Analytic Philosophy Vol. 58 No. 1 March 2017 pp. 71–85

PERCEPTUAL CONFIDENCE AND CATEGORIZATION

JOHN MORRISON
Barnard College, Columbia University

1. Introduction

In "Perceptual Confidence," I argue that our perceptual experiences assign degrees of confidence. In "Precision, not Confidence, Describes the Uncertainty of Perceptual Experience," Rachel Denison disagrees. I believe that most of our disagreements are merely terminological, because they just reflect differences in how we use the terms 'perceptual experiences', 'assign' and 'confidence'. If I'm right, only two of our disagreements are substantive, in particular whether perception involves automatic categorization, and whether there is an intrinsic difference between a blurry perception of a sharp image and a sharp perception of a blurry image.

In the next section I will introduce a framework for discussing all our disagreements.

2. Tasks

It will be helpful to consider three tasks. Let's start with a categorization task. Suppose there's a figure in the distance, and you're asked to categorize the figure as Isaac or as not Isaac. Your response will depend on a measurement in your brain. For example, it might depend on the activity in a certain cluster of neurons in the fusiform face area. As a result of this measurement, a subjective probability will be assigned to the possibility it's Isaac.[1] What subjective probability will be assigned? There isn't a straightforward answer. It depends on a number of factors, potentially including the subjective probabilities previously assigned to other possibilities, such as the possibility of getting that measurement given that it's Isaac, the possibility of getting that measurement given that it's not Isaac, and the possibility that it's Isaac, independent of any measurement. It also depends on how these factors are combined, such as whether they are combined in accordance with Bayes's theorem. There is disagreement about which factors are combined and how they are combined. But we can set these disagreements to the side. Suppose that, as a result of getting a certain measurement, a subjective probability of

[1] As I'm using 'subjective probability', it's a probability assigned for the purposes of making a decision, and it needn't be conscious.

Analytic Philosophy Vol. 58 No. 1

.55 is assigned to the possibility that it's Isaac, and that, as a result, there is a decision to categorize the figure as Isaac. We can represent this subjective probability using a horizontal bar graph:

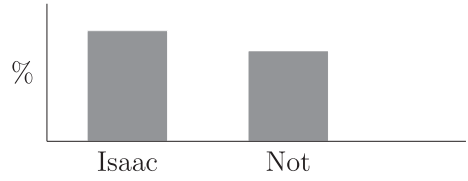

This subjective probability might inform other decisions as well, such as whether to wave at the approaching figure, and whether to accept certain bets, such as a bet that pays $1 if it's Isaac and costs $2 if it's not.

Let's next consider an estimation task. Suppose there's an oval in the periphery of your visual field, and you're asked to estimate its elongation. Once again, your response will depend on a measurement in your brain. For example, your response might depend on the activity in a certain cluster of neurons in the visual cortex. As a result of this measurement, subjective probabilities will be assigned to a continuous range of elongations. What subjective probabilities will be assigned? Once again, it depends on a number of factors, perhaps including the subjective probabilities already assigned to other possibilities. There are disagreements about which factors are combined and how they are combined. But let's again set these disagreements to the side. Suppose that, as a result of getting a certain measurement, subjective probabilities are assigned in a normal distribution, centered on a particular elongation, and that, as a result, there's a decision to estimate the oval's elongation as a certain value. We might represent this distribution of subjective probabilities using a line graph:

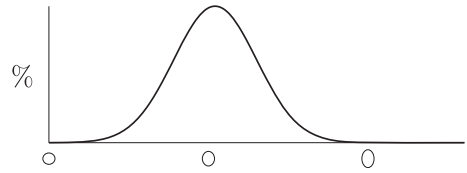

These subjective probabilities might inform other decisions as well, such as whether to reach for the object, and whether to accept certain bets, such as a bet that pays $1 if its elongation is greater than ◯ and cost $2 if it's not.

What's the difference between this task and the categorization task? In the categorization task, there were only two possibilities, the possibility that it's Isaac and the possibility that it's not. As a result,

subjective probabilities were assigned to just these possibilities. In the estimation task, there were infinitely many possibilities because there are infinitely many elongations. As a result, subjective probabilities were assigned to a continuous range of elongations. We might think of these tasks as the two limiting cases of a signal detection task, with all other signal detection tasks, such as those involving three possibilities, four possibilities, etc., falling in between. For the purposes of this paper, it will be helpful to think of decisions involving a finite number of possibilities as categorization tasks, and decisions involving an infinite number of possibilities as estimation tasks.

