Studying the neural representations of uncertainty

Edgar Y WALKER[*], Department of Physiology and Biophysics, Computational Neuroscience Center, University of Washington, Seattle, WA, eywalker@uw.edu

Stephan POHL[*], Department of Philosophy, New York University, stephan.pohl@nyu.edu

Rachel N DENISON, Department of Psychological & Brain Sciences, Boston University, Boston, MA, rdenison@bu.edu

David L BARACK, Departments of Neuroscience and Philosophy, University of Pennsylvania, Philadelphia, PA, dbarack@gmail.com

Jennifer LEE, Center for Neural Science, New York University, jll616@nyu.edu

Ned BLOCK, Department of Philosophy, New York University, ned.block@nyu.edu

Wei Ji MA[#] Center for Neural Science and Department of Psychology, New York University, weijima@nyu.edu

Florent MEYNIEL[#, §], Cognitive Neuroimaging Unit, INSERM, CEA, CNRS, Université Paris-Saclay, NeuroSpin center, Gif/Yvette, France, florent.meyniel@cea.fr

[*]: co-first authors; [#]: co-senior authors; [§]: corresponding author

Abstract

The study of the brain's representations of uncertainty is a central topic in neuroscience. Unlike other cases of representation, uncertainty is a property of an observer's representation of the world, posing specific methodological challenges. We analyze how the literature on uncertainty addresses those challenges and distinguish between "descriptive" and "process" approaches. Descriptive approaches treat uncertainty reported by subjects or inferred from stimuli as an independent variable used to test for a relationship to neural responses. By contrast, process approaches treat uncertainty derived from models of neural responses as a dependent variable used to test for a relationship to subjects' reports or stimuli. To compare those two approaches, we apply four criteria for neural representations: sensitivity, specificity, invariance, functionality. Experiments can be cataloged by their approach and whether they test for each criterion. Our analysis rigorously characterizes the study of neural representations of uncertainty, shaping research questions and guiding future experiments.

Introduction

Understanding how the brain represents its environment is one of the major goals of neuroscience and psychology. Another major goal is to understand the uncertainty of these representations. The uncertainty of representations is the topic of the Bayesian brain hypothesis (Knill & Pouget 2004) and probabilistic accounts of perception and cognition (Ballard 2015; Hoyer & Hyvärinen 2002; Lee & Mumford 2003; Ma & Jazayeri 2014). Taking into account uncertainty in perceptual processing can be crucial when interacting with the world. Imagine that while you are hiking through the mountains you have to decide whether to attempt to cross a steep slope. Your perception of the slope will determine whether to attempt to cross it. Given the stakes, it would be beneficial to take into account your uncertainty about the slope. Perhaps you should move closer to it in order to reduce your uncertainty before you decide whether or not to take a detour. A wide range of human behavior takes into account uncertainty, including decision making (Bach & Dolan 2012; Muller *et al.* 2019; Qamar *et al.* 2013; Rushworth & Behrens 2008; Tomov *et al.* 2020), learning (Behrens *et al.* 2007; McGuire *et al.* 2014; Meyniel *et al.* 2015a; O'Reilly 2013), perception (Kersten *et al.* 2004; Knill & Pouget 2004; Zhou *et al.* 2020) including multi-sensory fusion (Deroy *et al.* 2016; Ernst & Banks 2002), motor control (Todorov 2004; Trommershäuser *et al.* 2008), and memory (Devkar *et al.* 2017; Koriat *et al.* 2002; Rademaker *et al.* 2012). Similar observations have been made in non-human animals (Dekleva *et al.* 2016; Fiorillo 2003; Kepecs *et al.* 2008; Kiani & Shadlen 2009; Komura *et al.* 2013; Lak *et al.* 2014; Odegaard *et al.* 2018; Walker *et al.* 2020).

Many neuroscientists aim to understand how this uncertainty is represented in the brain. Studies of uncertainty often contain claims of the form "in a given brain region, neural activity *r* represents uncertainty about *s*". In practice, *r* can be measured with fMRI, EEG, intracranial recordings of local field potentials, spike trains, amongst others, and *s* can be the orientation of a stimulus, a reward probability, or some other feature. These claims about the representation of uncertainty should be assessed empirically based on the experiments and analyses used. The goal of this article is to provide a framework to assess the empirical validity of claims of uncertainty representation.

1. Defining uncertainty: its objects, origins, and measures

Uncertainty characterizes the representation of a world state by an observer

Consider again the example of a subject who has to perceive the slope of a field. If the slope is too steep to be passable, they have to take a strenuous detour through the valley. In order to appropriately weigh the risk of attempting the slope and the cost of taking the detour, the subject should take into account the uncertainty in their perceptual representation of the slope.

We will understand this situation in terms of generative models (see glossary) as depicted at the top of Figure 1. The subject is an observer who represents *s*, some feature of interest of the world state. Here *s* is the slope of the field. This feature *s* is typically not directly observable, and thus called a *latent* feature. Instead, the observer receives information about *s* from the more

proximal input state *I*. The brain processes this input to arrive at the neural response *r* which is a representation of *s*.

In an experimental context, *s* typically refers to some latent feature of the world state that is of interest in a task, while *I* is the input to the observer that is generated by the feature in a particular trial. In a visual task, *I* is the pattern of light that hits the observer's retina, which in practice is typically considered equivalent to the pattern of pixels presented on a screen (any discrepancy between the two is not relevant for the purpose of this article). As a standard example throughout this article, let *s* be an orientation value and *I* be an image on a screen: a Gabor patch with orientation *s* and added random pixel noise (see Figure 1, top). Subjects report the orientation *s* which they see in the noisy image of a Gabor patch *I*.

For simplicity, we will first discuss uncertainty about the world state *s* given the input *I*, and then uncertainty that depends on the internal neural state *r* of an observer.

We consider an observer who has access to the input state *I* and forms a representation of the world state *s* through a process that can be described as an inference from *I* by inverting a generative model (Helmholtz 1867). That is, the observer computes values for *s* given the observed *I* (I_{observed}) and the dependence of *I* on *s* in the generative model which is assumed in the inference.

However, often, one and the same input can be generated from multiple states of the world. That is, the state of the world is underdetermined given the input. When the inference an observer performs leaves the world state *s* underdetermined, the observer is uncertain about *s*.

Unlike other quantities that are represented in perception, uncertainty is not a property of the state of the world. Uncertainty is not another stimulus parameter besides *s*. Rather, uncertainty is a property of the relation between the world and an observer. It measures the lack of information the observer has about the world state *s* on the basis of an inference from a more proximal state like *I* (Shannon 1948). To emphasize this dependence on an inference, people sometimes speak of *epistemic* or inferential uncertainty.

Consider the uncertainty about *s* given I_{observed} on a given trial. We can visualize this uncertainty by looking at the posterior probability distribution $p(s|I_{\text{observed}})$, see Figure 1. This distribution describes how probable values of *s* are given the particular input I_{observed}. If there were no uncertainty, the value of *s* would be perfectly determined by some particular input I_{observed}, and all the probability mass in the distribution $p(s|I_{\text{observed}})$ would be assigned to just a single value of *s*. But since there is uncertainty about *s* given *I*, multiple values of *s* are possible given a particular input I_{observed} and the probability mass in the distribution $p(s|I_{\text{observed}})$ is spread out over those values.

The uncertainty about *s* given I_{observed} based on optimal inference and the true generative model (the one that captures the actual dependencies from *s* to *I*) is called the *ideal-observer uncertainty* (Barthelmé & Mamassian 2009; Kersten *et al.* 2004; Ma & Jazayeri 2014).

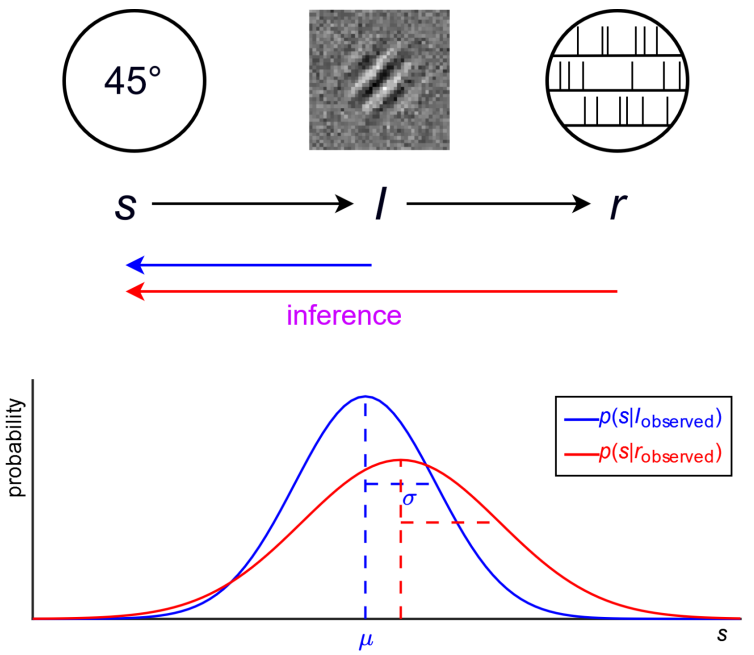

Figure 1 Uncertainty from a generative model. Top: Example generative model. The world state *s* is an orientation. The input *I* is an image of a Gabor patch with orientation *s* and some pixel noise. The neural response *r* is a train of spikes from neurons in some population. **Middle**: Black arrows indicate the dependencies in the generative model between the world state *s*, input *I*, and neural response *r*. Red and blue arrows indicate inferences an observer could make by inverting dependencies in the generative model from *I* and *r*, respectively . **Bottom**: Probability distributions of *s* given the particular input $I_{observed}$ or the particular neural response $r_{observed}$ in a given trial. σ is a measure of the uncertainty about *s* given $I_{observed}$, and μ is the expected value of *s* given $I_{observed}$.

Origins of uncertainty

Uncertainty about the world state *s* is present whenever *s* is under-determined given the inference an observer performs. In the structure of the generative model assumed in an observer's inference, this under-determination manifests as a many-to-one mapping from the world state to a later state in the generative model.

A typical source of under-determination is the ambiguity of the input as illustrated by the case of the Necker cube (Necker 1832). One and the same two dimensional image could be interpreted to be the result of different states of the three dimensional world. An observer who infers the state of the world from the image is uncertain between these different states.

