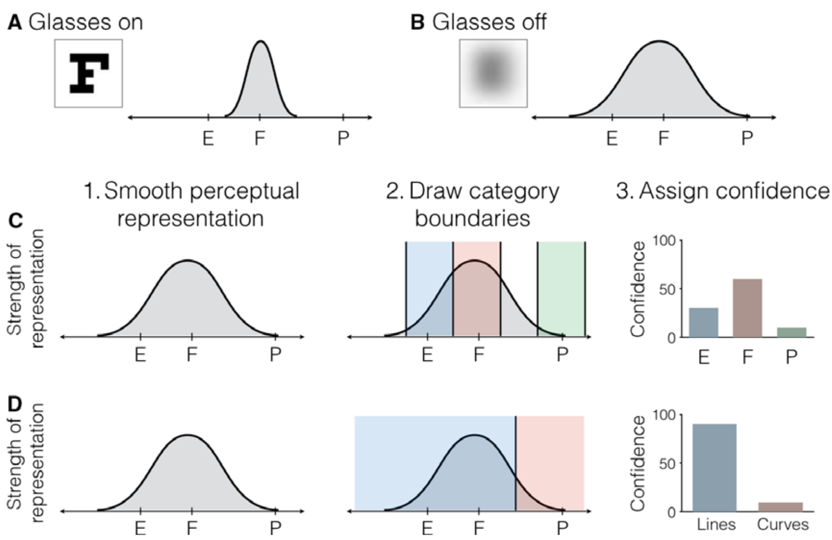

Figure 1: Perceptual precision across a feature space for letter identity. Feature spaces can be high-dimensional, but for visualization purposes, the letter space shown is simplified to one dimension. Nearby letters are meant to be more similar than distant letters. (A) High perceptual precision (narrow distribution) for a sharp image. (B) Low perceptual precision (broad distribution) for a blurry image. (C) A smooth perceptual representation can be transformed into degrees of confidence by drawing category boundaries for different possible decision outcomes, such as different letter identities. Perceptual confidence goes straight to step 3, mistakenly skipping steps 1 and 2. (D) A different question ("Is the letter made up only of straight lines, or does it contain curves?") elicits different category boundaries and confidence ratings. Note that the perceptual representation and corresponding perceptual experience is the same for C and D. [Colour figure can be viewed at wileyonlinelibrary.com]

boundaries can be drawn and confidences assigned in multiple ways, depending on the question asked about one's perception. This raises the question of whether, under Morrison's view, perceptual experiences must assign a massive number of confidences, corresponding to many possible perceptual categorizations, to be prepared for any question about one's experience that might come one's way. If Morrison instead wishes to argue that perceptual experience assigns confidence to only a subset of possible categories, then which perceptual categories come bundled with confidence? And what makes them special enough to receive this extra perceptual attribute?

2: How obligatory is the assignment of perceptual confidence? Morrison seems to claim that confidence is a ubiquitous attribute of perception, applying to all visual features. But perception takes place at many levels simultaneously, from the exact pattern of light and dark over a small region of Isaac's nose, to the overall shape of his nose, to the resemblance of his nose to those of his family members. Do all of these perceptual features simultaneously assign confidences? According to Morrison, perceptual confidence is "conscious, automatic, accessible, . . . and fast" (p. 20). So does each glance give us not only hundreds of perceptual features but hundreds of consciously accessible confidence ratings? If Morrison instead wishes to argue that only some features assign confidence at a time, then what determines which features have this special property?

3: What perceptual features are available for perceptual confidence? Here, I will give my own partial answer to the questions I just raised. Given that perception must be categorized to determine a confidence, only fast, automatic perceptual categorizations could support the generation of a perceptual confidence, as defined by Morrison. How often, and for what types of perceptual categorizations, do these requirements hold?

