Uncertainty in perceptual representations

Dobromir Rahnev

Department of Psychology, Georgia Institute of Technology, Atlanta, GA

Correspondence

Dobromir Rahnev

Department of Psychology

Georgia Institute of Technology

130 J.S. Coon Building

Atlanta, GA 30332

drahnev@gmail.com

Keywords: Uncertainty, perception, sensation, modeling, confidence, cue combination

ABSTRACT

How are perceptual decisions made? The answer to this seemingly simple question necessitates that we specify the nature of perceptual representations on which decisions are based. Some traditional models postulate that the perceptual representation consists of a simple point estimate of the stimulus. Such models do not allow the estimation of sensory uncertainty. On the other hand, recent models have proposed that the perceptual representation involves a full probability distribution over the possible stimulus values. Such models allow a precise estimation of sensory uncertainty. These two possibilities – point estimates vs. full distributions – are often seen as the only alternatives but they are not. Here I present five possible perceptual representation schemes that allow the extraction of different levels of sensory uncertainty. I explain where popular models fall within the five schemes and explore the relevant empirical evidence and theoretical arguments. The overwhelming evidence is at odds with both point estimates vs. full distributions. This conclusion is in stark contrast with current popular models in computational neuroscience built on such distributions. Instead, the most likely scheme appears to be one in which the perceptual representation features a point estimate coupled with a strength-of-evidence value.

INTRODUCTION

Uncertainty in perception

Everybody in the audience is quiet. The tennis player has a match point. The play turns into a long rally from the baseline until her opponent hits a shot that seems to land just out. However, the line judge remains quiet implying that the ball was inside the court. The player can stop the play and challenge the call using the replay system. The problem is that she will automatically lose the point if the ball was indeed not out. What should she do?

Situations that push our perceptual abilities to the edge are common in tennis. They are also more common in our daily life than we realize. We make judgments on when to cross a busy road that lacks traffic lights, who is the person talking outside our office door, and whether the "weird" smell indicates that the cheese is just right or completely spoiled.

The nature of the perceptual representation

The first step in understanding how we make the perceptual decisions described above is to determine the nature of the perceptual representation. When the tennis player ponders whether to challenge the call, what information is she basing her decision on?

Some traditional theories imply that the player only forms a point estimate of the most likely landing location of the ball (**Figure 1, left**). In contrast, many recent theories argue for the existence of a full probability distribution over landing locations (**Figure 1, right**). Deciding between these competing possibilities is not only fundamental for understanding the nature of our perceptual representation but would necessarily falsify a large number of popular theories that imply the wrong representation.

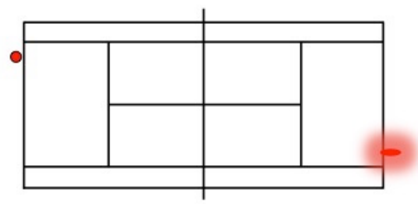

Figure 1. Perceptual representation of ball landing location in tennis.

Two extreme proposals for the nature of perceptual representations exist: the representation may consist of a simple point estimate (represented on the left as a ball landing just outside the court) or a full probability distribution (represented on the right as a ball landing within a certain area of the court).

A false dichotomy

The distinction between point estimates and full probability distributions represents a false dichotomy: there are a number of "intermediate" possibilities about the nature of the perceptual representation.

Here I present five different schemes for the nature of the perceptual representation. Point estimates and full distributions feature as the extreme possibilities. Further, I show that both of these extremes are strongly contradicted by the evidence. These considerations upend some of the most popular theories of perceptual decision making.

FIVE SCHEMES FOR THE NATURE OF THE PERCEPTUAL REPRESENTATION

The perceptual representation of a given stimulus is different in different stages of the visual hierarchy. Here we are interested in the perceptual representation that the brain uses at the decisional level (Pouget, Beck, Ma, & Latham, 2013). Even if full distributions *could* be extracted from the population response in V1 or MT, it does not follow that this information is in fact available at the decision stage.

The five schemes presented below are arranged in order of increasing complexity. To compare between them, I use the example of representing the direction of a single moving bar (**Figure 2A**). For simplicity, I will assume that the moving bar activates to a different degree a set of 36 neurons each having its receptive field centered 10° away from the previous one (**Figure 2B**). The question is what part of this population code is available for making the final decision.

Scheme 1. Single point estimate (no sensory uncertainty)

A first possibility is that the perceptual representation consists of a single point estimate (**Figure 2C**). For example, the direction a moving bar would be represented as a single orientation (e.g., 50°). In our neural example, the point estimate could be based on which of the 36 neurons has the highest activity.

However, more complex scenarios are also possible. For example, it is possible that maximum likelihood estimation (MLE) or another method is used to estimate the most likely direction of motion by taking into account all 36 neurons. The critical point, however, is that only the single value is passed onto the decision stage and therefore this first scheme does not allow the estimation of sensory uncertainty.

A popular model that falls within this scheme is the drift diffusion model (Ratcliff & McKoon, 2008). In its most common variant, a particle diffuses towards one of two boundaries representing each choice and the decision is made when the particle reaches one of the boundaries (**Figure 3A**). At the end of the process, the only available information is which of the two boundaries was reached – a single point estimate. Other examples from this category are low- and high-threshold models which postulate that stimuli give rise to a limited number – typically 2 – internal states (Krantz, 1969).

Scheme 2. Multiple point estimates (indirect sensory uncertainty)

The second possibility is that the perceptual representation consists of multiple point estimates (**Figure 2D**). For example, in estimating the direction of motion, people may additionally estimate the contrast of the bar, as well as their own decision time and attentional state. These additional estimates may provide clues about the reliability of the main point estimate. Therefore, this second scheme allows only for an indirect estimation of the sensory uncertainty. The drift diffusion model could also fall within this scheme since on each trial it explicitly represents the decision time (**Figure 3A**), which can then be used to judge the task difficulty (Kiani, Corthell, & Shadlen, 2014).