Another common source of uncertainty is aleatoric variability. Aleatoric variability is a measure of the indeterminacy of a process by which an earlier state in a generative model generates a later state. One important type of such variability is noise in the input the observer receives. Consider for instance pixel noise which corrupts an image of a Gabor patch with a certain orientation (see Figure 1). The image *I* is generated non-deterministically based on the orientation *s* and random noise. An observer is left uncertain when the image *I* under-determines the world state *s* because of noise.

Aleatoric variability increases uncertainty all the more the less data is available to the observer. For instance when there is pixel noise, the orientation task in Figure 1 is easier for higher pixel

resolutions of the image. Similarly, when *s* is not stationary but changes across time, the uncertainty depends on how long *s* can be observed and how much evidence about *s* can be accumulated (McGuire *et al.* 2014; Meyniel 2020; Meyniel & Dehaene 2017; O'Reilly *et al.* 2013; Payzan-LeNestour *et al.* 2013).

Another important type of aleatoric variability is internal noise. The firing behavior of neurons is only partially driven by the signal they receive and partially by further random factors (Faisal *et al.* 2008). Hence one and the same neural response *r* can be caused by different input states *I* which in turn depend on different world states *s*. An actual observer does not infer *s* from *I*, but from *r*. The red curve in the bottom of Figure 1 visualizes uncertainty about *s* given a particular r_{observed} in terms of the probability distribution $p(s|r_{\text{observed}})$. Uncertainty is a measure of the spread of this curve. The red curve is wider than the blue curve since it includes not only the uncertainty about *s* given *I*, but the additional uncertainty about *I* given *r*.

Lastly, uncertainty depends on the generative model assumed by the observer. An ideal observer performs an optimal inference based on the true generative model. However, an actual observer might not know the true generative model, or it might be so complex that the observer is not able to compute an optimal inference (Beck *et al.* 2012; Rahnev & Denison 2018). The assumed generative model might for instance exclude some variables and simplify the shapes of probability functions (Iglesias *et al.* 2013; Mathys *et al.* 2014; Norton *et al.* 2019). The uncertainty an observer actually has about *s* depends on the assumed generative model which actually underlies their inference to *s*.

Measures of uncertainty

Different formal measures are available to summarize uncertainty. Often, uncertainty is understood as the standard deviation of a random variable under a posterior distribution, especially for the frequently used Gaussian distribution (see Figure 1, bottom). The larger the standard deviation, i.e. the wider the spread of the probability distribution, the more uncertainty there is about this variable.

Another useful measure of uncertainty is entropy. The conditional entropy $H(s|r_{\text{observed}})$ is a measure for how much information or freedom of choice (hence, uncertainty) there is left about the variable *s*, after one already knows that the variable *r* takes the value $r_{observed}$ (Shannon 1948). The advantage of entropy as a measure of uncertainty is that it applies to probability distributions of any shape (categorical and numeric variables with some adjustments, and with one or more dimensions). However, the fact that entropy ignores ordinality can be a disadvantage: for instance, if two orientations have high probability and all other orientations have the same low probability, entropy (unlike standard deviation) will be the same no matter whether those two orientations are very close or very different.

Yet rather than summarizing uncertainty in a single quantity, one might also keep track of it implicitly in terms of the full probability distribution. If one were to represent the state of the world *s* in terms of, for instance, the posterior distribution $p(s|r_{\text{observerd}})$, uncertainty about *s* would be implicit in that representation (Ma & Jazayeri 2014). This uncertainty can be taken into account implicitly by performing computations over the full probability distribution. Whether observers use full distributions or summaries is an open empirical question (Fleming & Daw 2017; Meyniel *et al.* 2015b; Yeon & Rahnev 2020). In this article we do not commit to a specific measure and we use the standard deviation simply as an example.

Box 1: What is special about studying neural representations of uncertainty?

Several aspects of uncertainty make it special in comparison to other features studied in neuroscience and deserve consideration.

Uncertainty is not a world state. Instead, it is a property of the relation between a (real or hypothetical) observer's representation and the corresponding world state. Thus, studying uncertainty requires taking the observer's perspective to estimate their uncertainty, based on subjective reports or (potentially idealized) models of the process by which they arrive at their representation of the world state the uncertainty is about. The dominant tool, which we use in this article, is that of Bayesian inference.

Uncertainty can have multiple origins. Some are internal to the observer (e.g. neural noise, limitations and errors in information processing) the others are external (aleatoric variability, ambiguity, non-stationarity). Only the latter are under the experimenter's control.

Uncertainty can characterize representations of many different world states (see Table 1).

Uncertainty can be summarized in different ways: with different measures (variance, entropy, maximum of a probability distribution), with different models of the underlying inference, and different generative models assumed in that inference.

The main origin of uncertainty, the representation it characterizes and the way it is measured and modeled differ across studies. This diversity introduces confusion and makes it difficult to have a unified understanding of the neural representation of uncertainty.

Uncertainty is particularly prone to confounding factors reflecting its origins, specific aspects of the world state it refers to and behavior. For instance, the image resolution and contrast of a Gabor patch impact the uncertainty about the patch orientation; an apparent neural representation of uncertainty can actually be a representation of resolution or contrast.

2. Studying representations of uncertainty with descriptive and process approaches

How is uncertainty represented in the brain? The classic approach in neuroscience is to study the neural representation of some world state *s*, like the orientation of a bar (Hubel & Wiesel 1959, 1962), by experimentally inducing variations in *s* across trials, and then identifying candidate regions for the representation of *s* as those whose changes in activity *r* are related to changes in *s*. In such an analysis, *s* is the independent variable and *r* the dependent variable. The independent variable need not be under the experimenter's control (for instance, representation of location in space can be studied by letting an animal move freely), and instead for the purpose of the analysis of neural data it can be determined without knowing *r*.

This classic approach is not directly applicable to uncertainty because uncertainty not being a world state (see Box 1), the experimenter cannot determine it without knowing *r*. Other

approaches have been proposed in previous studies on the neural representation of uncertainty, which we categorize as descriptive and process approaches.

The descriptive approach is the one most inspired by the classic approach in neuroscience because it aims to treat the uncertainty *u* as an independent variable and test whether it is related to *r*. Doing so requires using proxies for *u* that can be determined without knowing *r*, from the subject's behavior, the stimuli, and the task. By contrast, the process approach differs more radically because it treats *u* as a dependent variable that the researcher derives from neural activity *r* (the data that the brain has at its disposal to represent the world). Doing so requires using models of how *r* represents the world, with the overarching goal of going from input to neural activity to behavior.

These approaches require different assumptions and methods, which we present in this section and summarize in Figure 2. Those differences have consequences regarding the type of neural representations of uncertainty that each approach can uncover and how the validity of their findings can be evaluated.

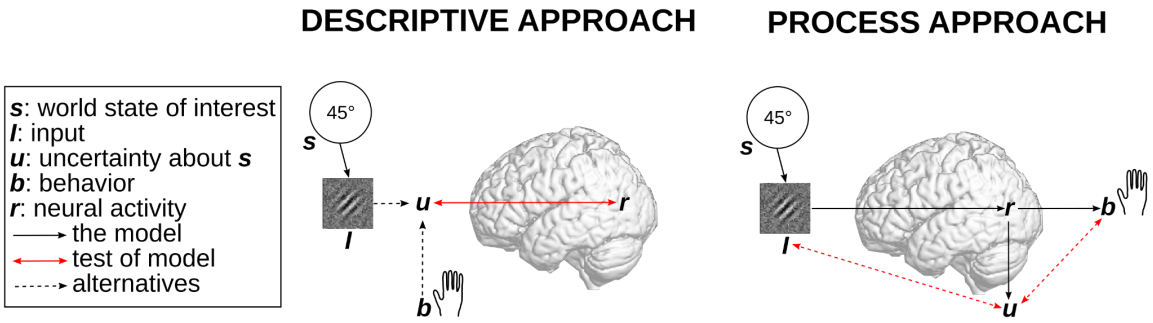

Models:	of ***r***: $f(\boldsymbol{I}, \boldsymbol{b})$ of ***I***: $p(\boldsymbol{I}\|\boldsymbol{s})$ (needed if ***b*** not available)	of ***r***: $p(\boldsymbol{r}\|\boldsymbol{s})$ in PPC sampled from $p(\boldsymbol{s}\|\boldsymbol{I})$ in SBC
Researchers derive ***u*** from:	$\boldsymbol{I}_{\text{observed}}$ (e.g. as SD$[\boldsymbol{s}\|\boldsymbol{I}_{\text{observed}}]$) or ***b*** (e.g. through report, RT)	$\boldsymbol{r}_{\text{observed}}$ as SD$[\boldsymbol{s}\|\boldsymbol{r}_{\text{observed}}]$ in PPC and SD$[\boldsymbol{r}_{\text{observed}}]$ in SBC
Test of model:	***u*** and ***r*** are related	***u*** and ***b*** are related or ***u*** and ***I*** are related

Figure 2: Comparison of the descriptive and process approaches. In both approaches, the subject is provided with a input *I* that is informative about a particular world state *s*. In the descriptive approach, researchers use a proxy for the uncertainty *u*, derived from the alternatives *I* itself (using an "external" generative model $p(I|s)$) or from the subject's behavior *b*. Then they test for a relationship between this proxy of *u* derived from *I* or *b*, and the neural activity *r,* which we summarize as $r=f(I, b)$. The function *f* models both how *u* can be derived from *I* or *b* in different studies, and the specific relation between *u* and *r* (e.g. linear, multivariate, etc.). In the process approach, researchers use a "neural" generative model of *r* to derive the uncertainty *u* about *s* given the observed *r* (e.g. by inverting the neural likelihood function $\mathcal{L}(s;r_{\text{observed}})=p(r_{\text{observed}}|s)$ in a probabilistic population code (PPC); or as the standard deviation of neural activity which is interpreted as samples from the posterior distribution $p(s|I)$ in sampling-based codes (SBC)). The process approach is then evaluated by testing for some expected relationship between the inferred *u* and the alternatives *b* or *I*. SD[***x***] denotes the standard deviation of the samples $x_1, \ldots, x_n$, and SD[*x*|*y*] denotes the standard deviation of the random variable *x* under a conditional probability distribution (conditioned on *y*) and RT denotes reaction times.