Finally, let's also consider a metacognitive task. Suppose that there's a decision to categorize the figure as Isaac, and you're then asked whether that decision is correct. What subjective probability will you assign to the possibility that decision is correct? The answer might seem obvious. If a subjective probability of .55 was assigned to the possibility it's Isaac, and that subjective probability was responsible for the decision that it's Isaac, it might seem obvious that you will assign a subjective probability of .55 to the possibility that the decision is correct. And that might be exactly what happens in most cases. But not all cases. In some cases, you might have access to the outcome of a decision, but not to the states responsible for that decision. In such cases, you might need to rely on your background beliefs about the reliability of whatever process made that decision for you. For example, if you believe that the process has a reliability of .7, you might assign a subjective probability of .7 to the possibility the decision is correct, even though in this case the process was relying on a state that assigned a subjective probability of .55. In other cases, you might not have control over the factors responsible for the state assigning a subjective probability of .55. In such cases, you might rely on background beliefs that try to compensate for whatever was non-ideal about the relevant factors. For example, you might believe that the relevant factors lead the relevant process to assign a subjective probability that's too high when the figure is far away, and thus compensate by assigning a lower subjective probability that the decision was correct. In still other cases, you might take into account evidence that wasn't taken into account by the process responsible for the decision. For example, you might take into account the amount of time it took to reach a decision, or testimonial evidence. For example, you might have access to the fact that the decision was based on a subjective probability of .55, but later assign a much lower subjective probability, given how long it took to reach that decision, or that you heard something to the contrary.

What's the difference between this task and the categorization and estimation tasks? In the categorization and estimation tasks, the assignment of subjective probabilities occurs *before* the initial decision that it's Isaac, or that the oval has a certain elongation. In the metacognition task, there's another assignment of subjective probabilities that occurs *after* that initial decision. Another related difference

is that in the categorization and estimation tasks, the relevant possibilities are about external objects, such as whether the figure is Isaac. In the metacognition task, the relevant possibilities are about a decision. In many cases, these aren't important differences, because, for example, the same subjective probability that was assigned to the possibility it's Isaac and to the possibility the decision is correct. But, in other cases, like the cases we just described, these are important differences.

My descriptions of these tasks should be familiar from any introductory psychophysics textbook (Kingdom and Prins 2010, Ch 2). But notice that I didn't say anything about *perceptual consciousness*. What role does it play in these tasks? This is an important question, because consciousness is a fundamental feature of the mind, and perhaps even a definitive feature of the mind. It also has potential implications for our access to, and control over, the underlying factors. Nonetheless, because consciousness is so hard to measure, neuroscientists are reluctant to speculate. This leaves a hole in their theories.

Let's consider two views about how to fill that hole. The first view is that, whenever perceptual consciousness includes one of the relevant possibilities, it just includes that possibility, and doesn't assign it a subjective probability. I call this POST-PERCEPTUAL CONFIDENCE. Among philosophers, it is the traditional view. It's important to be clear about what this view doesn't rule out. First, perceptual consciousness might still result from earlier decisions in the visual system, decisions that involve assignments of subjective probabilities. Your perceptual consciousness might result from an automatic process that assigns subjective probabilities to many possibilities, and then, at the end, selects one of those possibilities to include in your perceptual consciousness. The relevant process might be quite sophisticated, taking into account information about the figure's hair color, height, eye distance, and complexion. We can describe the result of this process as a *decision* insofar as the process relies on measurements and criteria. In this loose sense of 'decision', decisions needn't involve deliberation; they can be automatic. Second, even if your perceptual consciousness includes a representation of one of the relevant possibilities, but doesn't assign a subjective probability, you might later assign a subjective probability. For example, if it selects the possibility that the figure is Isaac, but you believe that the relevant process is unreliable at categorizing figures in the distance, you might later assign a low subjective probability to the possibility it's Isaac, even if that's what you perceive. Relatedly, if you know that Isaac is out of town, you might later assign zero subjective probability to the possibility it's Isaac, even if that's what you perceive. Thus, this view is compatible with the assignment of subjective probabilities both before and after perceptual consciousness. It just denies that perceptual consciousness itself ever assigns subjective probabilities.