Scheme 3. Strength of evidence (minimal sensory uncertainty)

The third possibility is that the perceptual representation consists of a point estimate complemented by a "strength-of-evidence" value (**Figure 2E**). For example, the direction of a moving bar would be represented with its mostly likely value (e.g., 50°) together with a judgment about the strength of evidence for this value. In our neural example, the motion representation may be based on which of the 36 neurons has the highest activity, as well as on the actual level of this activity. As in Scheme 1, more complex scenarios based on MLE or similar computations that take into account all neurons are also possible. This third scheme allows for what could be called "minimal" estimation of sensory uncertainty.

A popular model that falls within this scheme is signal detection theory (Green & Swets, 1966; Macmillan & Creelman, 2005). Signal detection theory assumes the existence of a single "evidence" axis in 2-choice tasks. On this axis, each trial produces a single point that represents the strength of evidence (**Figure 3B**). It should be noted that several variants of the drift diffusion model also fall under Scheme 3. For example, models with separate accumulators (Usher & McClelland, 2001; Vickers, 1970) can produce a strength-of-evidence value based on the comparison of the winning and losing accumulators.

Scheme 4. Partial distribution (partial sensory uncertainty)

The fourth possibility is that the perceptual representation consists of several moments of the sensory distribution (**Figure 2F**). For example, extracting the first four moments of 36-neuron distribution in our example will mean that the mean, variance, skewness, and kurtosis are explicitly represented. Other summary statistics such as the median are also possible. For simplicity, I will equate Scheme 4 with representing the mean and standard deviation. On a single trial, the direction of a moving bar may therefore be represented as $50° \pm 15°$. Therefore, this fourth scheme results in partial estimation of sensory uncertainty but still carries more information than Schemes 1-3. No currently popular model of perceptual decision making falls under this scheme. Nevertheless, the models discussed in Scheme 5 could be reduced to fit into this category (Ma, 2010).

Scheme 5. Full distribution (complete sensory uncertainty)

The final possibility is that the perceptual representation consists of a full probability distribution (**Figure 2G**). Thus, the direction of the moving bar is not summarized but is represented by a complete distribution and therefore allows the complete estimation of sensory uncertainty. This scheme is the only one that will work well in complex situations such as skewed or bimodal distributions. It is also the only scheme that allows for fully optimal decisions on every trial. These features have made this scheme very popular among computational neuroscientists (Beck, Ma, Pitkow, Latham, & Pouget, 2012; Berkes, Orbán, Lengyel, & Fiser, 2011; Drugowitsch & Pouget, 2012; Fiser, Berkes, Orbán, & Lengyel, 2010; Jazayeri &

Movshon, 2006; Knill & Pouget, 2004; Ma, 2010, 2012; Ma, Beck, Latham, & Pouget, 2006; Ma & Jazayeri, 2014; Pouget et al., 2013; Pouget, Dayan, & Zemel, 2000; Sahani & Dayan, 2003; Zemel, Dayan, & Pouget, 1998).

Two popular models that fall within this scheme are the probabilistic population codes (Ma et al., 2006) and neural sampling with a large number of samples (Fiser et al., 2010). Models based on probabilistic population codes propose that operations like cue combination can be performed using the whole 36-neuron distribution (**Figure 3C**). Indeed, under certain assumptions, simply adding the distributions produced by each cue results in optimal cue combination. Neural sampling models (**Figure 3D**), on the other hand, propose that neurons take discrete samples from the stimulus in small time intervals. These samples are then combined into a full distribution, provided that enough samples can be taken. Note that taking a single sample results in a Scheme 1 representation, while taking only a few samples may be more similar to a Scheme 3 representation.

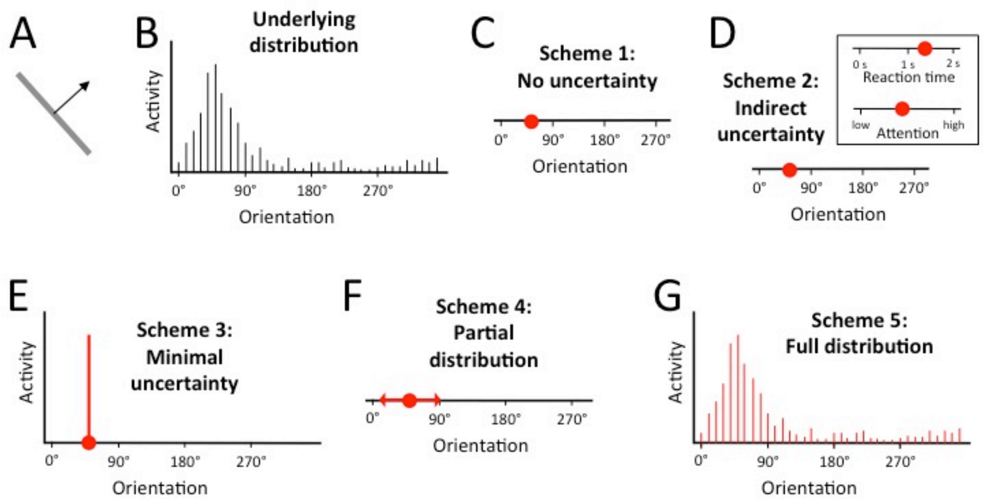

Figure 2. Five schemes for the nature of the perceptual representation.