Descriptive approach: Uncertainty as an independent variable

Studies that follow the descriptive approach can be divided into those that use the subject's behavior *b* and those that use the input *I* as proxies for uncertainty.

Different aspects of behavior can be used to study uncertainty. In some studies, human subjects report their uncertainty in the form of a rating or confidence judgment (Adler & Ma 2018; De Martino *et al.* 2013; Guggenmos *et al.* 2016; Hebart *et al.* 2014; Lebreton *et al.* 2015; Meyniel 2020; Meyniel *et al.* 2015b; Meyniel & Dehaene 2017; Peirce & Jastrow 1884; Pouget *et al.* 2016). This uncertainty can then be treated as an independent variable in the analysis of neural activity. Some researchers infer uncertainty by assuming that it regulates some specific aspects of subject's behavior, such as how fast to respond (Tzagarakis *et al.* 2010; Zylberberg *et al.* 2016), how long to wait for a reward (Kepecs *et al.* 2008; Masset *et al.* 2020; Schmack *et al.* 2021), or whether to opt-out of a bet (Gherman & Philiastides 2015; Hampton 2001; Kiani & Shadlen 2009; Komura *et al.* 2013; Middlebrooks & Sommer 2012; Odegaard *et al.* 2018). This strategy is common in studies on decision confidence, in particular those involving non-human animals, see (Kepecs & Mainen 2012; Meyniel *et al.* 2015b) for reviews.

Some researchers use the input *I* to study the uncertainty about a world state *s*. An "external" generative model describes the way a given *s* generates *I* in the task. When this process is stochastic, the relationship between *s* and *I* is described by the probability function $p(I|s)$. The ideal-observer uncertainty can be derived from this generative model (see section 1). For oriented Gabor patches (see Figure 1), ideal-observer uncertainty is commonly introduced by adding Gaussian noise at the pixel level (Cabrera *et al.* 2015; Kersten *et al.* 2004; Rahnev *et al.* 2011). Increasing the amount of pixel noise makes the resulting input I_{observed} compatible with a broader range of orientations, which increases uncertainty. Similarly, reducing the image contrast also increases uncertainty. The ideal-observer model quantifies this (externally generated) uncertainty for a given I_{observed}, which can be used as a proxy for *u* in the analysis of neural data. Ideal-observer models are useful to quantify uncertainty across a wide variety of task structures, notably when the relation between *u* and *I* is complex, e.g. in sequential learning (Badre *et al.* 2012; Meyniel 2020; Meyniel & Dehaene 2017; Muller *et al.* 2019). However, when simple qualitative relationships between *I* and *u* exist, such qualitative aspects of *I* (e.g. pixel noise or contrast in the oriented Gabor example) can be used as a proxy for *u* (Bang & Fleming 2018; Festa *et al.* 2021; Hénaff *et al.* 2020; Kepecs *et al.* 2008; Kiani & Shadlen 2009; Orbán *et al.* 2016).

No matter whether the proxy for *u* is derived from *b* or *I*, the descriptive approach seeks to relate this proxy to *r*, which we formalize as $r=f(I, b)$ (see Figure 2). Here *f* is both a model of how the proxy for *u* is derived from either *I* or *b* and of the relationship between *r* and the proxy for *u*. In practice, researchers test for this relationship by different means. They might simply search for correlations or use more complex models like multiple linear regressions and multivariate pattern decoding (Haxby *et al.* 2014). The strength of this relationship can be measured as a correlation coefficient, the significance of regression weights (Friston *et al.* 2007), or the cross-validated fraction of explained variance (Naselaris *et al.* 2011; Poldrack *et al.* 2020).

Process approach: Uncertainty as a dependent variable

In contrast to the descriptive approach, the process approach aims to preserve the fact that uncertainty is a property of the relation between the neural representation *r* and the corresponding world state *s*. A researcher following this approach thus treats uncertainty as a dependent variable in the data analysis where *u* is a function of *r*, *u*(*r*), not a function of *b* or *I*. This approach requires specifying how *r* represents *s* by means of a "neural" generative model. We present here two families of such models: those that rely on a neural likelihood function $\mathcal{L}(s;r)=p(r|s)$ (see glossary) derived from a generative model about *r* (broadly referred to as a Probabilistic Population Code) and those that rely on a neural code that samples from a posterior distribution over *s* (Sampling-Based Code). Other models relevant for the study of uncertainty exist (Barlow 1969; Deneve 2008; Jazayeri & Movshon 2006; Knill & Pouget 2004; Sahani & Dayan 2003; Sohn & Narain 2021), but have so far received less attention from experimenters.

In Probabilistic Population Codes, researchers formalize the uncertainty *u* conveyed by the neural activity observed on a particular trial, $r_{observed}$, as the posterior distribution $p(s|r_{observed})$. This posterior is derived by relying on a neural likelihood function $\mathcal{L}(s;r_{observed})$ and Bayes' rule. The construction of $\mathcal{L}(s;r)$ can be theory-driven or data-driven. The former constitutes the dominant approach in the literature (Dayan & Abbott 2005; Pouget *et al.* 2003). An influential example in the sensory domain posits that neurons have a stereotyped mean response to the input (known as their tuning curve) and some variability corresponding to the exponential family of distributions (e.g. Poisson distributions). Together with a few other assumptions, the log of $\mathcal{L}(s;r)$ becomes linear with respect to *r* and uncertainty about *s* is proportional to the average neural activity on a given trial (Deneve *et al.* 1999; Fetsch *et al.* 2012; Ma *et al.* 2006). By contrast, the data-driven approach requires fewer assumptions and estimates $\mathcal{L}(s;r)$ from the data itself. With the advent of large datasets and machine learning tools like artificial neural networks, even arbitrary shapes of $\mathcal{L}(s;r)$ can be estimated (Walker *et al.* 2020). With a smaller amount of data, further constraints are needed about the shape of $\mathcal{L}(s;r)$, e.g. assumptions about specific covariance matrices or noise distributions (Geurts *et al.* 2022; Li *et al.* 2021; van Bergen *et al.* 2015; van Bergen & Jehee 2019).

In Sampling-Based Codes, the neural activity (e.g., firing rate) is assumed to represent *s* in terms of samples stochastically drawn from the posterior distribution $p(s|I_{observed})$ (Fiser *et al.* 2010; Hoyer & Hyvärinen 2002; Lee & Mumford 2003; Orbán *et al.* 2016). Biologically plausible neural network models have been proposed for such a "neural" generative model (Echeveste *et al.* 2020). Under such a code, *u* is reflected in the spread of the distribution of *r* (e.g. the standard deviation) across time or across neurons.

In experimental studies that adhere to the process approach, uncertainty about *s* is not used to construct the neural generative model. Rather, researchers assume a neural generative model and use it to derive the uncertainty about *s* from *r, u(r)*. In order to be validated, the process approach needs to demonstrate that *u(r)* relates to the subject's uncertainty about *s*. One possibility is to test whether *u(r)* is reflected in the subject's actions, like choices based on uncertainty (Walker *et al.* 2020), behavioral variability (van Bergen *et al.* 2015), and reported confidence (Geurts *et al.* 2022). Another possibility inspired by the descriptive approach is to change *s* or *I* in a way that leads to predictable changes in *u*, e.g. change the contrast of the

image (Orbán *et al.* 2016; Walker *et al.* 2020) or use cardinal and oblique orientations (Orbán *et al.* 2016; van Bergen *et al.* 2015), and test whether *u(r)* changes accordingly.

Box 2: Relation to encoding and decoding approaches

Encoding and decoding are widely used notions in neuroscience (Lange *et al.* 2021) which are related to the distinction between descriptive and process approaches. However, because different notions of encoding and decoding exist, the relationship is multi-faceted.

In the data analysis domain (Huth *et al.* 2016; Naselaris *et al.* 2011), an encoding model corresponds to a specific relationship between *r* and *s* (or *I*). First, a mapping (often highly non-linear) from *s* (or *I*) to a list of latent features is assumed, then the relation to *r* is tested by means of a multiple linear regression of the latent features onto *r* (*linearizing encoding models* (Huth *et al.* 2016)) or a generalized linear model (Pillow *et al.* 2008). This notion of encoding is used in the descriptive approach: some features related to *u* are derived from *I* and regressed onto *r*. This notion of encoding has also been used together with the process approach. For instance, van Bergen and colleagues (van Bergen *et al.* 2015) used an encoding model and fit the residuals of the corresponding multiple linear regression (an unusual step in other fMRI studies that use encoding models) to turn their fit encoding model into a probabilistic neural generative model *p(r|s)*.

In theoretical and computational neuroscience (Dayan & Abbott 2005; Lange *et al.* 2021; Park *et al.* 2014), an encoding model captures the neural likelihood function $\mathfrak{L}(s;r)$. This notion then corresponds to the process approach.

In machine learning and its application to neuroimaging (Haxby *et al.* 2001, 2014; Haynes 2015), decoding models are trained based on a collection of pairs of data and labels (e.g. here, *r* and *u* respectively) and tested on new data. The label (*u*) being an independent variable, this decoding corresponds to the descriptive approach. In the context of uncertainty, those labels are derived from the input *I* or behavior.

In theoretical and computational neuroscience, decoding corresponds to the estimation of the cause of the observed neural activity (it goes from *r* to *s*) by inverting a neural generative model (which goes from *s* to *r*); when this model is probabilistic, the uncertainty *u* that accompanies this estimation is quantifiable. The key difference with the machine learning and descriptive approaches is that *u* is not an independent variable but instead a dependent variable that emerges through decoding, as in the process approach.

3. Evaluating claims about representations of uncertainty with general criteria of representation

We now turn to the evaluation of claims about the neural representation of uncertainty that can be found in studies that rely on either the descriptive or process approach. We propose to do so by applying criteria that are generally used in neuroscience to support claims of representations: sensitivity, specificity, invariance and functionality (defined in Box 3). This choice appears

legitimate since many studies on the neural representation of uncertainty presented below (Table 1) use methods and arguments that aim to fill those general criteria.