Morrison's examples seem to include cases that could fit this description of fast categorization, such as categorizing a person as Isaac [i.e. "seeing as" (Block 2014)]. But this intuitive example only works if identifying Isaac is relatively easy, so that categorization is rapid, or if you are already looking for Isaac, so the "Isaac" vs. "not Isaac" categorization has been defined in advance (in a cognitive, not perceptual, step). If you have not predefined "Isaac" as a category and the visual information is not good enough—Isaac is in the periphery, unattended, or poorly lit, for example—you may see Isaac without identifying him (or even the possibility that it might be him), and so without assigning confidence. Even when visual information is good, categorization may not be automatic. Objects seen from non-canonical viewpoints (e.g. a tea kettle viewed from below), for example, may be categorized slowly or not at all (Palmer, Rosch, and Chase 1981). Or if you encounter Isaac in an unexpected context—on the street in Toledo though you know him from Trento—you may fail to recognize him. Without a categorization, confidence cannot be assigned.

But most of Morrison's example features are not rapidly categorized for a more fundamental reason: they do not take on discrete values, but instead vary continuously along some physical dimension, or in magnitude. These features include "illumination, color, shape, and distance" (p. 16), as well as size, line orientation, and many other basic perceptual features. Do we automatically categorize a distance as 7 m as opposed to 8 m? Why not as opposed to 7.1 m? Confidence will be different for different decision boundaries (higher confidence for the comparison to 8 m vs. 7.1 m), so these choices and their automaticity matter (see Figure 1). Much of our perceptual experience consists of continuous features, and confidence is ill-suited to capture our phenomenal experience of these.

Therefore, even if we accept Morrison's assertion that if an assignment of confidence is fast, automatic, conscious, etc., then we should call it perceptual, the requirement for rapid categorization seriously limits the possible scope of the perceptual confidence view. Morrison's overly broad application of perceptual confidence stems from the fact that he blurs together the four separate processing steps required to determine a confidence. So while we can say that our perceptual experience is uncertain (an old idea) and also that we have degrees of confidence about our perceptual experiences (another old idea), confidence is not a general way to characterize the uncertainty of phenomenal experience, such that we should say it comes part and parcel with perception itself.

4. Scientific Approaches to Perceptual Confidence

Morrison claims that his account of perceptual confidence "fills a hole in our best scientific theories of perception" (p. 15). However, the notion that confidence is obligatorily assigned by perception is actually at odds with the scientific literature. Scientists studying this topic also use the term "perceptual confidence" to refer to the subjective likelihood of a perceptual decision outcome (e.g., Hebart et al. 2016; Koizumi, Maniscalco, and Lau 2015). But they generally adopt a much more flexible model for how people determine perceptual confidence, which respects the processing steps of representation, categorization, and likelihood assessment.

In a typical laboratory study on perceptual confidence (Peirce and Jastrow 1884), a participant is presented with a sensory stimulus and asked two questions: (i) to which of two categories does the stimulus belong? and (ii) how confident are you that it belongs to that category?[1] Observers might adjust a slider or give a rating on a discrete scale (e.g. 1–4) to report their confidence. Note that the experimenter determines the categories (which can be completely arbitrary), and the participant gives a confidence report that depends on these category definitions. If the experimenter changes the category definitions, the participant will report different levels of confidence, decoupling perception and confidence reports.

Studies like these have led to the view that confidence is flexibly "read out" from neural representations (Fleming and Dolan 2012;

[1] The way confidence is defined in these studies is not exactly the same as the way Morrison defines it. Participants report confidence in a perceptual decision—the subjective probability that their categorical decision was correct—as opposed to separate confidences for each decision outcome; though the two notions of confidence are obviously related. Interestingly, in practice, they may be even more similar than they first appear: people's confidence reports are tightly linked to the evidence for the category they selected, without much dependence on the other category (Zylberberg, Barttfeld, and Sigman 2012).