A. A person is judging the direction of motion of a bar. B. The moving bar creates a distribution in motion-sensitive neurons tuned to different orientations. In this example, 36 neurons are used with preferred directions tuned in multiples of 10°. C. Scheme 1: Single point estimate. Only a single point estimate is extracted. D. Scheme 2: Multiple point estimates. Several point estimates are extracted for variables relevant to the task (e.g., the inset shows estimates of decision time and attentional state). E. Strength of evidence. A point estimate and a strength-of-evidence value are extracted. F. Partial probability distribution. The mean and standard deviation of the distribution of neuron activity are extracted. G. Full probability distribution. The whole distribution of neuron activity is used.

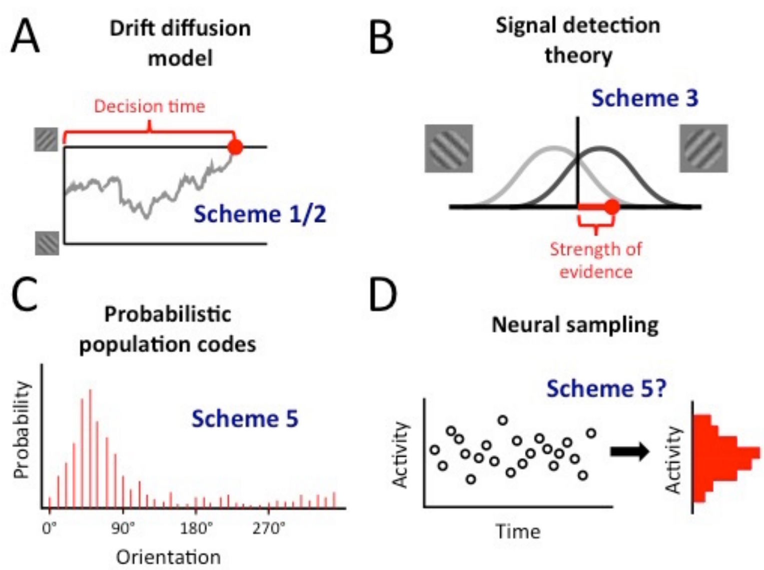

Figure 3. Four popular models of perceptual decision making.

A. Drift diffusion models assume that decisions are made by a particle diffusing towards one of two boundaries, representing the possible stimulus categories (in this case, clockwise and counterclockwise oriented gratings). The decision is based on which boundary was crossed first (Scheme 1). Decisions can also take into account the decision time (Scheme 2). B. Signal detection theory assumes that decisions are made based on which side of a criterion the evidence on a particular trial falls. The distance from the criterion can be used as an estimate of the strength of evidence (Scheme 3). C. Probabilistic population codes assume the existence of a complete distribution built by a population of neurons (Scheme 5). D. Neural sampling models assume that a single neuron takes samples over time. Sufficient number of samples results in a complete distribution (Scheme 5). However, a single sample results in a point estimate (Scheme 1), while a situation with a small number of samples may be best described by Scheme 3.

EMPIRICAL EVIDENCE

As already briefly touched upon, each of the five schemes comes with strong implications about the nature of perceptual decision making. Here I make these implications explicit and compare them with the available empirical evidence. I specifically discuss findings related to confidence ratings and cue combination tasks (**Table 1**). The reason for focusing on these two phenomena is that they directly test people's ability to estimate the uncertainty in their perceptual representations. Confidence ratings reflect our uncertainty about perceptual judgments for a single stimulus, while cue combination studies test how the uncertainty from two different sensory signals is compared and combined.

Confidence

In virtually every perceptual task subjects automatically and naturally produce confidence ratings (Metcalfe & Shimamura, 1994). These confidence ratings reflect the likelihood of being correct (Baranski & Petrusic, 1994; Fleming & Lau, 2014). This ability to naturally produce meaningful confidence ratings is strong evidence against Scheme 1. Indeed, this scheme does not feature any information on which appropriate confidence judgments could be based. Scheme 2 also has difficulties but partly meaningful confidence is still possible to the extent to which additional parameters, such as decision time, correlate with performance. Schemes 3-5 easily explain meaningful confidence ratings: higher confidence is associated with higher strength of evidence (Scheme 3), lower standard deviations (Scheme 4), or narrower distributions (Scheme 5).

On the other hand, a number of suboptimalities exist in confidence ratings (Rahnev & Denison, 2016). For example, subjects are typically either under- or over-confident and are rarely able to calibrate their confidence properly (Adams, 1957; Baranski & Petrusic, 1994; Björkman, Juslin, & Winman, 1993; Dawes, 1980; Harvey, 1997; Keren, 1988; Koriat, 2011; Winman & Juslin, 1993). Even more importantly, a large number of studies have reported conditions matched on accuracy that produce different levels of confidence (Baranski & Petrusic, 1994; de Gardelle & Mamassian, 2015; Kiani et al., 2014; Koizumi, Maniscalco, & Lau, 2015; Navajas, Sigman, & Kamienkowski, 2014; Rahnev et al., 2011; Rahnev, Bahdo, de Lange, & Lau, 2012; Rahnev, Koizumi, McCurdy, D'Esposito, & Lau, 2015; Samaha, Barrett, Sheldon, LaRocque, & Postle, 2016; Song, Koizumi, & Lau, 2015; Spence, Dux, & Arnold, 2015; Vickers & Packer, 1982; Vlassova, Donkin, & Pearson, 2014; Wilimzig, Tsuchiya, Fahle, Einhäuser, & Koch, 2008; Zylberberg, Roelfsema, & Sigman, 2014). Finally, the trial-to-trial relationship between confidence and accuracy, while usually positive, rarely reaches its theoretical maximum (Maniscalco & Lau, 2015; Maniscalco, Peters, & Lau, 2016; Massoni, 2014; McCurdy et al., 2013; Schurger, Kim, & Cohen, 2015; Sherman, Seth, Barrett, & Kanai, 2015; Vlassova et al., 2014). These findings strongly suggest the existence of heuristics in confidence computation. This conclusion is at odds with Scheme 5 since the presence of complete probability distributions should allow for fully Bayesian, rather than heuristic, computations. The extent to which

Scheme 4 is consistent with these findings depends on whether subjects can transform the standard deviation of the underlying distribution into probability of being correct. This computation is challenging for humans (Zhang & Maloney, 2012) and may explain some of the biases above. On the other hand, Schemes 2-3 fit well with these findings of suboptimal confidence ratings. Indeed, the probability of being correct is likely a highly non-linear function of the strength of evidence in Scheme 3 and is only imperfectly correlated with parameters such as decision time in Scheme 2. These considerations explain why confidence biases are to be expected in Schemes 2-3. Since Scheme 1 cannot be used to produce confidence ratings, it also cannot explain any of the above biases.