Box 3: General criteria for neural representations

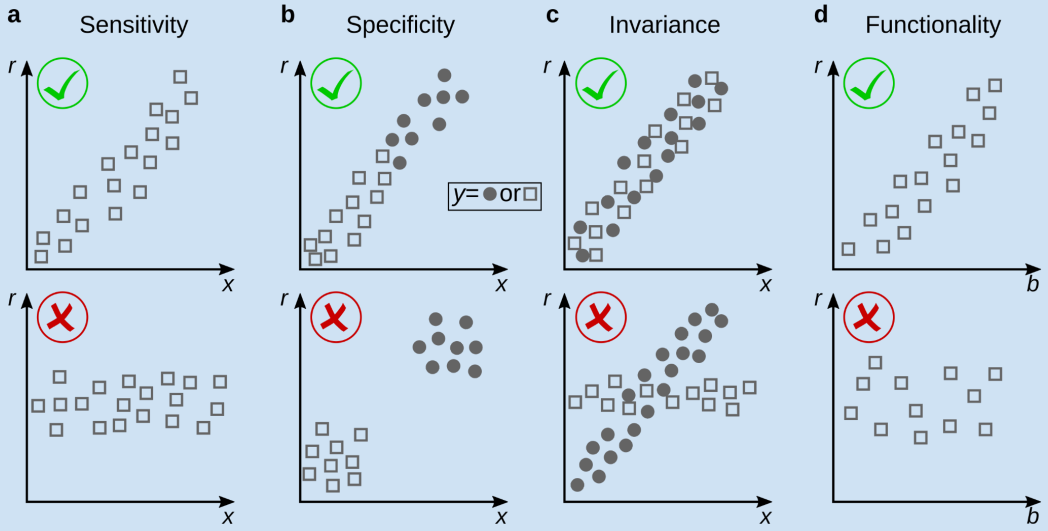

Figure 3: Simple examples of successes and failures of different criteria of neural representation. The criteria for testing for a representational relationship between neural activity *r* and a feature of interest *x* (like uncertainty in this article) or behavior *b*, whether by itself and in comparison to another feature *y* (which is categorical here for simplicity, it could also be continuous). The top row shows examples that pass a given criteria and the bottom row examples that fail. Note that the horizontal axis corresponds to *x* in **a-c** and *b* in **d**.

Sensitivity: *r* is sensitive to a feature *x* if changes in *r* are related to changes in *x*. For instance, a neuron is sensitive to the orientation of a bar if different activity patterns are recorded when different orientations are presented.

Specificity: *r* represents *x* specifically with respect to another feature *y* if changes in *r* are related to changes in *x* even when controlling for *y*. This criterion enables the researcher to rule out the effect of confounding variables. For instance, uncertainty about orientation depends on the image contrast: a neural representation of the uncertainty about orientation therefore ought to be sensitive to contrast. However, to be a representation of uncertainty *per se* rather than contrast, *r* should reflect uncertainty even when the image contrast is kept fixed.

Invariance: the representation of *x* by *r* is invariant to *y* if changes in *r* are not related to changes in *y* when controlling for *x*. For instance, the representation of orientation (our *x*) by V1 neurons (our *r*) is not invariant to position since different *r* are observed for a given orientation when changing position in the visual field (Hubel & Wiesel 1959). By contrast, the representation of object identity in the infero-temporal cortex is invariant to the position and orientation of objects (DiCarlo *et al.* 2012).

Functionality: *r* plays a functional role in producing *b* if changes in *r* are related to changes in *b*. This notion of functionality corresponds to a form of sensitivity; we can also apply the

notions of specificity and invariance to functionality. For instance, if *x* causes *r* which itself causes *b*, then the mediating role of *r* can be assessed with a test of specificity (Are changes in *r* related to changes in *b* even when holding *x* fixed?) and invariance (Is *b* the same for a given *r* when *x* changes?).

Claims about representations, ultimately, have to be claims about the causal structure of information processing in the brain. Nonetheless, we express the criteria in terms of information theoretic relationships between variables rather than in causal terms because researchers often use correlative (not causal) methods. Previous studies on uncertainty have used correlation (Fiorillo 2003; Geurts *et al.* 2022; van Bergen *et al.* 2015), differences between conditions (Kiani & Shadlen 2009; Komura *et al.* 2013; Vilares *et al.* 2012), linear regression (Hebart *et al.* 2014; Meyniel 2020; Meyniel & Dehaene 2017; Payzan-LeNestour *et al.* 2013), or decoding (Cortese *et al.* 2016; Gherman & Philiastides 2018; Hebart *et al.* 2014; Monosov & Hikosaka 2013; Nastase *et al.* 2018) to establish representations of uncertainty.

In practice, the criteria are graded rather than all-or-none (sensitivity might for instance be measured in terms of the strength of the correlation between *r* and *x*). Moreover, while linear relationships such as those illustrated in Figure 3 are often simplest to understand and most common in experimental studies, many models, especially process models, posit non-linear relationships; these can still be assessed in terms of the mutual information between variables.

How do the two approaches compare in terms of testing general criteria?

We focus the discussion on epistemic uncertainty because it is the target of most process models, making this case best suited for comparing the descriptive and the process approaches. The descriptive approach is more flexible and Table 1 provides broader coverage, including examples on aleatoric variability and decision confidence.

Sensitivity. The descriptive approach treats the uncertainty *u* as an independent variable and tests for sensitivity, i.e. the degree to which the response *r* depends on *u*. The higher the sensitivity, measured for instance in terms of the strength of correlation or decoding accuracy, the more plausible it is that a given *r* is a representation of *u*. Some studies that follow the descriptive approach *only* address sensitivity, especially when they are among the first of their kind or when uncertainty is not central in the study (Bach *et al.* 2011; O'Reilly *et al.* 2013; Payzan-LeNestour *et al.* 2013).

By contrast, the process approach derives *u* from *r*. The derived representation of uncertainty *u(r)* is thus trivially sensitive to *r* and this relationship cannot constitute a test of this approach. Instead, as shown in Figure 2, studies following the process approach focus on the relationship between *u(r)* and *b* (see functionality below), or *u(r)* and *I*. In the latter case, researchers assess the sensitivity of *u(r)* to some feature of *I* that is expected to impact uncertainty, thereby showing some similarity with the descriptive approach that uses *I* as a proxy for *u*. To illustrate, uncertainty about the orientation of a Gabor patch is expected to decrease with increasing contrast (Echeveste *et al.* 2020; Orbán *et al.* 2016; Walker *et al.* 2020), close to cardinal orientations (van Bergen *et al.* 2015), or when embedded in a higher-level context (Bányai *et al.* 2019).

Specificity. In the descriptive approach, the researcher uses some aspects of *I* or *b* as a proxy for *u*. Studies following this approach are thus vulnerable to the problem of confounding variables: *r* may not represent *u* but the aspect of *I* or *b* from which *u* has been derived, such as contrast in the orientation task. It is still possible to test for specificity of *r* to *u* if several features of *I* or *b* are related to *u*. In that case, specificity of *r* to *u* with respect to each feature in isolation can be tested by holding each feature fixed and testing for the dependence of *r* on other features of *I* or *b* (Bang & Fleming 2018; Dekleva *et al.* 2016; Festa *et al.* 2021). For example, Dekleva and colleagues (Dekleva *et al.* 2016) manipulated uncertainty about a reaching direction through the current trial's cue and the cue history, and Bang and colleagues (Bang & Fleming 2018) manipulated uncertainty about the direction of motion by changing the strength of motion evidence and the distance to the category boundary. In both studies, *r* continued to track *u* when either feature was kept fixed. Some previous studies (see Table 1) include tests for potential confounding variables like reaction times (Grinband *et al.* 2006), attention (Grinband *et al.* 2006), exploration (Muller *et al.* 2019; Trudel *et al.* 2021), and task difficulty (Bang & Fleming 2018; Gherman & Philiastides 2015; Kepecs *et al.* 2008; Kiani & Shadlen 2009; Masset *et al.* 2020).

When uncertainty is not related to a simple feature of the input but derives from a more complex model, like an ideal-observer model, several confounding variables might still undermine the specificity of *r* to *u*. For instance, in the context of sequential learning, *u* is often negatively correlated with recently surprising outcomes (Meyniel 2020; Meyniel & Dehaene 2017; Strange *et al.* 2005). Confounding variables of epistemic uncertainty also include constructs presented in other studies, which we don't explain here, like the likelihood of a change point (McGuire *et al.* 2014), expected uncertainty (Tan *et al.* 2016), total uncertainty (Badre *et al.* 2012; Tomov *et al.* 2020), outcome uncertainty (Hsu *et al.* 2005; Meyniel 2020; Meyniel & Dehaene 2017), and expected reward (Monosov *et al.* 2015; Monosov & Hikosaka 2013; Preuschoff *et al.* 2006; So & Stuphorn 2016).

With regard to process models, once again similar to the descriptive approach, researchers who compare *u(r)* to some feature of *I* that is expected to impact *u* can control for specificity with respect to other features. For instance, uncertainty about local features in an image is expected to decrease when they are embedded in a higher-level structure; one can test for this effect while controlling for the spectral content of the image, which is often confounded with the presence of high-level structure (Bányai *et al.* 2019)).

Invariance. Some researchers following the descriptive approach tested for invariance. For instance, Michael and colleagues (Michael *et al.* 2015) used a categorization task and inputs with two features: shape and color. The relevant feature used for the categorization task changed across trials and a common neural representation *r* of the categorization uncertainty was found in both conditions. In the shape condition, *r* tracked the uncertainty related to shape, not color (and vice versa in the color condition), demonstrating that *r* coded for uncertainty abstracted away from low-level features. Using a similar logic, Lebreton and colleagues (Lebreton *et al.* 2015) found a general neural representation of the uncertainty associated with estimating the value of paintings, objects, and prospects. Other researchers have tested invariance with respect to sensory modality (Masset *et al.* 2020; Meyniel & Dehaene 2017; Nastase *et al.* 2018).

Studies that follow the process approach and test for a relation between *u(r)* and some feature of *I* could also test for invariance with respect to another feature of *I*. While possible in principle, we are not aware of such examples (see Table 1).