The second view is that perceptual consciousness itself assigns subjective probabilities. I call this PERCEPTUAL CONFIDENCE. Once again, it's important to be clear about what this view doesn't rule out. First, perceptual consciousness might still result from earlier decisions in the visual system. In particular, your perceptual consciousness might result from an automatic process that decides what subjective probabilities to include in your perceptual consciousness. For example, if the relevant process relies on activity in the fusiform face area, it might decide to assign a subjective probability of .55 to the possibility it's Isaac whenever the activity in the region responsible for identifying Isaac is in a certain range. If the activity in that region drops enough, the relevant process will decide to include a lower probability, and if the activity in that region increases enough, the relevant process will decide to include a higher probability. The relevant process might also be quite sophisticated, taking into account other information about the figure. Second, even if your perceptual consciousness assigns subjective probability to a possibility, you can subsequently adjust that amount. In an extreme case, like when you know that Isaac is out of town, or that you're in a wax museum, you might later assign zero subjective probability to the possibility that it's Isaac, despite the fact that your conscious perception assigned a high probability to that possibility.

An analogy might help clarify the difference between these two views. According to POST-PERCEPTUAL CONFIDENCE, perceptual consciousness is like a weatherman who tells you that it will rain, or that it won't rain, but never tells you how much subjective probability he assigns to those possibilities. As a result, when deciding how much subjective probability to assign to the possibility it will rain tomorrow, you can't just rely on his report. You must also rely on your background beliefs, including beliefs about his reliability. According to PERCEPTUAL CONFIDENCE, your experience is like a weatherman who tells you how much subjective probability he assigns to the possibility that it will rain tomorrow. As a result, you could just rely on his report. But you don't need to. Even if he assigns a high probability to the possibility it will rain tomorrow, you might assign much less, because you believe he's unreliable.

As I said, I think that most of my disagreements with Denison are merely terminological. In the next section I will use this framework to clarify how I'm using the terms 'perceptual experience', 'assign', and 'confidence'. In the remaining sections I will address what I take to be our only substantive disagreements, namely whether perception involves automatic categorization, and whether there is an intrinsic difference between a blurry perception of a sharp image and a sharp perception of a blurry image. Because I believe the first of these disagreements is more fundamental, I will spend more time discussing it.

3. Terminology

The same term is often used in different ways in different disciplines, making it hard to communicate across disciplinary boundaries. For example, some neuroscientists use the term 'belief' to describe encodings of probability distributions that are exploited in early sensory processing and that are inaccessible by us (such as our "perceptual posteriors"). Philosophers, in contrast, use the term 'belief' to describe a mental state with a closer connection to report and deliberate action, and thus won't say that early sensory processing exploits beliefs (Schwitzgebel 2015). Given the different ways that neuroscientists and philosophers use the term 'belief', it is easy to image a neuroscientist and a philosopher finding themselves in a disagreement about whether beliefs play a role in early sensory processing, not realizing that their disagreement is merely terminological. I think that most of my disagreements with Denison are similar. I'm therefore grateful for the opportunity to clarify how I'm using 'perceptual experience', 'assign', and 'confidence'.

Let's start with 'confidence'. I use 'confidence' to describe *any* assignment of subjective probability. In the categorization task, in which a subjective probability is assigned to the possibility that it's Isaac and the possibility that it's not, I'd say that confidence is assigned (Morrison 2016, p.15). In the estimation task, in which a subjective probability is assigned to a continuous range of elongations, I'd again say that confidence is assigned (Morrison 2016; p.19, 32–33). In the metacognition task, in which a subjective probability is assigned to the possibility a decision is correct, I'd yet again say that confidence is assigned.

Denison uses it more restrictively, to describe only the assignment of subjective probability to the possibility that a decision is correct or incorrect. She writes that, "confidence... refers to the subjective probability of decision outcomes" (Denison 2017, p.59). Thus, in the categorization and estimation tasks, she wouldn't say that confidence is assigned. She would, however, say that confidence is assigned in the metacognition task, because in that task a subjective probability is assigned to the possibility that a decision is correct.