Hilgenstock, Weiss, and Witte 2014; Kepecs and Mainen 2012; Yeung and Summerfield 2012; Zizlsperger et al. 2014), rather than being a basic perceptual attribute. This arrangement, critically, allows degrees of confidence to be determined for any possible question and set of decision outcomes. Scientists use the term "metacognition" (Fleming, Dolan, and Frith 2012) to refer to the process of determining one's confidence about a perceptual experience; the use of "cognition" reflects the fact that some decision has to be made—at the very least, possible decision outcomes must be specified—before confidence can be assigned.

Empirically, the representations that support the determination of confidence can be both perceptual and nonperceptual, with nonperceptual factors including the time it takes to make the response (Kiani, Corthell, and Shadlen 2014) and action-related neural activity (Fleming et al. 2015). Manipulations of brain activity using neurofeedback in frontal and parietal cortical areas can alter confidence reports without affecting accuracy, even though these areas seem to carry little information about the perceptual signal (Cortese et al. 2016). The fact that people's confidence reports are not based only on perceptual factors should lead us to question Morrison's central claim that "you endorse your experience" (p. 27) when reporting confidence, even when instructed to report about perceptual information alone. Empirical work also casts doubt on Morrison's notion that confidence is as rapid as perception: when observers are asked to make a fast perceptual decision, they need more time to report their confidence (Baranski and Petrusic 1998). In the sizeable and growing body of scientific literature on perceptual confidence, it has proven more useful to conceptualize confidence as the outcome of a flexible, higher-level decision process than as a basic perceptual attribute (Fleming and Daw 2017).

5. Perceptual Precision Better Characterizes the Uncertainty of Perceptual Experience

There is an important sense in which perceptual experience feels uncertain—and more so with glasses off than with glasses on. We would like a way to capture this uncertainty that does not depend on a particular perceptual question, but instead is linked to what our confidence would be across all possible decision outcomes to all possible questions. The importance of distinguishing between confidence and certainty has been recognized in the scientific literature (Pouget, Drugowitsch, and Kepecs 2016). If we think of confidence as being read out, at least in part, from a perceptual representation, then the uncertainty of perceptual experience is best described by characterizing the perceptual representation itself. In the spirit of providing a positive alternative to perceptual confidence, I will suggest one way to characterize the uncertainty of a perceptual representation: namely, by quantifying what I will call the "perceptual precision".

Perceptual precision refers to the discriminability of perceptual feature values over some feature dimension (Figure 1A,B). With glasses off, perceptual experience has low precision across space, making the world look blurry. As a result, it also has low precision across letter identity, making it difficult for us to distinguish Es from Fs (Figure 1C). It might be easier to make coarser distinctions between letters: for example, to distinguish letters made up only of straight lines from letters with curves (Figure 1D). Confidence about either of these sets of decision outcomes could be read out from this single underlying perceptual representation. This flexibility is possible because the representational space is smooth, unlike the categorical representational space needed to assign confidence (Figure 1C,D; also reflected in Morrison's bar graphs). Similarly, in the candlelit room example, our perceptual precision across the color dimension depends on the light level. At a medium light level, we can discriminate colors that are very different (red and blue) but not colors that are very similar (scarlet red and crimson red). At a higher light level, our perceptual precision improves, and we can make finer discriminations. Perceptual precision is an established concept in psychology and neuroscience. There are well-developed methods to measure perceptual precision using tasks in which observers discriminate between two similar items or estimate the value of a particular stimulus feature (Kingdom and Prins 2010).

Scientists often characterize sensory representations as smooth distributions across a feature space (representing line orientation, spatial location, etc.). Given such a smooth perceptual representation, as shown in Figure 1 for letter identity, the discriminability of two feature values (e.g., E and F) is directly related to the difference between their representational strengths. This difference depends on the perceptual precision. When precision is high (Figure 1A), the difference between the representational strengths of the two letters is large (F has a high strength and E has zero strength). When precision is low (Figure 1B), the difference is reduced (the strength of E now has an intermediate value), so discriminability is also reduced. For a one-dimensional feature, like line orientation, we can think (simplistically) of the representational strength of a particular orientation as the firing rate of a neuron tuned to that orientation. In general, the strength will be determined by a population code, combining information across many neurons.