Functionality. Although the descriptive approach primarily tests for sensitivity of *r* to *u* (and the related criteria of specificity, invariance), some aspects of this approach are also relevant for functionality. When the behavior *b* is used as a proxy for *u*, then sensitivity of *r* to *u* implies sensitivity of *r* to *b*, and thus functionality (Lebreton *et al.* 2015). Compelling evidence of functionality also comes from descriptive studies that identify a neural representation of *u* based on *I* (not *b*) and then find relations between this neural representation of *u* and *b*. Based on this idea, some studies reported correlations between neural representations of *u* and the reported uncertainty (Meyniel 2020), trial-to-trial variability during reaching (Dekleva *et al.* 2016) and learning (McGuire *et al.* 2014); other studies reported correlations across subjects between neural representations of *u* and aspects of behavior that should in principle depend on uncertainty, like risk attitude (Blankenstein *et al.* 2017), exploration (Badre *et al.* 2012; Tomov *et al.* 2020), and in prior-likelihood combination (Ting *et al.* 2015).

In the context of the process approach, the functionality criterion often plays a key role (Figure 2 and Table 1). It is the case when the validity of the process model used to infer *u* from *r* is evaluated as the relation between *u*(*r*) and *b*. For instance, for probabilistic population codes, van Bergen and colleagues (van Bergen *et al.* 2015) showed that the uncertainty about the orientation of a Gabor patch inferred from V1 fMRI activity correlated with behavioral variability in orientation reports. The same group also found that this fMRI readout of uncertainty correlated across trials with both the uncertainty reported by the subject (Geurts *et al.* 2022) and the strength of sequential effects in their perceptual decisions (van Bergen & Jehee 2019). Using another task, Walker and colleagues (Walker *et al.* 2020) found that the uncertainty read out from V1 in monkeys accounted for their decisions in an uncertainty-based categorization task.

In contrast to studies that assume probabilistic population codes, functionality is less often tested in studies that assume sampling-based codes. Such studies thus appear closer to the descriptive approach in the way neural representations of uncertainty are tested. An exception is when Haefner and colleagues reported that the structure of covariance among artificial neurons in a network reflected the uncertainty about the task-relevant orientation during visual categorization in a way that correlated with performance in the task (Haefner *et al.* 2016).

The notion of functionality can be coupled with specificity. Such a test assesses whether *u(r)* still correlates with behavior *b* even when all features of the input *I* are held constant. Uncovering such an effect suggests that it is indeed uncertainty, and not any confounding feature of the input, that is represented. The process approach lends itself to such a test because the researcher infers *u* from *r* (not *I*) making it possible to find correlation between *u(r)* and *b* when *I* is fixed (van Bergen *et al.* 2015; Walker *et al.* 2020). The descriptive approach seems not suited for this test when it derives *u* from *I* but some analyses inspired by this test are still informative. For instance, McGuire and colleagues derived *u* from *I* based on an ideal observer and they identified a neural representation *r* of this *u*; they then regressed this ideal-observer uncertainty out of *r* (which is analogous to keeping the effect of *I* fixed) and showed that *r* still correlated with some aspect of behavior (McGuire *et al.* 2014).

Field	Type	Study	Context	Sensitivity	Specificity	Invariance	Functionality
epistemic uncertainty	process (PPC)	van Bergen et al 2015	uncertainty about orientation of Gabor. Human, fMRI	■	■		■
		van Bergen & Jehee 2019	uncertainty about orientation of Gabor. Human, fMRI				■
		Geurts et al 2021	uncertainty about orientation of Gabor. Human, fMRI	■			■
		Walker et al 2020	uncertainty about orientation of Gabor. Monkey, ephys	■	■		■
		Li et al 2021	uncertainty about the content of spatial working memory. Human, fMRI	■	■		■
		Dabney et al 2020	uncertainty about reward probability. Mice, ephys	■			■
	process (SBC)	Hénaff et al 2020	uncertainty about a (visual) feature. Macaque, ephys	■			
		Orbán et al 2016	uncertainty about visual stimulus. Macaque and cat ephys	■			
		Festa et al 2021	uncertainty about a local feature in an image. Macaque, ephys	■	■		
		Echeveste et al 2020	uncertainty about the orientation of Gabor patches. artificial neural networks	■			
		Haefner et al 2016	uncertainty about the task-relevant orientation, visual task. Artificial neural networks	■			■
		Bányai et al 2019	uncertainty about local features in an image. Macaque ephys	■	■		
	descriptive	Dekleva et al 2016	uncertainty about reaching direction. Macaque, ephys	■	■		■
		Vilares et al 2012	uncertainty about location of source generating dots. Human, fMRI	■			
		Bach et al 2011	uncertainty about proposed gambles. Human, fMRI	■			
		Meyniel 2020	uncertainty about generative probability of binary stimulus. Human, MEG	■	■		■
		Meyniel & Dehaene 2017	uncertainty about generative probability of binary stimulus. Human, fMRI	■	■	■	■
		Payzan-LeNestour et al 2013	uncertainty about reward probabilities. Human, fMRI	■			
		O'Reilly et al 2013	uncertainty about lending location of moving source generating dots. Human, fMRI	■			
		McGuire et al 2014	uncertainty about location of source generating dots. Human, fMRI	■	■		■
		Sedley et al 2016	uncertainty about mean pitch in sound sequence. Human, ephys	■		■	
		Michael et al 2015	uncertainty about the average stimulus (visual, categorization). Human, fMRI	■		■	
		Tomov et al.2020	uncertainty about reward probabilities. Human, fMRI	■	■		■
		Badre et al 2012	uncertainty about reward probability. Human, fMRI	■	■		
		Hsu et al 2005	uncertainty about gambles. Human, fMRI	■	■		
		Lebreton et al 2015	uncertainty about subjective value. Human, fMRI	■	■	■	
		Trudel et al 2021	uncertainty about the variance of stimuli (visual). Human, fMRI	■			
		Blankenstein et al 2017	uncertainty about gambles. Human, fMRI	■			■
		Grinband et al 2006	uncertainty about categorization boundary (visual). Human, fMRI	■	■		
		Tan et al 2016	uncertainty about the forward sensory-motor model. Human, EEG	■			
		Ting et al 2015	uncertainty about the generative proportion of stimuli (visual). Human, fMRI	■			■
aleatoric variability	descriptive	FitzGerald et al 2010	uncertainty about reward delivery (either learned or described). Human, fMRI	■			
		Strange et al 2005	uncertainty about next observation (shapes). Human, fMRI	■	■		
		Monosov & Hikosaka 2013	uncertainty about reward delivery. Macaque, ephys	■	■	■	
		Monosov 2017	uncertainty about reward delivery. Macaque, ephys	■	■	■	
		Monosov et al 2015	uncertainty about reward delivery or reward size. Macaque, ephys	■	■	■	
		Preuschoff et al 2006	uncertainty about reward delivery. Human, fMRI	■		■	
		Fiorillo 2003	uncertainty about reward delivery. Macaque, ephys	■			
		Huettel 2005	uncertainty about next observation (shapes). Human, fMRI	■			
		Muller et al 2019	uncertainty about which option is best (4-arm bandit task). Human, fMRI	■	■		
		Nastase et al 2018	uncertainty about next observation (visual and auditory modalities). Human, fMRI	■		■	
		Tzagarakis et al 2010	uncertainty about which target to reach. MEG human	■			
decision confidence	Process	Hebart et al 2014	motion discrimination task. Human, fMRI	■			■
	descriptive	Gherman & Philiastides 2018	motion discrimination task. Human, EEG	■			■
		Guggenmos et al 2016	orientation task. Human, fMRI	■			■
		Kiani & Shadlen 2009	motion discrimination task. Monkey, ephys	■	■		■
		Kepecs et al 2008	odor discrimination task. Rat, ephys	■	■	■	■
		Lak et al 2014	odor discrimination task. Rat, ephys	■			■
		Komura et al 2013	motion discrimination task. Monkey, ephy	■			■
		Cortese et al 2016	motion discrimination. Human, fMRI	■	■		■
		Gherman & Philiastides 2015	visual object categorization. Human, EEG	■	■		
		So & Stuphorn 2016	confidence about reward delivery. Monkey ephys	■	■		
		Odegaard et al 2018	motion discrimination. Monkey, ephys	■	■		
		Masset et al 2020	Perceptual categorization task (olfactory and auditory stimuli). Rat, ephys	■	■	■	■
		Stern et al 2010	"urns and marbles" task. Human, fMRI	■	■		■
		de Martino et al 2013	confidence about utility-based decision. Human, fMRI	■	■		■
		Bang & Fleming 2018	motion direction categorization. Human, fMRI	■	■		■

Table 1: Example studies on the neural representation of uncertainty assessed with general criteria. Green: tested and passed. Categorization of studies that mention both epistemic uncertainty and aleatoric variability (or decision confidence) is arbitrary and reflects our chosen

emphasis. For the process approach, we distinguish those that assume a probabilistic population code (“PPC”) and sampling-based code (“SBC”).

How do the two approaches compare in terms of satisfying general criteria?

Table 1 reports the analysis of example previous studies on the neural representation of uncertainty. We focus on epistemic uncertainty but also include for comparison studies in which aleatoric variability makes some outcome (e.g. reward delivery) uncertain and studies on decision confidence (Fleming *et al.* 2012; Kepecs & Mainen 2012; Meyniel *et al.* 2015b; Pouget *et al.* 2016). Table 1 enables the reader to inspect how the different approaches (descriptive, process) and the different types (epistemic uncertainty, aleatoric variability, decision confidence) compare in terms of satisfying the criteria. Based on those examples, functionality appears to be currently a strong point of the process approach (when it relies on probabilistic population codes, not on sampling-based codes) in comparison to the descriptive approach, whereas invariance, and to a lesser extent specificity, are more often tested in the descriptive approach than the process approach.

4. Caveats of current approaches and future directions

Equipped with the definitions of process and descriptive approaches (Figure 2) and the evaluation of several studies following each approach (Table 1), we now summarize how the approaches compare in terms of what they enable researchers to do.

Comparing descriptive and process approaches

Prior assumptions. Descriptive and process approaches differ in how much prior knowledge they require. Constructing a neural likelihood function $\mathcal{L}(s;r)$ requires substantial prior knowledge and theoretical motivation. The process approach is thus more easily applicable to brain systems whose functioning is better known, such as the primary visual system (most studies following the process approach involve simple visual inputs and the primary visual system; see Table 1). In comparison, the descriptive approach often requires fewer assumptions. The process approach is consequently better suited to advance research on already sufficiently well known brain systems, while the descriptive approach can be used for more exploratory research.