	Finding	**Scheme 1**	**Scheme 2**	**Scheme 3**	**Scheme 4**	**Scheme 5**
Confidence	Confidence "meaningful" in most (all?) tasks	No	Maybe	Yes	Yes	Yes
	Confidence biases are ubiquitous	No	Yes	Yes	Maybe	No
Cue combination	Cue combination near optimal in many tasks	No	Maybe	Yes	Yes	Yes
	Cue combination clearly suboptimal in some tasks	Yes	Yes	Yes	No	No

Table 1*. Can Schemes 1-5 account for various findings related to confidence and cue combination?*

Cue combination

Cue combination is needed when two or more pieces of information need to be combined to form a single decision (Trommershäuser, Körding, & Landy, 2011). For example, the length of a bar could be estimated based on a combination of visual and haptic (touch) information (Ernst & Banks, 2002). When the information from each of these sensory modalities is noisy, the evidence from each is combined in order to arrive at a better estimate than either sense can afford by itself.

Cue combination studies have often found near optimal integration (Alais & Burr, 2004; Ernst & Banks, 2002; Gu, Angelaki, & DeAngelis, 2008; van Beers, Sittig, & Denier van der Gon, 1996). Such findings have been cited as providing the strongest support for the existence of a full probability distribution as in Scheme 5 (Beck et al., 2012; Berkes et al., 2011; Drugowitsch & Pouget, 2012; Fiser et al., 2010; Knill & Pouget, 2004; Ma, 2010, 2012; Ma et al., 2006; Ma & Jazayeri, 2014; Pouget et al., 2013). Indeed, the existence of a full probability distribution easily explains optimal cue combination (Ma et al., 2006). Importantly, Scheme 4 also naturally fits with optimal cue combination since such combination only requires the representation of the distributions' mean and standard deviation.

What is less appreciated is that Scheme 3, and to a lesser extent even by Scheme 2, can also explain near optimal performance in cue combination studies. Scheme 3 requires subjects to weight each stimulus' point estimate by the strength-of-evidence value associated with it. In many cases this strategy would result in near optimal performance. Scheme 2 has more difficulties with near optimal cue combination because of its indirect representation of sensory uncertainty. Still, occasional near optimal performance is possible when parameters such as decision time or attentional state are strongly correlated with performance. Thus only Scheme 1 is completely inconsistent with near optimal cue combination.

On the other hand, discussions of cue combination studies often ignore the fact that many such studies have found substantial suboptimalities (Battaglia, Jacobs, & Aslin, 2003; Battaglia, Kersten, & Schrater, 2011; Burr, Banks, & Morrone, 2009; Fetsch, Pouget, Deangelis, & Angelaki, 2012; Knill & Saunders, 2003; Maiworm & Röder, 2011; Prsa, Gale, & Blanke, 2012; Rosas, Wagemans, Ernst, & Wichmann, 2005; Rosas, Wichmann, & Wagemans, 2007) (reviewed in (Rahnev & Denison, 2016)). These studies typically report that one of the cues was weighted more than its reliability, relative to the other cue. Such findings are extremely surprising if indeed the brain represents full probability distributions (Scheme 5) or has direct access to the stimulus reliability through the standard deviation of the distribution (Scheme 4). Therefore, these findings argue strongly against Schemes 4 and 5. On the other hand, Schemes 2-3 can easily explain findings of suboptimal cue combination just as they could explain biases in confidence ratings. Scheme 1 would result in random weighting of the cues, which will almost always be suboptimal.

Neural evidence

Little neural evidence has been used to distinguish between schemes for perceptual representation. A notable exception is a recent paper by van Bergen et al. (van Bergen, Ma, Pratte, & Jehee, 2015). The authors showed that the degree of sensory uncertainty could be directly decoded from activity in the visual cortex. van Bergen et al. interpreted these results as consistent with the existence of full probability distributions (Scheme 5). However, these results do not directly show what information subjects actually *used* when making their decisions. In fact, the results can easily be explained by any of Schemes 2-5. Nonetheless, van Bergen et al.'s findings provide strong evidence against Scheme 1. More generally, the considerations above demonstrate the difficulty of using neural data to distinguish between Schemes 2-5.

THEORETICAL ARGUMENTS AGAINST FULL DISTRIBUTIONS

As noted above, full probability distributions (as in Scheme 5) have been accepted among computational neuroscientists almost to the exclusion of other alternatives (Beck et al., 2012; Berkes et al., 2011; Drugowitsch & Pouget, 2012; Fiser et al., 2010; Jazayeri & Movshon, 2006; Knill & Pouget, 2004; Ma, 2010, 2012; Ma et al., 2006; Ma & Jazayeri, 2014; Pouget et al., 2013, 2000; Sahani & Dayan, 2003; Zemel et al., 1998). There seem to be three main reasons for this acceptance. First, computational neuroscientists often prefer to build normative models of how the visual system should or could deal with uncertainty. Full distributions are best for normative computations but there is no a priori reason to expect that the brain implements normative solutions, especially in complex situations (Gigerenzer & Brighton, 2009; Juslin, Nilsson, & Winman, 2009; Simon, 1956). Second, discussions of empirical findings typically focus on findings of optimality, while an extensive literature on suboptimalities in perception (Rahnev & Denison, 2016) is often ignored. Finally, Scheme 1 and 5 are typically considered as the only alternatives. In this view, any finding that the brain can extract uncertainty estimates from sensory representations is taken as evidence for full probability distributions. As demonstrated by Schemes 2-4, many other options exist. The issue is further complicated by the fact that terms such as "Bayesian" and "probabilistic" have are often used with different meanings (**Box 1**). Full probability distributions are further contradicted by several theoretical considerations discussed below.