For what it's worth, philosophers standardly use 'confidence' in my way (e.g., Christensen 2004; Hájek 2012; Joyce 1999; Schwitzgebel 2015; Steele and Stefánsson 2016). Some neuroscientists use 'confidence' in my way as well. For example, Meyniel, Sigman and Mainen (2015) use 'distributional confidence' to describe subjective probability distributions over many possibilities, and 'summary confidence' to describe a single number that's derived from that distribution, such as the mean, and the subjective probability a decision is correct (p.79). Other neuroscientists, and perhaps even most other psychologists, use it Denison's way. But even these neuroscientists acknowledge

that it is used in contradictory ways. For example, as Denison notes (p.63), Pouget, Drugowitsch, and Kepecs (2016) use 'confidence' to describe the subjective probability that a decision is correct or incorrect, and 'certainty' to describe all other probability distributions (p.366). But even they present themselves as offering a *proposal* about how these terms should be used in the future, to minimize terminological confusion, given that these terms are currently used in contradictory ways (p.367). Thus, I think that Denison is wrong to insist that any other use of 'confidence' is "confusing" and "at odds with the way the term is normally used" (p.65).

From this point forward, I'll use 'confidence' in my way, in the hope that those who prefer another term, such as 'certainty', will know how to translate what I say back into their preferred terminology.

Let's next address 'perceptual experience'. I use 'perceptual experience' to describe a state that's conscious, automatic, accessible, dissociable from doxastic states, directed toward perceived objects and properties, and fast enough that we can't detect any delay (p.20). Thus, when I say that our perceptual experiences assign confidence, I am talking about a state that's conscious. And when I talk about perceptual confidence, I am talking about the confidence assigned by a state that's conscious.

Some neuroscientists use 'perception' differently, so that, by definition, perceptions occur before *any* decision has been made about the stimulus (for a critical review, see Witt et al. 2015). If I were using 'perception' in this way, then Denison would be right that it is old news that perceptions assign confidence (p.69). For example, consider the experiments involving cue combination that I discuss in the paper (Morrison 2016, p.24–27). Those experiments might establish that there are states that assign confidence before any decision has been made about the stimulus. That's old news. But what those experiments don't establish is that there are *conscious* states that assign confidence. As noted above, neuroscientists are reluctant to speculate about the role of consciousness, because their experiments don't seem to indicate the role of consciousness. Thus, it *would* be news if perceptual experiences assigned confidence; it would fill a hole in our best scientific theories.

Let's next address 'assigns'. I'm using this term so that to say that confidence is *assigned* is to describe a relation to the relevant possibility (Morrison 2016, p.21, 37–8), and doesn't imply anything about how confidences are encoded in the brain. For example, it doesn't imply that they are encoded linearly, and it doesn't imply that they are encoded apart from all other variables, such as signal strength or utility.

I just clarified how I use the terms 'perceptual experience', 'assign', and 'confidence'. Let's now consider one of Denison's terms, 'precision'. She uses it to describe a subjective probability that's distributed over a continuous range of possibilities, and that's encoded together

with signal strength (p.64–65). I hope it's now clear that, if this subjective probability is included in perceptual consciousness, it's just a special instance of what I call 'perceptual confidence'. And if this subjective probability isn't included in perceptual consciousness, it's not an alternative to what I'm calling 'perceptual confidence', because I'm trying to describe what's included in perceptual consciousness. Thus, in any case, the disagreement about whether perception involves confidence or precision is merely terminological.

In the next two sections I'll address what I take to be our only substantive disagreements. But I first want to mention another of our agreements: that perceptual experiences don't assign degrees of confidence to decisions. Denison lists third-personal evidence, including that people are faster at reporting how they categorized a stimulus than at reporting a degree of confidence in their categorization (Baranski and Petrusic 1998). I also think there's first-personal evidence. When you open your eyes, you perceive people, lights, tablecloths, letters, basketballs, ovals, and so on. You don't perceive *decisions*. For this reason, if our perceptual experiences assign degrees of confidence, it is just to possibilities involving people, lights, tablecloths, letters, basketballs, ovals, and so on, not to decisions. This isn't to deny that decisions play an important role. As noted above, your perceptual experiences might assign a particular degree of confidence to the possibility it's Isaac *as a result of* an earlier decision to assign that much confidence. That is, your perceptual experience's assigning confidence to that possibility might be a *decision outcome*. But that doesn't imply that our perceptual experiences assign a degree of confidence *to that decision*. Analogously, a weatherman might decide to assign a certain confidence to the possibility it will rain tomorrow, without thereby assigning a degree of confidence to his own decision. It might not even occur to him to assign confidence to his own decision.[2]