Critically, perceptual precision depends not only on the representational strength of a feature value, but also on the variance, or spread, of the representation across values. A representation could have a small magnitude but relatively high precision if its distribution is very narrow. In this case, you might say that the stimulus is not very strong, but you are quite sure what it is. Conversely, a representation could have a large magnitude but low precision if its distribution is broad. In this case, you might say that you can easily see the stimulus, but it's hard to say exactly what it is. Variance and signal-to-noise ratio (the magnitude divided by the standard deviation) are standard ways to characterize the uncertainty of sensory representations. These metrics characterize perceptual uncertainty more parsimoniously than confidence: a single number summarizes uncertainty across the feature dimension, whereas confidence requires a separate number for each category bin.

6. Comparisons with Perceptual Confidence

We might ask, is "representational strength" just another term for Morrison's perceptual confidence? The answer is no—or if it is, that would be a confusing use of the word "confidence" at odds with the way the term is normally used. In reporting one's confidence that the letter on the eye chart is an E, an observer would not merely report the strength of the representation at the point corresponding to the "ideal" E (marked with a tick and "E" label on the x-axes of the distribution plots in Figure 1). Instead, the observer must draw category boundaries specifying what counts as an E, which can then be used to compute confidence (Figure 1C). Two people could have the same perceptual experience but draw their category boundaries in different ways, leading to different levels of confidence. This is a classic issue of criterion-setting (Macmillan and Creelman 2005). Furthermore, different types of perceptual decisions will lead to different ways of drawing category boundaries and different reports of confidence, while the underlying perceptual representation and corresponding perceptual experience remain constant (compare Figure 1C,D).

Unlike perceptual confidence, which reflects a degree of belief about a decision outcome, perceptual precision describes uncertainty across an entire feature dimension. It defines how precisely possible decision outcomes could be distinguished even before a particular perceptual question is specified. It also corresponds naturally to the likely underlying neural representation of perceptual features. For example, our lower spatial precision in the periphery corresponds to the smaller cortical territory devoted to processing peripheral compared to central vision. Or in the eye chart example, with glasses off, light is literally smeared across the retina (compared to with glasses on, when light is well-focused on the retina). As a result, the representation of the visual scene is smeared across the cortex. This physical blurring of the underlying neural representation decreases our perceptual precision in an intuitive way.

What type of neural representation does perceptual confidence have, in Morrison's view? Rather than focus on specific brain regions or any biological details of the neural implementation, we can ask about both the representational format of confidence and the computational ingredients required to assign confidence. I have suggested that the representation of confidence, based on its definition, must be categorical and discrete across a feature space. This interpretation fits with Morrison's use of bar graphs to illustrate perceptual confidence. I have also suggested that the computational ingredients include not only sensory information, but flexible, context-dependent decision boundaries. Should we call neurons that compute likelihoods from these types of inputs "perceptual"?

7. Perceptual Experience Need Not Feel Probabilistic

A central assertion of the perceptual confidence view is that perceptual experience is probabilistic in the sense that it represents the likelihoods of different perceptual possibilities. In the perceptual precision view, this need not be the case. The smooth distributions over feature values proposed to underlie perceptual precision could, of course, be formally treated as probability distributions by scaling them so that their integrals sum to one. Indeed, it might be useful for various neural computations to treat sensory representations as probability distributions (Ma et al. 2006). But even if one can interpret a perceptual representation as probabilistic, the associated perceptual experience need not be[2] . Here are a couple of examples:

7.1 Dots and Smudges

Consider looking at a black dot on a white piece of paper (Figure 2, left). With glasses off, you see a blurry dot. In the underlying representation, the representational strength is highest at the location of the dot and falls off gradually across a somewhat larger region of space. According to a probabilistic interpretation, the dot is most likely to be in the center of the blur and slightly less likely to be at positions further from the center. According to a nonprobabilistic interpretation, there is a black smudge on the piece of paper (not a crisp dot). These are equally valid interpretations, and perceptual experience cannot distinguish between them. To make this clear, put the glasses back on and look at an actual black smudge on a white piece of paper (Figure 2, right). The underlying perceptual representation and associated perceptual experience are *identical* to that of the blurry black dot. But now that the smudge is physical, a probabilistic interpretation seems much less appealing. In both cases—the blurry dot and the smudge—it is most straightforward to say simply that the experience is imprecise. We need not, and in some cases clearly should not, invoke probabilities.

[2] In fact, I think it is notable how nonprobabilistic perceptual experience feels—things look exactly the way they look.

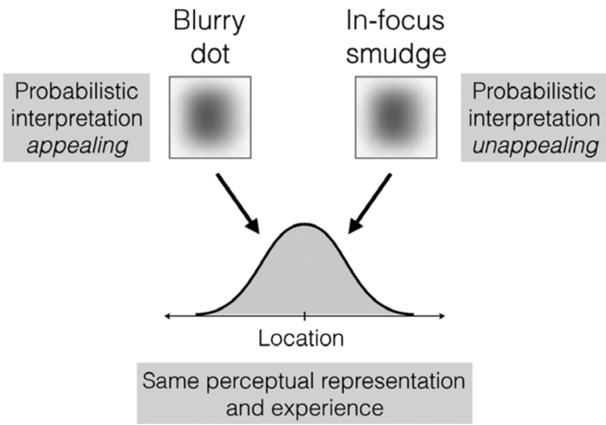

Figure 2: The perceptual representation and associated experience are identical for a crisp black dot that is blurry because your glasses are off (left) and a black smudge that is perfectly in focus with glasses on (right). While a probabilistic interpretation of the blurry dot's location might be appealing, a probabilistic interpretation of the location of the smudge is not—the smudge is actually spread across a range of locations. Perceptual precision, unlike confidence, does not commit to probabilistic interpretations of perceptual experience.

7.2 Bistability

If there is any situation in which we might expect perception to give us a probabilistic experience, it is when viewing a bistable image. Bistable images have two possible interpretations that are more or less equally likely but are mutually incompatible. If these probabilities translated into perceptual experience, then we should experience both perceptual interpretations, with 50% confidence for each. But this is not what happens. Instead we see only one interpretation at a time. Why does the blurry eye chart feel uncertain while bistable images feel certain—such that even a small bias toward one interpretation results in a fully certain experience?

Perceptual confidence cannot explain these different types of perceptual experiences. Because this view blurs together separate stages of processing, the implied representation is the same for blurry and bistable images (Figure 3). Perceptual precision, though, helps us think about the representation of a bistable image as a bimodal distribution over some feature space (Figure 3). Each peak of the distribution can be quite narrow—meaning high perceptual precision, and a feeling of high certainty. The presence of two peaks results in a selection process, such that only one interpretation is perceived at a time. The blurry image, on the other hand, has a broad, unimodal