Real-world tasks are exceedingly complex

Real-life perception comes with an explosion in computational complexity. Such complexity virtually guarantees that decisions will be based on heuristics rather than fully principled computations (Gigerenzer & Brighton, 2009; Juslin et al., 2009; Simon, 1956). Even supporters of full probability distributions admit that complex situations call for simplified computations (Pouget et al., 2013). Perception evolved to serve us in real life rather than in the laboratory. Thus complex conditions are the norm rather than the exception for the brain. If heuristics are necessary anyway, then perceptual representations that allow the estimation of more than minimal sensory uncertainty (as in Schemes 4 and 5) and are mostly applicable in very simple situations may be an unnecessary luxury.

Discrete judgments

The example used to introduce Schemes 1-5 used a continuous quantity: the motion direction of a bar. However, many decisions involve discrete judgments: who is this person in your high school reunion or what species is the bird on the distant tree? It is unclear how full probability distributions can be constructed and meaningfully used in such situations, especially since relevant possibilities may not be available (e.g., we may have forgotten about Ricardo or not know about the existence of brown-headed cowbirds). Scheme 4 would be completely impossible in such situations since discrete judgments do not allow the computation of a mean and a standard deviation.

Improvements with practice

Perceptual judgments are known to become more optimal with practice (Balci et al., 2011; Baranski & Petrusic, 1994; Maddox & Bohil, 2005). Such findings fit well with Schemes 2 and 3. The reason is that in both of these schemes, the additional quantities (beyond the point estimate) such as decision time and strength of evidence may predict one's accuracy differently for different tasks. Thus, both Schemes 2 and 3 require learning to calibrate how these additional quantities should be used. However, if full probability distributions are present on every trial, it is less clear how and why learning would make judgments more optimal.

Box 1. Are perceptual decisions Bayesian? Are they probabilistic?

The notion that perceptual representations at the decision stage do not consist of full probability distributions has relevance to theories about "Bayesian computation," "probabilistic computation," "probabilistic approach," and "probabilistic brains" (Beck et al., 2012; Drugowitsch & Pouget, 2012; Knill & Pouget, 2004; Ma, 2010, 2012; Ma et al., 2006; Ma & Jazayeri, 2014; Pouget et al., 2013). These concepts sound similar but are not necessarily synonymous. Here I explore the implications of rejecting Scheme 5 for these theories.

Are perceptual decisions Bayesian? Decisions are Bayesian as long as they follow Bayes' rule. Much evidence suggests that they do (Rahnev & Denison, 2016). Importantly, many Bayesian models require a single point estimate per stimulus and build the Bayesian machinery around inferring how the point estimates vary over trials. Thus all five schemes from Figure 2 and all four models from Figure 3 are fully consistent with the notion that our perceptual decisions are Bayesian.

Are perceptual decisions probabilistic? The term "probabilistic" is more challenging since it has been used in different ways. Ma (Ma, 2010) distinguishes between *probabilistic models* (in which trial-to-trial observations are stochastic; such models can feature representations consistent with all five schemes) and *models of probabilistic computation* (which require the representation of at least two moments of the sensory distribution; such models are only consistent with Schemes

4 and 5). Nevertheless, Ma's terminology is not widely used. Other papers equate phrases such as the "probabilistic approach" (Pouget et al., 2013) and representing stimuli in a "probabilistic manner" (Fiser et al., 2010) with the existence of the full probability distributions from Scheme 5. Thus, rejecting Scheme 5 means rejecting the notion of probabilistic decisions in some but not other definitions of the term.

Referring to Schemes 1-5 should clarify the exact concept researchers seek to advance. It may also help avoid common logical traps such as stating that the brain computes in a Bayesian manner (true) and concluding that full probability distributions are necessarily needed (false).

WHICH IS THE CORRECT SCHEME?

The empirical evidence is strongly against Scheme 5, which posits full probability distributions. This conclusion is in stark contrast to most of the recent work on computational models of perception that have assumed the presence of complete distributions (Beck et al., 2012; Berkes et al., 2011; Drugowitsch & Pouget, 2012; Fiser et al., 2010; Jazayeri & Movshon, 2006; Knill & Pouget, 2004; Ma, 2010, 2012; Ma et al., 2006; Ma & Jazayeri, 2014; Pouget et al., 2013, 2000; Sahani & Dayan, 2003; Zemel et al., 1998). This conclusion questions the plausibility of probabilistic population codes (Ma et al., 2006) and neural sampling with a large number of samples (Fiser et al., 2010). Scheme 4 – which proposes that the brain computes partial distribution information such as mean and standard deviation – is also unlikely, especially in the light of the numerous examples of suboptimal cue combination (**Table 1**).

On the other hand, the empirical evidence also strongly contradicts Scheme 1, which posits a single point estimate. This conclusion questions the plausibility of the simplest version of the drift diffusion model where the subject only knows which boundary has been crossed (Ratcliff & McKoon, 2008). Scheme 2 – which only allows an indirect estimate of sensory uncertainty – has trouble with findings of meaningful confidence or optimal cue combination. Nevertheless, it is possible the brain computes sufficient number of related parameters to achieve high performance in these tasks. Still, potential proponents of Scheme 2 have the tall task of determining what these parameters are and demonstrating that they are sufficient for near optimal behavior in a number of tasks.