4. Categorization

If I've understood her correctly, Denison and I agree that perceptual experiences assign degrees of confidence over continuous ranges of possibilities, such as continuous ranges of colors, elongations, and line orientations (p.61, 65; Morrison 2016 p.20–21, 32–33). But we

[2] Given PERCEPTUAL CONFIDENCE, there's another reason to deny that perceptual experiences assign confidence to decisions: there hasn't yet been a decision of the right kind. For example, if your perceptual experience assigns confidence to the possibility that it's Isaac, there was presumably an earlier decision to assign confidence. But there presumably wasn't an earlier decision that it's Isaac, or your perceptual experience would have assigned full confidence to that possibility. Thus, your assignment of confidence to the possibility it's Isaac can't be the assignment of confidence to the possibility you made a correct decision, because there hasn't yet been a decision of the right kind.

disagree about whether perceptual experiences distribute confidence over a finite number of possibilities, such as that it's Isaac or that it's not Isaac. At bottom, this is a disagreement about the processes responsible for our perceptual experiences, in particular the tasks they are trying to complete. According to Denison, they are just trying to complete estimation tasks, and as a result our perceptual experiences just assign degrees of confidence over continuous ranges of possibilities. According to me, they are also trying to complete categorization tasks, and as a result also assign degrees of confidence over finite numbers of possibilities.

Denison's argument is straightforward: In order to complete a categorization task, we first need to know what possibilities are relevant, and we don't know that until a question is asked. In our initial example, the question was whether it's Isaac or not. But there are many other questions that could have been asked, including whether the figure is Isaac or Aaron, tall or short, dark skinned or light skinned, and so on. Thus, we have to wait until a question is asked. According to Denison, these questions are never asked automatically, before our perceptual experiences. Instead, they are always asked later, after our perceptual experiences, and thus we can't try to complete a categorization task until after our perceptual experiences (p.60–61).[3]

Importantly, Denison's argument would establish that our perceptual experiences never include the possibility that it's Isaac, regardless of whether our perceptual experiences assign confidence to that possibility. Thus, in what follows, let's temporarily set aside the question of whether our perceptual experiences assign degrees of confidence, and just focus on whether our perceptual experiences involve automatic categorization.

To make our disagreement more concrete, let's start with several examples of categorization. Keep in mind that we're using 'categorization' to describe any decision between finitely many possibilities. The relevant possibilities needn't involve so-called high-level features, such as faces. The relevant decision needn't depend on representations that have the same format as the representations used in cognition. And the relevant decision needn't involve deliberation; it need only involve measurements and criteria. With this in mind, let's focus on the following examples: Recognizing Isaac involves categorization, because it involves the decision that it's Isaac, rather than someone else. A decision is required, because Isaac can produce many different proximal stimulations, depending on his distance, clothes, angle, and so on. Color constancy involves categorization, because it involves the decision that two objects are the same shade, rather than different

[3] This isn't Denison's only argument. She also argues: Categorization often fails when an object is in the periphery, unattended, poorly lit, in an atypical context, or seen from a non-canonical perspective. She thinks it follows that categorization occurs only if someone asks a question (p.61). But that doesn't follow. All that follows is that categorization is less reliable when the object is in the periphery, unattended, etc.

shades. A decision is required because objects with the same shade can be illuminated differently, as when one is in a shadow, and thus produce different proximal stimulations. Color grouping involves categorization, because it involves the decision that two objects have the same color type, or different color types (e.g., red). A decision is required because objects that have the same color type can have different shades, and thus produce different proximal stimulations (e.g., scarlet and crimson). Object tracking involves categorization because it involves the decision that it's the same object you perceived a moment ago, rather than a new object. A decision is required because the object might change location, color, and shape, and thus produce different proximal stimulations. Size constancy involves categorization because it involves the decision that two objects have the same size, rather than different sizes. A decision is required because objects that are the same size can take up different proportions of the visual field. Phoneme perception involves categorization because it involves the decision that a sound is a "ba," rather than a "pa." A decision is required because the acoustic properties of some "ba" sounds are more similar to the acoustical properties of "pa" sounds than other "ba" sounds. The detection of animals involves categorization because it involves the decision that an object is an animal, rather than not. A decision is required because different animals can produce many different proximal stimulations. Likewise for the detection of flowers faces, and other evolutionary significant categories.