Moreover, the process approach studies neural representations in which a neural population r jointly represents both a world state s and uncertainty about s. The descriptive approach does not require such joint representations. This restriction to joint representations and the need for much prior knowledge in the process approach is likely to explain different findings between the two approaches. The process approach more often identifies representations of u in sensory regions like the early visual cortex, which are well known for representing visual features (Geurts *et al.* 2022; Hénaff *et al.* 2020; Orbán *et al.* 2016; van Bergen *et al.* 2015; van Bergen & Jehee 2019; Walker *et al.* 2020), whereas the descriptive approach often identifies representations of u in regions that are further from sensory input and its representation, closer to the decision or reporting mechanisms, whether from subcortical structures (Grinband *et al.*

2006; Payzan-LeNestour *et al.* 2013; Ting *et al.* 2015; Vilares *et al.* 2012), prefrontal cortex (Badre *et al.* 2012; Lebreton *et al.* 2015; McGuire *et al.* 2014; Meyniel 2020; Michael *et al.* 2015; Payzan-LeNestour *et al.* 2013; Trudel *et al.* 2021; Vilares *et al.* 2012), parietal cortex (Badre *et al.* 2012; McGuire *et al.* 2014; Meyniel & Dehaene 2017; Michael *et al.* 2015; Payzan-LeNestour *et al.* 2013; Trudel *et al.* 2021) or temporal cortex (Badre *et al.* 2012; Blankenstein *et al.* 2017).

The fact that a joint neural representation of a world state and the associated uncertainty is assumed only in the process approach results in different notions of *representation of uncertainty* across the two approaches. The process approach studies the uncertainty of the neural representation of a world state *s* and the way this uncertainty is coded. The descriptive approach also studies the way the uncertainty about *s* is coded but without assuming that it is coded together with the neural representation of *s*: both aspects can be detached from one another, possibly through downstream computations in the brain.

Functionality and specificity. The process and descriptive approaches also differ in their reliance on behavioral data. The central criterion in the descriptive approach is the sensitivity of *r* to *u*; since *u* can be derived from *I*, descriptive models do not have to include behavioral responses at all (see Figure 2). By contrast, it is essential for process models to explain how the brain transforms a sensory input into a representation of a world state and into a behavioral response; functionality is consequently more central in this approach (see Table 1). Studies in which behavior is not available or difficult to collect are thus better suited to the descriptive approach.

Studies under the process and descriptive approaches differ in their susceptibility to confounding factors. The descriptive approach is more at risk, because the researcher uses *I* or *b* as a proxy for *u* (see Figure 2), raising the concern that a representation of *u* may be spurious, in that *r* is a representation of *I* or *b* rather than of *u*. By contrast, the process approach has an elegant method to demonstrate functional specificity, namely by testing whether *u(r)* makes a difference to the behavioral response *b*, even while controlling for *I*. Yet, concerns about specificity remain for the process approach even when parameters of the input *I* can be ruled out as confounding variables. Representations of uncertainty may still be confounded by behavioral features or processes internal to the brain, such as attention.

Overall, the two approaches lend themselves to different interpretations. The process model is causal in spirit, going from the *I* to *r* to *b*. Functionality plays a key role in this approach, making causal interpretations more legitimate than with the descriptive approach which is essentially correlative.

Sources of uncertainty. The process and descriptive approach are also sensitive to different sources of uncertainty. Descriptive studies derive *u* from *I* based on risky assumptions that may not be true of the brain and limit themselves to external (observer-independent) sources of uncertainty. When the descriptive approach is used in combination with reports of uncertainty (or confidence), uncertainty becomes observer-dependent. However, the reported uncertainty may differ from the uncertainty used at earlier stages of information processing in the brain (Pouget *et al.* 2016; Zylberberg *et al.* 2014). By contrast, the process approach always derives *u* from *r* itself, capturing observer-dependent computations and internal sources of uncertainty like neural noise.

Neural code. The two approaches also differ in assumptions about the complexity of the neural code of uncertainty. In the descriptive approach, codes are usually assumed to be simple: many studies of this kind only look for linear codes (monotonic changes in average activity or in a weighted sum of activities as a function uncertainty); representations with such linear codes have been termed explicit representations (DiCarlo *et al.* 2012; Kriegeskorte & Diedrichsen 2019). Studies following the process approach are open to the use of non-linear computations, which are often employed to derive *u* from *r*, e.g. when reading out the standard deviation of the decoded distribution (Geurts *et al.* 2022; van Bergen *et al.* 2015; van Bergen & Jehee 2019), or the standard deviation of neural activity (Orbán *et al.* 2016), or when using artificial neural networks (Walker *et al.* 2020).

The two approaches are also differently informative about the way uncertainty is computed by the brain. In the descriptive approach the researcher uses uncertainty as an independent variable, not derived from brain activity. The process approach makes a step toward unraveling the neural computations of uncertainty by proposing a generative model of *r* and in many studies, making a connection between neural representations of uncertainty and downstream computations or behavior. Note that in the process approach, the uncertainties estimated by the researcher and the brain will differ if there are discrepancies between the actual neural computations and the ones assumed by the researcher, and between the *r* used by the brain and measured by the researcher.

Setting goals for future research

Given that the process and descriptive approaches have different limitations and advantages, they could be used in synergy. One possibility is to leverage our knowledge of early sensory cortices to decode uncertainty in a perceptual task using the process approach, and then use the decoded uncertainty as an input to the descriptive approach in order to unravel other parts of the brain that could represent this uncertainty. Such combined analysis could shed light on how uncertainty-based computations unfold in the brain, from early sensory processing to subsequent decision making or learning.

Understanding how the brain extracts and uses uncertainty can also be achieved by further investigation of the functional aspect of representation. If uncertainty is used only in a given context (e.g. uncertainty about color, not shape, is relevant for color-based categorization (Michael *et al.* 2015)) or for different goals (e.g. guide the decision to wager (Kiani & Shadlen 2009) or to update prior estimates (Meyniel & Dehaene 2017)), then some aspects of its representation are expected to change. Manipulating the task relevance of uncertainty is thus a promising avenue to explore the function of the representation of uncertainty. In particular, it would be useful to distinguish representations of uncertainty that are automatic and occur independently of task demands from those that are task-dependent.

As the field matures, a switch from single-model testing to the comparison of different models (i.e. different *f*s and *p*s in Fig 2) would be valuable to narrow down the neural codes of uncertainty. For instance, Walker and colleagues (Walker *et al.* 2020) showed that assuming independent Poisson distribution of neural responses provides a good account of the neural coding of uncertainty, but one that is worse than a less constrained neural likelihood function (estimated with a deep neural network).

We have stressed that uncertainty can be about different things (e.g. orientation of a grating (Walker *et al.* 2020), color (Michael *et al.* 2015), the next outcome (Monosov & Hikosaka 2013), probability of a event (Meyniel & Dehaene 2017)) and have multiple origins (e.g. prior knowledge, current input). Open questions regard both the extent to which any representation of uncertainty is invariant to what uncertainty is about and to its origin.

Manipulating prior expectations could help to tackle the pervasive issue of specificity: posterior uncertainty depends on both the current input and the prior, but most studies focus only on the former. Manipulating the priors thus enables researchers to partly decorrelate posterior uncertainty from the current input. Some previous studies included the manipulation of priors in their experimental design (Ting *et al.* 2015; Vilares *et al.* 2012) but with the aim of comparing the encoding of the prior and current likelihood. Beyond the mere methodological interest regarding specificity, systematic manipulation of priors (as in previous behavioral studies (Acerbi *et al.* 2014)) would also be useful in order to study at which stage prior and current uncertainties are combined in the brain when processing the current input.

In conclusion, we propose that studies on the neural representation of uncertainty can be distinguished based on whether the researcher treats uncertainty as a variable derived from the input or behavior (descriptive approach) or neural activity (process approach). This distinction results in an emphasis on different general criteria used to assess empirically neural representations of uncertainty. The descriptive approach primarily tests for relations between neural activity and the independently-derived uncertainty (relying more on sensitivity, specificity and invariance) whereas the process approach seeks to relate neural representations of uncertainty to downstream computations and behavior (relying more on functionality). Consequently, the two approaches also differ in the assumptions they require and the type of findings they uncover, and could both have a great potential if used synergistically.

Glossary

Uncertainty: measure of how indeterminate the world state *s* is given the information an observer has about *s* (under high uncertainty, many different world states seem plausible)

Aleatoric variability: randomness in the outcome of a process (synonyms: indeterminacy, stochasticity)

Neural likelihood function: refers to the probability of an observed neural activity pattern $r_{observed}$ given some world state *s*; s in the argument of this function that $r_{observed}$ parameterizes.

Generative model: a model that specifies how effects follow their causes. It is often characterized only in terms of statistical dependencies, e.g. $p(r|s)$. We often distinguish between *external* and *neural* generative models to highlight that the former is about the input and the latter is about neural activity.

Descriptive approach: the researcher treats uncertainty as an independent variable (derived from the input and the subject's behavior) and tests for a given relationship between uncertainty and some neural activity.

Process approach: the researcher treats uncertainty as a dependent variable by assuming a particular neural code for *s* in *r* and inferring the uncertainty about *s* from the observed neural

activity *r*; the researcher then tests for a relationship between this uncertainty and the subject's behavior (and sometimes the input).

Neural representation of s: neural activity pattern used by the brain to convey information about *s*.

Ideal observer: a model of how the world state *s* is inferred from the input *I* by relying optimally (i.e. with Bayes' rule) on the true generative model of *I*.

References (including highlights)

Acerbi, L., Vijayakumar, S., & Wolpert, D. M. (2014). On the Origins of Suboptimality in Human Probabilistic Inference. *PLoS Comput Biol*, **10**(6), e1003661.

Adler, W. T., & Ma, W. J. (2018). Comparing Bayesian and non-Bayesian accounts of human confidence reports. *PLOS Computational Biology*, **14**(11), e1006572.

Bach, D. R., & Dolan, R. J. (2012). Knowing how much you don't know: a neural organization of uncertainty estimates. *Nature Reviews Neuroscience*, **13**(8), 572–586.

Bach, D. R., Hulme, O., Penny, W. D., & Dolan, R. J. (2011). The Known Unknowns: Neural Representation of Second-Order Uncertainty, and Ambiguity. *The Journal of Neuroscience*, **31**(13), 4811–4820.