The evidence as a whole therefore provides greatest support for Scheme 3 (**Table 1**). It should be noted that it is possible to extend Scheme 3 so that different related parameters, such as decision time and attention state, are also estimated (as in Scheme 2). These parameters would still contribute to the decision and can each feature a strength-of-evidence value. Clearly, models that feature this type of uncertainty need to specify explicitly how the strength of evidence is computed and how people learn to relate it to the probability of being correct. Still, the ability of Scheme 3 to naturally account for both findings of optimality and suboptimality makes it a particularly promising model of perceptual decision making.

CONCLUSION

The nature of the perceptual representation at the decision level is still a mystery. Proposals vary from single point estimates to full probability distributions. This paper presents five possible schemes for the nature of the perceptual representation. Evidence from confidence and cue combination studies is at odds with both single point estimates and full distributions. The latter are the preferred schemes in most popular models in computational neuroscience. Instead, the most likely scheme features a point estimate accompanied by a strength-of-evidence value.

REFERENCES

Adams, J. K. (1957). A confidence scale defined in terms of expected percentages. *The American Journal of Psychology*, *70*(3), 432–6.

Alais, D., & Burr, D. (2004). The ventriloquist effect results from near-optimal bimodal integration. *Current Biology*, *14*(3), 257–62. http://doi.org/10.1016/j.cub.2004.01.029

Balci, F., Simen, P., Niyogi, R., Saxe, A., Hughes, J. A., Holmes, P., & Cohen, J. D. (2011). Acquisition of decision making criteria: reward rate ultimately beats accuracy. *Attention, Perception & Psychophysics*, *73*(2), 640–57. http://doi.org/10.3758/s13414-010-0049-7

Baranski, J. V, & Petrusic, W. M. (1994). The calibration and resolution of confidence in perceptual judgments. *Perception & Psychophysics*, *55*(4), 412–28. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/8036121

Battaglia, P. W., Jacobs, R. A., & Aslin, R. N. (2003). Bayesian integration of visual and auditory signals for spatial localization. *Journal of the Optical Society of America A, Optics and Image Science*, *20*(7), 1391–1397.

Battaglia, P. W., Kersten, D., & Schrater, P. R. (2011). How haptic size sensations improve distance perception. *PLoS Computational Biology*, *7*(6), e1002080. http://doi.org/10.1371/journal.pcbi.1002080

Beck, J. M., Ma, W. J., Pitkow, X., Latham, P. E., & Pouget, A. (2012). Not noisy, just wrong: the role of suboptimal inference in behavioral variability. *Neuron*, *74*(1), 30–9. http://doi.org/10.1016/j.neuron.2012.03.016

Berkes, P., Orbán, G., Lengyel, M., & Fiser, J. (2011). Spontaneous cortical activity reveals hallmarks of an optimal internal model of the environment. *Science*, *331*(6013), 83–7. http://doi.org/10.1126/science.1195870

Björkman, M., Juslin, P., & Winman, A. (1993). Realism of confidence in sensory discrimination: the underconfidence phenomenon. *Perception & Psychophysics*, *54*(1), 75–81. Retrieved from http://www.ncbi.nlm.nih.gov/pubmed/8351190

Burr, D., Banks, M. S., & Morrone, M. C. (2009). Auditory dominance over vision in the perception of interval duration. *Experimental Brain Research*, *198*(1), 49–57. http://doi.org/10.1007/s00221-009-1933-z

Dawes, R. M. (1980). Confidence in intellectual vs. confidence in perceptual judgments. In E. D. Lantermann & H. Feger (Eds.), *Similarity and choice: Papers in honor of Clyde Coombs* (pp. 327–345). Bern: Han Huber.

de Gardelle, V., & Mamassian, P. (2015). Weighting Mean and Variability during Confidence Judgments. *PloS One*, *10*(3), e0120870. http://doi.org/10.1371/journal.pone.0120870

Drugowitsch, J., & Pouget, A. (2012). Probabilistic vs. non-probabilistic approaches to the neurobiology of perceptual decision-making. *Current Opinion in Neurobiology*, *22*(6), 963–9. http://doi.org/10.1016/j.conb.2012.07.007

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, *415*(6870), 429–33. http://doi.org/10.1038/415429a

Fetsch, C. R., Pouget, A., Deangelis, G. C., & Angelaki, D. E. (2012). Neural correlates of reliability-based cue weighting during multisensory integration. *Nature Neuroscience*, *15*(1), 146–154. http://doi.org/10.1038/nn.2983

Fiser, J., Berkes, P., Orbán, G., & Lengyel, M. (2010). Statistically optimal perception and learning: from behavior to neural representations. *Trends in Cognitive Sciences*, *14*(3), 119–30. http://doi.org/10.1016/j.tics.2010.01.003

Fleming, S. M., & Lau, H. (2014). How to measure metacognition. *Frontiers in Human Neuroscience*, *8*. http://doi.org/10.3389/fnhum.2014.00443

Gigerenzer, G., & Brighton, H. (2009). Homo Heuristicus: Why Biased Minds Make Better Inferences. *Topics in Cognitive Science*, *1*(1), 107–143. http://doi.org/10.1111/j.1756-8765.2008.01006.x

Green, D. M., & Swets, J. A. (1966). *Signal detection theory and psychophysics*. New York: John Wiley & Sons Ltd.