These also seem to be examples of *automatic* categorization. In each case, categorization is extremely fast. There's disagreement about exactly how fast (100 ms? 200 ms?) (see e.g., Liu et al. 2009). But, in any case, there doesn't seem to be enough time for deliberation; categorization seems automatic. As Mandelbaum (forthcoming) points out, this is what we'd expect from an evolutionary point of view, because fast responses often make the difference between life and death. If categorizing a tiger as a tiger requires you to deliberately ask "Is that a tiger?" you might not survive very long.

Finally, these seem to be examples in which the relevant categorization is included in our *perceptual experiences*. Recall that we're using 'perceptual experiences' to describe states that are conscious, automatic, accessible, dissociable from doxastic states, directed toward perceived objects and properties, and fast enough that we can't detect any delay. Introspection alone seems like enough to convince us that recognition, color constancy, color grouping, object tracking, etc., are included in our perceptual experiences. Thus, introspection alone seems like enough to convince us that some questions are asked before our perceptual experiences, and the answers included in our perceptual experiences. Others use 'perceptual experience' with additional restrictions, and might be able to convince us that what they call perceptual experiences don't include *some* of these categories. For example, they might be able to convince us that we don't recognize

Isaac until after what they call perceptual experiences (e.g., Burge 2014). But, as long as these philosophers are talking about a state that's conscious, I doubt they'll be able to convince us that *none* of these categories are included in our perceptual experiences. It would take a long time to properly consider all the possible restrictions on 'perceptual experience' and how they apply to each of our examples. Fortunately, this is a cluster of issues that others have already written a lot about (e.g., Block 2014; Block draft; Cohen 2015; Hatfield 2012; Hilbert 2005; Mandelbaum forthcoming; O'Callaghan 2015).

It's worth noting that even Denison seems to think that some perceptual experiences involve automatic categorization. In particular, she writes about our perception of the Necker Cube that, "the presence of two peaks results in a selection process, such that only one interpretation is perceived at a time" (p.67). Thus, she seems to think that our perceptual experience of the Necker Cube is the result of an automatic decision between two possibilities.

Suppose I'm right, and that perceptual experiences involve automatic categorization. This would leave us with a number of questions. We'd like to know: Into what categories are objects automatically sorted? Just so-called low-level categories? When, if ever, are new categories introduced? When, if ever, are old categories eliminated? What impact, if any, do our beliefs have? We don't yet have complete answers to all these questions. Denison suggests that this casts doubt on PERCEPTUAL CONFIDENCE (p.60–61). But these questions arise for anyone who thinks that perceptual experiences involve automatic categorization, regardless of whether they prefer PERCEPTUAL CONFIDENCE or POST-PERCEPTUAL CONFIDENCE. And, for the reasons mentioned above, I take it that we have reason to think that perceptual experiences involve automatic categorization. In addition, I'm hopeful that we will eventually have complete answers to all these questions. Neuroscientists like Shadlen et al. (2008) are similarly optimistic: "perception, like decision making, arises by asking and answering questions that bear on specific propositions... What is it that establishes the question that the brain is asking about the data? What establishes the set of hypotheses?... We do not know the answers to these questions, but we suspect that they will turn out to be more tractable than the problems they replace" (p.72 and 93).

Moreover, suppose I'm right, not only that perceptual experiences involve automatic categorization, but also that categorization is sometimes graded, so that confidence is distributed over more than one possible categorization. This would leave us with another question. We'd also like to know: Why isn't categorization always graded, so that, for example, our perceptual experience of the Necker Cube assigns confidence to both interpretations? We don't yet have a complete answer to this question beyond that it presumably has something to do with the neural structures responsible for the relevant categorizations. Once again, Denison thinks that this casts doubt on