Badre, D., Doll, B. B., Long, N. M., & Frank, M. J. (2012). Rostrolateral Prefrontal Cortex and Individual Differences in Uncertainty-Driven Exploration. *Neuron*, 73(3), 595–607.

Example of descriptive approach that uses an ideal-observer model of the learning process to infer uncertainty in a task. Finds evidence of a functional role for uncertainty (here, in terms of exploration).

Ballard, D. H. (2015). *Brain Computation as Hierarchical Abstraction*, MIT Press. Retrieved from https://mitpress.mit.edu/books/brain-computation-hierarchical-abstraction

Bang, D., & Fleming, S. M. (2018). Distinct encoding of decision confidence in human medial prefrontal cortex. *Proceedings of the National Academy of Sciences of the United States of America*, **115**(23), 6082–6087.

Bányai, M., Lazar, A., Klein, L., … Orbán, G. (2019). Stimulus complexity shapes response correlations in primary visual cortex. *Proceedings of the National Academy of Sciences*, 116(7), 2723–2732.

Example of process approach that uses a sampling-based code and shows that the co-variance of neural activity in a population of neurons can be explained by hierarchical inference with a prominent impact of the image higher level features even in regions tuned to local features like the primary visual cortex.

Barlow, H. B. (1969). Pattern Recognition and the Responses of Sensory Neurons*. *Annals of the New York Academy of Sciences*, **156**(2), 872–881.

Barthelmé, S., & Mamassian, P. (2009). Evaluation of objective uncertainty in the visual system. *PLoS Computational Biology*, **5**(9), e1000504.

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), 30–39.
Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, **10**(9), 1214–1221.
Blankenstein, N. E., Peper, J. S., Crone, E. A., & van Duijvenvoorde, A. C. K. (2017). Neural Mechanisms Underlying Risk and Ambiguity Attitudes. *Journal of Cognitive Neuroscience*, **29**(11), 1845–1859.
Cabrera, C. A., Lu, Z.-L., & Dosher, B. A. (2015). Separating decision and encoding noise in signal detection tasks. *Psychological Review*, **122**(3), 429–460.
Cortese, A., Amano, K., Koizumi, A., Kawato, M., & Lau, H. (2016). Multivoxel neurofeedback selectively modulates confidence without changing perceptual performance. *Nature Communications*, **7**(1), 13669.
Dabney, W., Kurth-Nelson, Z., Uchida, N., … Botvinick, M. (2020). A distributional code for value in dopamine-based reinforcement learning. *Nature*, **577**(7792), 671–675.
Dayan, P., & Abbott, L. F. (2005). *Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems*, 1st edn, The MIT Press.
De Martino, B., Fleming, S. M., Garrett, N., & Dolan, R. J. (2013). Confidence in value-based choice. *Nature Neuroscience*, **16**(1), 105–110.
Dekleva, B. M., Ramkumar, P., Wanda, P. A., Kording, K. P., & Miller, L. E. (2016). Uncertainty leads to persistent effects on reach representations in dorsal premotor cortex. *ELife*, **5**, e14316.
Deneve, S. (2008). Bayesian Spiking Neurons I: Inference. *Neural Computation*, **20**(1), 91–117.
Deneve, S., Latham, P. E., & Pouget, A. (1999). Reading population codes: a neural implementation of ideal observers. *Nature Neuroscience*, **2**(8), 740–745.
Deroy, O., Spence, C., & Noppeney, U. (2016). Metacognition in Multisensory Perception. *Trends in Cognitive Sciences*, **20**(10), 736–747.
Devkar, D., Wright, A. A., & Ma, W. J. (2017). Monkeys and humans take local uncertainty into account when localizing a change. *Journal of Vision*, **17**(11), 4.
DiCarlo, J. J., Zoccolan, D., & Rust, N. C. (2012). How does the brain solve visual object recognition? *Neuron*, **73**(3), 415–434.
Echeveste, R., Aitchison, L., Hennequin, G., & Lengyel, M. (2020). Cortical-like dynamics in recurrent circuits optimized for sampling-based probabilistic inference. *Nature Neuroscience*, 23(9), 1138–1149.
Shows that an artificial neural network can be trained to emit spikes that correspond to samples from a posterior distribution of some feature of the inpu. Although not trained to so do, the artificial network shows dynamics similar to actual neural networks.
Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, **415**(6870), 429–433.
Faisal, A. A., Selen, L. P. J., & Wolpert, D. M. (2008). Noise in the nervous system. *Nature Reviews Neuroscience*, **9**(4), 292–303.
Festa, D., Aschner, A., Davila, A., Kohn, A., & Coen-Cagli, R. (2021). Neuronal variability reflects probabilistic inference tuned to natural image statistics. *Nature Communications*, **12**(1), 3635.
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.
Fiorillo, C. D. (2003). Discrete Coding of Reward Probability and Uncertainty by Dopamine Neurons. *Science*, **299**(5614), 1898–1902.
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), 119–130.
FitzGerald, T. H. B., Seymour, B., Bach, D. R., & Dolan, R. J. (2010). Differentiable Neural Substrates for Learned and Described Value and Risk. *Current Biology*, **20**(20), 1823–1829.
Fleming, S. M., & Daw, N. D. (2017). Self-evaluation of decision-making: A general Bayesian framework for metacognitive computation. *Psychological Review*, **124**(1), 91–114.
Fleming, S. M., Dolan, R. J., & Frith, C. D. (2012). Metacognition: computation, biology and function. *Philosophical Transactions of the Royal Society B: Biological Sciences*, **367**(1594), 1280–1286.
Friston, K., Ashburner, J., Kiebel, S., Nichols, T., & Penny, W. (2007). *Statistical parametric mapping: the analysis of functional brain images*, Academic Press, London.
Geurts, L. S., Cooke, J. R. H., van Bergen, R. S., & Jehee, J. F. M. (2022). Subjective confidence reflects representation of Bayesian probability in cortex. *Nature Human Behaviour*, 1–12.
Gherman, S., & Philiastides, M. G. (2015). Neural representations of confidence emerge from the process of decision formation during perceptual choices. *NeuroImage*, **106**, 134–143.
Gherman, S., & Philiastides, M. G. (2018). Human VMPFC encodes early signatures of confidence in perceptual decisions. *ELife*, **7**, e38293.
Grinband, J., Hirsch, J., & Ferrera, V. P. (2006). A Neural Representation of Categorization Uncertainty in the Human Brain. *Neuron*, **49**(5), 757–763.
Guggenmos, M., Wilbertz, G., Hebart, M. N., & Sterzer, P. (2016). Mesolimbic confidence signals guide perceptual learning in the absence of external feedback. *ELife*, **5**, e13388.
Haefner, R. M., Berkes, P., & Fiser, J. (2016). Perceptual Decision-Making as Probabilistic Inference by Neural Sampling. *Neuron*, **90**(3), 649–660.
Hampton, R. R. (2001). Rhesus monkeys know when they remember. *Proceedings of the National Academy of Sciences*, **98**(9), 5359–5362.
Haxby, J. V., Connolly, A. C., & Guntupalli, J. S. (2014). Decoding Neural Representational Spaces Using Multivariate Pattern Analysis. *Annual Review of Neuroscience*, **37**(1), 435–456.
Haxby, J. V., Gobbini, M. I., Furey, M. L., Ishai, A., Schouten, J. L., & Pietrini, P. (2001). Distributed and Overlapping Representations of Faces and Objects in Ventral Temporal Cortex. *Science*, **293**(5539), 2425–2430.
Haynes, J.-D. (2015). A Primer on Pattern-Based Approaches to fMRI: Principles, Pitfalls, and Perspectives. *Neuron*, **87**(2), 257–270.
Hebart, M. N., Schriever, Y., Donner, T. H., & Haynes, J.-D. (2014). The Relationship between Perceptual Decision Variables and Confidence in the Human Brain. *Cerebral Cortex*

(New York, N.Y.: 1991). doi:10.1093/cercor/bhu181

Helmholtz, H. (1867). *Handbuch der Physiologischen Optik*. (G. Karsten, Ed.), Leipzig: Leopold Voss.

Hénaff, O. J., Boundy-Singer, Z. M., Meding, K., Ziemba, C. M., & Goris, R. L. T. (2020). Representation of visual uncertainty through neural gain variability. *Nature Communications*, **11**(1), 2513.

Hoyer, P. O., & Hyvärinen, A. (2002). Interpreting Neural Response Variability as Monte Carlo Sampling of the Posterior. In *Advances in Neural Information Processing Systems*, MIT Press, p. 2002.

An influential article that proposed that neural activity could be explained with a sampling-based code.

Hsu, M., Bhatt, M., Adolphs, R., Tranel, D., & Camerer, C. F. (2005). Neural systems responding to degrees of uncertainty in human decision-making. *Science (New York, N.Y.)*, 310(5754), 1680–1683.

Presented a distinction between epistemic uncertainty and aleatoric variability (referred to as ambiguity and risk respectively in behavioral economics), whose fMRI correlates are anatomically segregated in the human brain.

Hubel, D. H., & Wiesel, T. N. (1959). Receptive fields of single neurones in the cat's striate cortex. *The Journal of Physiology*, **148**(3), 574–591.

Hubel, D. H., & Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *The Journal of Physiology*, **160**(1), 106–154.

Huettel, S. A. (2005). Decisions under Uncertainty: Probabilistic Context Influences Activation of Prefrontal and Parietal Cortices. *Journal of Neuroscience*, **25**(13), 3304–3311.

Huth, A. G., de Heer, W. A., Griffiths, T. L., Theunissen, F. E., & Gallant, J. L. (2016). Natural speech reveals the semantic maps that tile human cerebral cortex. *Nature*, **532**(7600), 453–458.

Iglesias, S., Mathys, C., Brodersen, K. H., … Stephan, K. E. (2013). Hierarchical prediction errors in midbrain and basal forebrain during sensory learning. *Neuron*, **80**(2), 519–530.

Jazayeri, M., & Movshon, J. A. (2006). Optimal representation of sensory information by neural populations. *Nature Neuroscience*, **9**(5), 690–696.

Kepecs, A., & Mainen, Z. F. (2012). A computational framework for the study of confidence in humans and animals. *Philosophical Transactions of the Royal Society B: Biological Sciences*, **367**(1594), 1322–1337.