Gu, Y., Angelaki, D. E., & DeAngelis, G. C. (2008). Neural correlates of multisensory cue integration in macaque MSTd. *Nature Neuroscience*, *11*(10), 1201–1210. http://doi.org/10.1038/nn.2191

Harvey, N. (1997). Confidence in judgment. *Trends in Cognitive Sciences*, *1*(2), 78–82. http://doi.org/10.1016/S1364-6613(97)01014-0

Jazayeri, M., & Movshon, J. A. (2006). Optimal representation of sensory information by neural populations. *Nature Neuroscience*, *9*(5), 690–6. http://doi.org/10.1038/nn1691

Juslin, P., Nilsson, H., & Winman, A. (2009). Probability theory, not the very guide of life. *Psychological Review*, *116*(4), 856–874. http://doi.org/10.1037/a0016979

Keren, G. (1988). On the ability of monitoring non-veridical perceptions and uncertain knowledge: Some calibration studies. *Acta Psychologica*, *67*(2), 95–119. http://doi.org/10.1016/0001-6918(88)90007-8

Kiani, R., Corthell, L., & Shadlen, M. N. (2014). Choice Certainty Is Informed by Both Evidence and Decision Time. *Neuron*, *84*(6), 1329–1342. http://doi.org/10.1016/j.neuron.2014.12.015

Knill, D. C., & Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends in Neurosciences*, *27*(12), 712–9. http://doi.org/10.1016/j.tins.2004.10.007

Knill, D. C., & Saunders, J. A. (2003). Do humans optimally integrate stereo and texture information for judgments of surface slant? *Vision Research*, *43*(24), 2539–2558.

Koizumi, A., Maniscalco, B., & Lau, H. (2015). Does perceptual confidence facilitate cognitive control? *Attention, Perception & Psychophysics*. http://doi.org/10.3758/s13414-015-0843-3

Koriat, A. (2011). Subjective confidence in perceptual judgments: a test of the self-consistency model. *Journal of Experimental Psychology: General*, *140*(1), 117–139. http://doi.org/10.1037/a0022171

Krantz, D. H. (1969). Threshold theories of signal detection. *Psychological Review*, *76*(3), 308–324. http://doi.org/10.1037/h0027238

Ma, W. J. (2010). Signal detection theory, uncertainty, and Poisson-like population codes. *Vision Research*, *50*(22), 2308–2319. http://doi.org/10.1016/j.visres.2010.08.035

Ma, W. J. (2012). Organizing probabilistic models of perception. *Trends in Cognitive Sciences*, *16*(10), 511–8. http://doi.org/10.1016/j.tics.2012.08.010

Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*, *9*(11), 1432–8. http://doi.org/10.1038/nn1790

Ma, W. J., & Jazayeri, M. (2014). Neural Coding of Uncertainty and Probability. *Annual Review of Neuroscience*, *37*, 205–220. http://doi.org/10.1146/annurev-neuro-071013-014017

Macmillan, N. A., & Creelman, C. D. (2005). *Detection Theory: A User's Guide* (2nd ed.). Mahwah, NJ: Erlbaum.

Maddox, W. T., & Bohil, C. J. (2005). Optimal classifier feedback improves cost-benefit but not base-rate decision criterion learning in perceptual categorization. *Memory & Cognition*, *33*(2), 303–19.

Maiworm, M., & Röder, B. (2011). Suboptimal Auditory Dominance in Audiovisual Integration of Temporal Cues. *Tsinghua Science & Technology*.

Maniscalco, B., & Lau, H. (2015). Manipulation of working memory contents selectively impairs metacognitive sensitivity in a concurrent visual discrimination task. *Neuroscience of Consciousness*, *2015*(1), niv002. http://doi.org/10.1093/nc/niv002

Maniscalco, B., Peters, M. A. K., & Lau, H. (2016). Heuristic use of perceptual evidence leads to dissociation between performance and metacognitive sensitivity. *Attention, Perception & Psychophysics*, *78*(3), 923–37. http://doi.org/10.3758/s13414-016-1059-x

Massoni, S. (2014). Emotion as a boost to metacognition: How worry enhances the quality of confidence. *Consciousness and Cognition*, *29*, 189–198. http://doi.org/10.1016/j.concog.2014.08.006

McCurdy, L. Y., Maniscalco, B., Metcalfe, J., Liu, K. Y., de Lange, F. P., & Lau, H. (2013). Anatomical Coupling between Distinct Metacognitive Systems for Memory and Visual Perception. *The Journal of Neuroscience*, *33*(5), 1897–906. http://doi.org/10.1523/JNEUROSCI.1890-12.2013

Metcalfe, J., & Shimamura, A. P. (1994). *Metacognition: Knowing about Knowing*. Cambridge, MA: MIT Press.

Navajas, J., Sigman, M., & Kamienkowski, J. E. (2014). Dynamics of visibility, confidence, and choice during eye movements. *Journal of Experimental Psychology: Human Perception and Performance*, *40*(3), 1213–1227. http://doi.org/10.1037/a0036321

Pouget, A., Beck, J. M., Ma, W. J., & Latham, P. E. (2013). Probabilistic brains: knowns and unknowns. *Nature Neuroscience*, *16*(9), 1170–8. http://doi.org/10.1038/nn.3495

Pouget, A., Dayan, P., & Zemel, R. (2000). Information processing with population codes. *Nature Reviews Neuroscience*, *1*(2), 125–132. http://doi.org/10.1038/35039062

Prsa, M., Gale, S., & Blanke, O. (2012). Self-motion leads to mandatory cue fusion across sensory modalities. *Journal of Neurophysiology*, *108*(8), 2282–2291. http://doi.org/10.1152/jn.00439.2012

Rahnev, D., Bahdo, L., de Lange, F. P., & Lau, H. (2012). Prestimulus hemodynamic activity in dorsal attention network is negatively associated with decision confidence in visual perception. *Journal of Neurophysiology*, *108*(5), 1529–36. http://doi.org/10.1152/jn.00184.2012

Rahnev, D., & Denison, R. (2016). Suboptimality in perception. *bioRxiv*. Retrieved from http://biorxiv.org/content/early/2016/06/22/060194.abstract