PERCEPTUAL CONFIDENCE (p.67). But this is just a special instance of a more general question that arises for everyone. In particular, we'd like to know, more generally, why the brain sometimes engages in yes–no categorization and other times engages in graded categorization. We don't yet have a complete answer to this question. Nonetheless, I take it that we have good reason to think that our brain is capable of both kinds of categorization, and thus should be hopeful that we'll eventually know the answer. I think we should have the same attitude toward the more specific question of why some of the processes responsible for our perceptual experiences engage in yes-no categorization, while others engage in graded categorization. This is especially true given that we seem to have examples of perceptual experiences that result from yes–no categorization, such as our perceptual experience of the Necker Cube, and examples of perceptual experiences that result from graded categorization, such as many of the examples I consider in the paper. Note that when our perceptual experiences result from yes–no categorization, and categorization occurs, I think that they assign full confidence to just one possibility (Morrison 2016, p.20).

There are many other questions that we'd like to answer, including the list of questions at the end of my paper (Morrison 2016, p.45–46). I regard these questions as opportunities for future research, rather than as reasons to give up on PERCEPTUAL CONFIDENCE.

Finally, keep in mind that PERCEPTUAL CONFIDENCE is true even if our perceptual experiences just involve automatic estimation, as Denison seems to think. Thus, our disagreement about automatic categorization is perhaps best understood as a disagreement about how PERCEPTUAL CONFIDENCE is implemented.

5. Blurriness

As far as I can tell, our only other substantive disagreement is about whether there's ordinarily a difference between a blurry perception of a sharp image, and a sharp perception of a blurry image. In the paper, I claim that there is ordinarily a difference, because when one has a blurry perception of a sharp image, the blurriness ordinarily seems to be a feature of one's relation to the image—a feature of one's perspective on the image—rather than a feature of the image itself (Morrison 2016, p.17 fn1). I think that this difference is reflected in the way one's perceptions assign confidence. When one has a blurry perception of a sharp image, one's perceptual experience ordinarily assigns confidence to the shapes with sharp boundaries that might be on the page, in Denison's example to differently sized dots with sharp boundaries. As a result, one might feel inclined to bring the image closer, or to refocus one's eyes. In contrast, when one has a sharp perception of a blurry image, one's perception assigns

confidence to different shapes, in Denison's example to a smudge with a roughly circular shape. As a result, one won't feel any inclination to move the image closer, or to refocus one's eyes. Thus, unlike Denison, I think that there's an intrinsic difference between these perceptions (p.66).

How can the visual system tell when the blurriness belongs to the perception, rather than to the image? Ordinarily, it relies on visual cues, such as depth cues, and whether anything in the background or periphery is in focus. But these cues are fallible, and it might be possible to trick someone into having the wrong kind of perception. For example, it might be possible to design a blurry image that fills the entire visual field, so that there are no visual cues indicating that one is perceiving a sharp image. In that case, one might have the same kind of perception as when one has a blurry perception of a sharp image. That is, one's perception might assign confidence to shapes with sharp boundaries. Importantly, this wouldn't establish that there's no difference between a blurry perception of a sharp image and a sharp perception of a blurry image. It would just establish that one is sometimes tricked into assigning the wrong confidences.

It might help to consider a related point about the difference between perceiving a horse and perceiving an image of a horse. A perception of a horse might represent the horse as 5 m away. An ordinary perception of an image of a horse, say in a magazine, will represent the image as much closer. How can the visual system tell whether it's perceiving a horse or an image of a horse? Ordinarily, it relies on visual cues, such as cues involving depth, size, and illumination. But it might still be possible to design an image of a horse that fills one's entire visual field, and that tricks one into perceiving a horse as 5 m away. Once again, this wouldn't establish that there's no difference between perceiving a horse and perceiving an image of a horse. It would just establish that one is sometimes tricked into representing the wrong distances.

6. Conclusion

In "Perceptual Confidence" I argued for PERCEPTUAL CONFIDENCE using first-personal evidence. When I wrote that paper, I wasn't aware of any third-personal evidence to help us choose between PERCEPTUAL CONFIDENCE and POST-PERCEPTUAL CONFIDENCE. But the situation has since changed, because there's now third-personal evidence that I think can be used to support PERCEPTUAL CONFIDENCE (in particular, van Bergen et al. 2015; Gherman and Philiastides 2015). I'm hopeful that more evidence will emerge in the near future. I'm not sure that this evidence will ever be enough to establish PERCEPTUAL CONFIDENCE by itself, without also taking into account first-personal evidence. But,

regardless, this is a great opportunity for collaboration between philosophers and neuroscientists.[4]

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.

van Bergen, Ruben, Wei Ji Ma, Michael S. Pratte, and Janneke F. M. Jehee. "Sensory Uncertainty Decoded from Visual Cortex Predicts Behavior." *Nature Neuroscience* 18 (2015): 1728–1730.