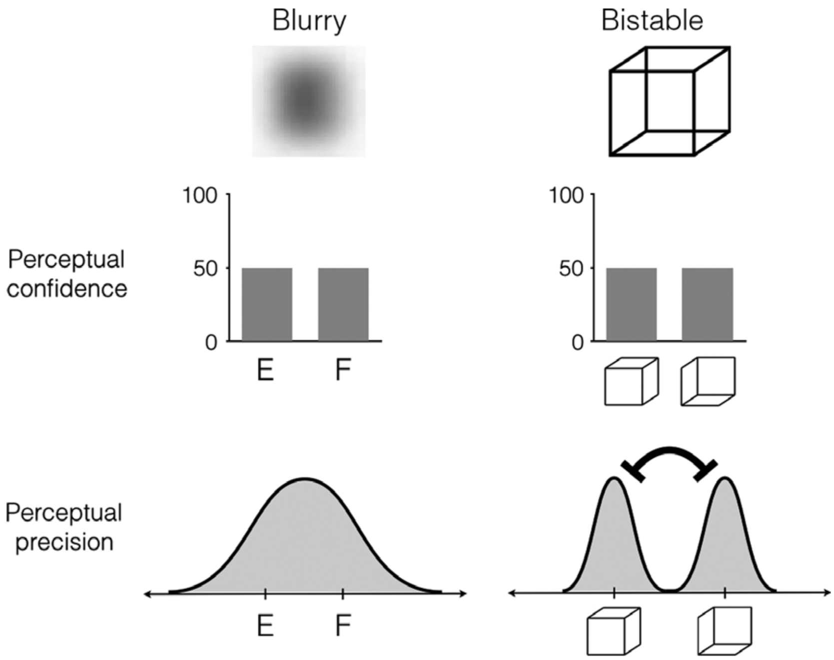

Figure 3: Two cases in which an image has multiple possible perceptual interpretations (top row): the blurry eye chart letter (left) and bistable Necker cube (right). The Necker cube can be seen as though from two different perspectives. Perceptual confidence (middle row) treats the blurry and bistable images similarly. Perceptual precision (bottom row) captures their different perceptual representations (unimodal for the blurry image and bimodal for the bistable image), explaining why they are associated with different selection processes and feelings of perceptual certainty. The unimodal peak between the two letters indicates that the blurry image looks like something in between E and F. The thick curve with end stops indicates perceptual competition between the two interpretations of the bistable image, such that only one is seen at a time (Blake and Logothetis 2002).

distribution over the relevant feature space (Figure 3). The broad peak reflects the low perceptual precision and corresponding feeling of uncertainty.

8. Reconciling Perceptual Precision and Perceptual Confidence

While perceptual precision captures the uncertainty of perceptual experience better than perceptual confidence for the reasons I have described, the general concept of perceptual confidence should not be discarded. Once we undertake a perceptual decision, we can sensibly say that we have degrees of confidence in the decision outcomes.

At that point, separating perceptual confidence from doxastic confidence is important, for all the reasons laid out by Morrison. However, doxastic confidence about our perceptual experience has to come from somewhere, and it does not seem particularly new to say that it mostly comes from the uncertainty in our perceptual representations. It is Morrison's stronger claim that those perceptual uncertainties are confidences that I dispute.

In particular, I have questioned the generality of the perceptual confidence view. I have argued that conscious experience can and does occur before perceptual decision outcomes are specified—we do not have to categorize our experience in order to have it. I have also argued that perceptual experience is not inherently probabilistic (the dot vs. the smudge and bistable images). However, I do not wish to say that confidence is never perceptual, in Morrison's sense that it can be rapidly generated, concern features of the environment, and be independent of other beliefs. Rather I suggest, consistent with the scientific literature, that perceptual representations across feature spaces are used to make categorical decisions as well as to determine our confidence in the possible decision outcomes (see Figure 1). These different types of conscious representations may (though do not always) coexist, creating a many-layered perceptual experience.

Acknowledgments

I thank Ned Block, Petra Vetter, and members of the NYU Department of Philosophy for critical feedback and helpful discussion.

References

Baranski, Joseph V., and William M. Petrusic. "Probing the Locus of Confidence Judgments: Experiments on the Time to Determine Confidence," *Journal of Experimental Psychology: Human Perception and Performance* 24 (1998): 929–945.

Blake, Randolph, and Nikos K. Logothetis. "Visual Competition," *Nature Reviews Neuroscience* 3 (2002): 13–21.

Block, Ned. "Seeing as in the Light of Vision Science," *Philosophy and Phenomenological Research* 89 (2014): 560–572.

Cortese, Aurelio, Kaoru Amano, Ai Koizumi, Mitsuo Kawato, and Hakwan Lau. "Multivoxel Neurofeedback Selectively Modulates Confidence without Changing Perceptual Performance," *Nature Communications* 7 (2016): 13669.