Rahnev, D., Koizumi, A., McCurdy, L. Y., D'Esposito, M., & Lau, H. (2015). Confidence Leak in Perceptual Decision Making. *Psychological Science*, *26*(11), 1664–1680. http://doi.org/10.1177/0956797615595037

Rahnev, D., Maniscalco, B., Graves, T., Huang, E., De Lange, F. P., & Lau, H. (2011). Attention induces conservative subjective biases in visual perception. *Nature Neuroscience*, *14*(12), 1513–1515. http://doi.org/10.1038/nn.2948

Ratcliff, R., & McKoon, G. (2008). The diffusion decision model: theory and data for two-choice decision tasks. *Neural Computation*, *20*(4), 873–922. http://doi.org/10.1162/neco.2008.12-06-420

Rosas, P., Wagemans, J., Ernst, M. O., & Wichmann, F. A. (2005). Texture and haptic cues in slant discrimination: reliability-based cue weighting without statistically optimal cue combination. *Journal of the Optical Society of America A, Optics and Image Science*, *22*(5), 801–809.

Rosas, P., Wichmann, F. A., & Wagemans, J. (2007). Texture and object motion in slant discrimination: failure of reliability-based weighting of cues may be evidence for strong fusion. *Journal of Vision*, *7*(6), 3. http://doi.org/10.1167/7.6.3

Sahani, M., & Dayan, P. (2003). Doubly distributional population codes: Simultaneous representation of uncertainty and multiplicity. *Neural Computation*, *15*(10), 2255–2279.

Samaha, J., Barrett, J. J., Sheldon, A. D., LaRocque, J. J., & Postle, B. R. (2016).

Dissociating Perceptual Confidence from Discrimination Accuracy Reveals No Influence of Metacognitive Awareness on Working Memory. *Frontiers in Psychology*, *7*, 851. http://doi.org/10.3389/fpsyg.2016.00851

Schurger, A., Kim, M.-S., & Cohen, J. D. (2015). Paradoxical Interaction between Ocular Activity, Perception, and Decision Confidence at the Threshold of Vision. *PloS One*, *10*(5), e0125278. http://doi.org/10.1371/journal.pone.0125278

Sherman, M. T., Seth, A. K., Barrett, A. B., & Kanai, R. (2015). Prior expectations facilitate metacognition for perceptual decision. *Consciousness and Cognition*, *35*, 53–65. http://doi.org/10.1016/j.concog.2015.04.015

Simon, H. A. (1956). Rational choice and the structure of the environment. *Psychological Review*, *63*(2), 129–138.

Song, A., Koizumi, A., & Lau, H. (2015). A Behavioral Method to Manipulate Metacognitive Awareness Independent of Stimulus Awareness. In M. Overgaard (Ed.), *Behavioral Methods in Consciousness Research*. Oxford: Oxford University Press.

Spence, M. L., Dux, P. E., & Arnold, D. H. (2015). Computations Underlying Confidence in Visual Perception. *Journal of Experimental Psychology. Human Perception and Performance*. http://doi.org/10.1037/xhp0000179

Trommershäuser, J., Körding, K. P., & Landy, M. S. (Eds.). (2011). *Sensory Cue Integration*. New York: Oxford University Press.

Usher, M., & McClelland, J. L. (2001). The time course of perceptual choice: the leaky, competing accumulator model. *Psychological Review*. http://doi.org/10.1037/0033-295X.108.3.550

van Beers, R. J., Sittig, A. C., & Denier van der Gon, J. J. (1996). How humans combine simultaneous proprioceptive and visual position information. *Experimental Brain Research*, *111*(2), 253–61.

van Bergen, R. S., Ma, W. J., Pratte, M. S., & Jehee, J. F. M. (2015). Sensory uncertainty decoded from visual cortex predicts behavior. *Nature Neuroscience*, *18*(12), 1728–1730. http://doi.org/10.1038/nn.4150

Vickers, D. (1970). Evidence for an accumulator model of psychophysical discrimination. *Ergonomics*, *13*(1), 37–58. http://doi.org/10.1080/00140137008931117

Vickers, D., & Packer, J. (1982). Effects of alternating set for speed or accuracy on response time, accuracy and confidence in a unidimensional discrimination task. *Acta Psychologica*, *50*(2), 179–97.

Vlassova, A., Donkin, C., & Pearson, J. (2014). Unconscious information changes decision accuracy but not confidence. *Proceedings of the National Academy of Sciences*, *111*(45), 16214–16218. http://doi.org/10.1073/pnas.1403619111

Wilimzig, C., Tsuchiya, N., Fahle, M., Einhäuser, W., & Koch, C. (2008). Spatial attention increases performance but not subjective confidence in a discrimination task. *Journal of Vision*, *8*(5), 1–10. http://doi.org/10.1167/8.5.7

Winman, A., & Juslin, P. (1993). Calibration of sensory and cognitive judgments: Two different accounts. *Scandinavian Journal of Psychology*, *34*(2), 135–148. http://doi.org/10.1111/j.1467-9450.1993.tb01109.x

Zemel, R. S., Dayan, P., & Pouget, A. (1998). Probabilistic Interpretation of Population Codes. *Neural Computation*, *10*(2), 403–430.

http://doi.org/10.1162/089976698300017818

Zhang, H., & Maloney, L. T. (2012). Ubiquitous log odds: a common representation of probability and frequency distortion in perception, action, and cognition. *Frontiers in Neuroscience*, *6*, 1. http://doi.org/10.3389/fnins.2012.00001

Zylberberg, A., Roelfsema, P. R., & Sigman, M. (2014). Variance misperception explains illusions of confidence in simple perceptual decisions. *Consciousness and Cognition*, *27*, 246–253. http://doi.org/10.1016/j.concog.2014.05.012