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

Block, Ned. "A Joint in Nature between Cognition and Perception without Modularity of Mind." November 27, 2016. Unpublished draft.

Burge, Tyler. "Reply to Block: Adaptation and the Upper Border of Perception." *Philosophy and Phenomenological Research* 89 (2014): 573–583.

Christensen, David. *Putting Logic in its Place: Formal Constraints on Rational Belief*. New York, NY: Oxford University Press, 2004.

Cohen, Jonathan. "Perceptual Constancy." *Oxford Handbook of the Philosophy of Perception*. Ed. Mohan, Matthen. New York, NY: Oxford University Press, 2015. 621–639.

Denison, Rachel. "Precision, not Confidence, Describes the Uncertainty of Perceptual Experience." *Analytic Philosophy* 58 (2017): 58–70.

Gherman, Sabina, and Marios G. Philiastides. "Neural Representations of Confidence Emerge from the Process of Decision Formation during Perceptual Choices." *NeuroImage* 106 (2015): 134–143.

Hájek, Alan. "Interpretations of Probability". *The Stanford Encyclopedia of Philosophy* (Winter 2012 Edition), Ed. Edward N. Zalta. 2012. https://plato.stanford.edu/archives/win2012/entries/probability-interpret/.

Hatfield, Gary. "Phenomenal and Cognitive Factors in Spatial Perception." *Visual Experience: Sensation, Cognition, and Constancy*. Eds. Gary Hatfield, and Sarah Allred. Oxford, UK: Oxford University Press, 2012. 35–62.

Hilbert, David R. "Color Constancy and the Complexity of Color." *Philosophical Topics* 33 (2005): 141–158.

Joyce, James M. *The Foundations of Causal Decision Theory*. New York, NY: Cambridge University Press, 1999.

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

Liu, Hesheng, Yigal Agam, Joseph R. Madsen, and Gabriel Kreiman. "Timing, Timing, Timing: Fast Decoding of Object Information from Intracranial Field Potentials in Human Visual Cortex." *Neuron* 62 (2009): 281–290.

Mandelbaum, Eric. "Seeing and Conceptualizing: Modularity and the Shallow Contents of Vision." *Philosophy and Phenomenological Research* DOI: 10.1111/phpr.12368.

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

O'Callaghan, Casey. "Speech Perception". *Oxford Handbook of the Philosophy of Perception*. Ed. Mohan Matthen. New York, NY: Oxford University Press, 2015. 475–494.

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

[4] Thanks to the members of Mike Shadlen's journal club, E. J. Green, Eric Mandelbaum, Jorge Morales, Jon Simon, and especially Chris Fetsch for comments on earlier drafts.

Schwitzgebel, Eric. "Belief". The Stanford Encyclopedia of Philosophy (Summer 2015 Edition). Ed. Edward N. Zalta. 2015. https://plato.stanford.edu/archives/sum2015/entries/belief/.
Michael, Shadlen N., Roozbeh Kiani, Timothy D. Hanks, and Anne K. Churchland. "Neurobiology of Decision Making: An Intentional Framework." *Better than Conscious? Decision Making, the Human Mind, and Implications for Institutions*. Eds. Christoph Engel, and Wolf Singer. Cambridge: MIT Press, 2008. 71–101.
Sigman, Mariano, and Zachary F. Mainen. "Confidence as Bayesian Probability: From Neural Origins to Behavior." *Neuron* 88 (2015): 78–92.
Steele, Katie, and H. Orri Stefánsson. "Decision Theory". The Stanford Encyclopedia of Philosophy (Winter 2016 Edition). Ed. Edward N. Zalta. 2016. https://plato.stanford.edu/archives/win2016/entries/decision-theory/.
Witt, Jessica K., J. Eric, T. Taylor, Mila Sugovic, and John T. Wixted. "Signal Detection Measures Cannot Distinguish Perceptual Biases from Response Biases." *Perception* 44 (2015): 289–300.