Fleming, Stephen M., and Raymond J. Dolan. "The Neural Basis of Metacognitive Ability," *Philosophical Transactions of the Royal Society B: Biological Sciences* 367 (2012): 1338–1349.

Fleming, Stephen M., Raymond J. Dolan, and Christopher D. Frith. "Metacognition: Computation, Biology and Function," *Philosophical Transactions of the Royal Society B: Biological Sciences* 367 (2012): 1280–1286.

Fleming, Stephen M. et al. "Action-Specific Disruption of Perceptual Confidence," *Psychological Science* 26 (2015): 89–98.

Fleming, Stephen M., and Nathaniel D. Daw. "Self-evaluation of Decision-making: A General Bayesian Framework for Metacognitive Computation," *Psychological Review* 124 (2017): 91–114.

Hebart, Martin N., Yoren Schriever, Tobias H. Donner, and John-Dylan Haynes. "The Relationship between Perceptual Decision Variables and Confidence in the Human Brain," *Cerebral Cortex* 26 (2016): 118–130.

Hilgenstock, Raphael, Thomas Weiss, and Otto W. Witte. "You'd Better Think Twice: Post-Decision Perceptual Confidence," *NeuroImage* 99 (2014): 323–331.

Kepecs, Adam, and Zachary F. Mainen. "A Computational Framework for the Study of Confidence in Humans and Animals," *Philosophical Transactions of the Royal Society B: Biological Sciences* 367 (2012): 1322–1337.

Kiani, Roozbeh, Leah Corthell, and Michael N. Shadlen. "Choice Certainty Is Informed by Both Evidence and Decision Time," *Neuron* 84 (2014): 1329–1342.

Kingdom, Frederick A. A., and Nicolaas Prins. *Psychophysics: A Practical Introduction*. London: Elsevier, 2010.

Koizumi, Ai, Brian Maniscalco, and Hakwan Lau. "Does Perceptual Confidence Facilitate Cognitive Control?," *Attention, Perception, and Psychophysics* 77 (2015): 1295–1306.

Ma, Wei Ji, Jeffrey Beck, Peter Latham, and Alexandre Pouget. "Bayesian Inference with Probabilistic Population Codes," *Nature Neuroscience* 9 (2006): 1432–1438.

Macmillan, Neil A., and Douglas C. Creelman. *Detection Theory: A User's Guide*, 2nd ed. Mahwah, NJ: Lawrence Erlbaum Associates, 2005.

Morrison, John. "Perceptual Confidence," *Analytic Philosophy* 57 (2016): 15–48.

Palmer, Stephen, Eleanor Rosch, and Paul Chase. "Canonical Perspective and the Perception of Objects." *Attention and performance IX*. Eds. J. Long, A. Baddeley. Hillsdale, NJ: Erlbaum, 1981. 135–151.

Peirce, Charles, and Joseph Jastrow. "On Small Differences of Sensation," *Memoirs of the National Academy of Sciences* 3 (1884): 75–83.

Pouget, Alexandre, Jan Drugowitsch, and Adam Kepecs. "Confidence and Certainty: Distinct Probabilistic Quantities for Different Goals," *Nature Neuroscience* 19 (2016): 366–374.

Yeung, Nick, and Christopher Summerfield. "Metacognition in Human Decision-Making: Confidence and Error Monitoring," *Philosophical Transactions of the Royal Society B: Biological Sciences* 367 (2012): 1310–1321.

Zizlsperger, Leopold, Thomas Sauvigny, Barbara Händel, and Thomas Haarmeier. "Cortical Representations of Confidence in a Visual Perceptual Decision," *Nature Communications* 5 (2014): 3940.

Zylberberg, Ariel, Pablo Barttfeld, and Mariano Sigman. "The Construction of Confidence in a Perceptual Decision," *Frontiers in Integrative Neuroscience* 6 (2012): 79.