Kepecs, A., Uchida, N., Zariwala, H. A., & Mainen, Z. F. (2008). Neural correlates, computation and behavioural impact of decision confidence. *Nature*, **455**(7210), 227–231.

Kersten, D., Mamassian, P., & Yuille, A. (2004). Object Perception as Bayesian Inference. *Annual Review of Psychology*, **55**(1), 271–304.

Kiani, R., & Shadlen, M. N. (2009). Representation of confidence associated with a decision by neurons in the parietal cortex. *Science (New York, N.Y.)*, **324**(5928), 759–764.

Knill, D. C., & Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends in Neurosciences*, **27**(12), 712–719.

Komura, Y., Nikkuni, A., Hirashima, N., Uetake, T., & Miyamoto, A. (2013). Responses of

pulvinar neurons reflect a subject's confidence in visual categorization. *Nature Neuroscience*, **16**(6), 749–755.

Koriat, A., Sheffer, L., & Ma'ayan, H. (2002). Comparing objective and subjective learning curves: Judgments of learning exhibit increased underconfidence with practice. *Journal of Experimental Psychology: General*, **131**(2), 147–162.

Kriegeskorte, N., & Diedrichsen, J. (2019). Peeling the Onion of Brain Representations. *Annual Review of Neuroscience*, **42**(1), 407–432.

Lak, A., Costa, G. M., Romberg, E., Koulakov, A. A., Mainen, Z. F., & Kepecs, A. (2014). Orbitofrontal cortex is required for optimal waiting based on decision confidence. *Neuron*, **84**(1), 190–201.

Lange, R. D., Shivkumar, S., Chattoraj, A., & Haefner, R. M. (2021). *Bayesian Encoding and Decoding as Distinct Perspectives on Neural Coding*. Retrieved from https://www.biorxiv.org/content/10.1101/2020.10.14.339770v2

Lebreton, M., Abitbol, R., Daunizeau, J., & Pessiglione, M. (2015). Automatic integration of confidence in the brain valuation signal. *Nature Neuroscience*, **18**(8), 1159–1167.

Lee, T. S., & Mumford, D. (2003). Hierarchical Bayesian inference in the visual cortex. *Journal of the Optical Society of America. A, Optics, Image Science, and Vision*, **20**(7), 1434–1448.

Li, H.-H., Sprague, T. C., Yoo, A. H., Ma, W. J., & Curtis, C. E. (2021). Joint representation of working memory and uncertainty in human cortex. *Neuron*, **109**(22), 3699-3712.e6.

Ma, W. J., Beck, J. M., Latham, P. E., & Pouget, A. (2006). Bayesian inference with probabilistic population codes. *Nature Neuroscience*, **9**(11), 1432–1438.

Ma, W. J., & Jazayeri, M. (2014). Neural coding of uncertainty and probability. *Annual Review of Neuroscience*, 37, 205–220.

Introduced the concept of probabilistic population code as the idea that the representation of probability distribution over a latent world state by a population of neurons, conferred by an internal model of neural variability, allows certain Bayesian computations to be implemented by simple neural operations.

Masset, P., Ott, T., Lak, A., Hirokawa, J., & Kepecs, A. (2020). Behavior- and Modality-General Representation of Confidence in Orbitofrontal Cortex. *Cell*, 182(1), 112-126.e18.

Studies decision confidence in rats using waiting times as a proxy and identifies a neural representation of decision confidence in the orbitofrontal cortex that passes the tests of sensitivity, specificity (with respect to the features of the input), invariance (to the sensory modality) and functionality (correlation with learning).

Mathys, C. D., Lomakina, E. I., Daunizeau, J., … Stephan, K. E. (2014). Uncertainty in perception and the Hierarchical Gaussian Filter. *Frontiers in Human Neuroscience*, **8**. doi:10.3389/fnhum.2014.00825

McGuire, J. T., Nassar, M. R., Gold, J. I., & Kable, J. W. (2014). Functionally Dissociable Influences on Learning Rate in a Dynamic Environment. *Neuron*, **84**(4), 870–881.

Meyniel, F. (2020). Brain dynamics for confidence-weighted learning. *PLOS Computational Biology*, **16**(6), e1007935.

Meyniel, F., & Dehaene, S. (2017). Brain networks for confidence weighting and hierarchical inference during probabilistic learning. *Proceedings of the National*

***Academy of Sciences*, 201615773.**

Example of descriptive approach that uses an ideal-observer model of the input to derive uncertainty about a probability. The study reports fMRI correlates of this uncertainty distinct from correlates of confounding factors like aleatoric variability and surprise.

Meyniel, F., Schlunegger, D., & Dehaene, S. (2015a). The Sense of Confidence during Probabilistic Learning: A Normative Account. *PLoS Comput Biol*, **11**(6), e1004305.

Meyniel, F., Sigman, M., & Mainen, Z. F. (2015b). Confidence as Bayesian Probability: From Neural Origins to Behavior. *Neuron*, **88**(1), 78–92.

Michael, E., de Gardelle, V., Nevado-Holgado, A., & Summerfield, C. (2015). Unreliable Evidence: 2 Sources of Uncertainty During Perceptual Choice. *Cerebral Cortex*, 25(4), 937–947.

Example of descriptive approach that uses a categorization task based on either shape or color from trial-to-trial and identifies representations of uncertainty about the decision that are invariant to the perceptual feature decision is based on (shape or color).

Middlebrooks, P. G., & Sommer, M. A. (2012). Neuronal correlates of metacognition in primate frontal cortex. *Neuron*, **75**(3), 517–530.

Monosov, I. E. (2017). Anterior cingulate is a source of valence-specific information about value and uncertainty. *Nature Communications*, **8**(1), 134.

Monosov, I. E., & Hikosaka, O. (2013). Selective and graded coding of reward uncertainty by neurons in the primate anterodorsal septal region. *Nature Neuroscience*, **16**(6), 756–762.

Monosov, I. E., Leopold, D. A., & Hikosaka, O. (2015). Neurons in the Primate Medial Basal Forebrain Signal Combined Information about Reward Uncertainty, Value, and Punishment Anticipation. *Journal of Neuroscience*, **35**(19), 7443–7459.

Muller, T. H., Mars, R. B., Behrens, T. E., & O'Reilly, J. X. (2019). Control of entropy in neural models of environmental state. *Elife*, **8**, e39404.

Naselaris, T., Kay, K. N., Nishimoto, S., & Gallant, J. L. (2011). Encoding and decoding in fMRI. *NeuroImage*, **56**(2), 400–410.

Nastase, S. A., Davis, B., & Hasson, U. (2018). Cross-modal and non-monotonic representations of statistical regularity are encoded in local neural response patterns. *NeuroImage*, **173**, 509–517.

Necker, L. A. (1832). Observations on some remarkable optical phænomena seen in Switzerland; and on an optical phænomenon which occurs on viewing a figure of a crystal or geometrical solid. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, **1**(5), 329–337.

Norton, E. H., Acerbi, L., Ma, W. J., & Landy, M. S. (2019). Human online adaptation to changes in prior probability. *PLOS Computational Biology*, **15**(7), e1006681.

Odegaard, B., Grimaldi, P., Cho, S. H., Peters, M. A. K., Lau, H., & Basso, M. A. (2018). Superior colliculus neuronal ensemble activity signals optimal rather than subjective confidence. *Proceedings of the National Academy of Sciences*, **115**(7), E1588–E1597.

Orbán, G., Berkes, P., Fiser, J., & Lengyel, M. (2016). Neural Variability and Sampling-Based Probabilistic Representations in the Visual Cortex. *Neuron*, 92(2),

530–543.

Example of process approach that uses a sampling-based code and finds that neural variability (in spiking activity and membrane potential) changes along features of visual input related to uncertainty (e.g. it quenches at the stimulus onset, decreases with contrast and aperture).

O'Reilly, J. X. (2013). Making predictions in a changing world—inference, uncertainty, and learning. *Frontiers in Decision Neuroscience*, **7**, 105.

O'Reilly, J. X., Jbabdi, S., Rushworth, M. F. S., & Behrens, T. E. J. (2013). Brain Systems for Probabilistic and Dynamic Prediction: Computational Specificity and Integration. *PLoS Biol*, **11**(9), e1001662.

Park, I. M., Meister, M. L. R., Huk, A. C., & Pillow, J. W. (2014). Encoding and decoding in parietal cortex during sensorimotor decision-making. *Nature Neuroscience*, **17**(10), 1395–1403.

Payzan-LeNestour, E., Dunne, S., Bossaerts, P., & O'Doherty, J. P. (2013). The neural representation of unexpected uncertainty during value-based decision making. *Neuron*, **79**(1), 191–201.

Peirce, C. S., & Jastrow, J. (1884). On Small Differences in Sensation. *Memoirs of the National Academy of Sciences*, **3**, 75–83.

Pillow, J. W., Shlens, J., Paninski, L., … Simoncelli, E. P. (2008). Spatio-temporal correlations and visual signalling in a complete neuronal population. *Nature*, **454**(7207), 995–999.

Poldrack, R. A., Huckins, G., & Varoquaux, G. (2020). Establishment of Best Practices for Evidence for Prediction: A Review. *JAMA Psychiatry*, **77**(5), 534–540.

Pouget, A., Dayan, P., & Zemel, R. S. (2003). Inference and computation with population codes. *Annual Review of Neuroscience*, **26**(1), 381–410.

Pouget, A., Drugowitsch, J., & Kepecs, A. (2016). Confidence and certainty: distinct probabilistic quantities for different goals. *Nature Neuroscience*, **19**(3), 366–374.

Preuschoff, K., Bossaerts, P., & Quartz, S. R. (2006). Neural Differentiation of Expected Reward and Risk in Human Subcortical Structures. *Neuron*, **51**(3), 381–390.

Qamar, A. T., Cotton, R. J., George, R. G., … Ma, W. J. (2013). Trial-to-trial, uncertainty-based adjustment of decision boundaries in visual categorization. *Proceedings of the National Academy of Sciences*, **110**(50), 20332–20337.

Rademaker, R. L., Tredway, C. H., & Tong, F. (2012). Introspective judgments predict the precision and likelihood of successful maintenance of visual working memory. *Journal of Vision*, **12**(13), 21.

Rahnev, D., & Denison, R. N. (2018). Suboptimality in perceptual decision making. *Behavioral and Brain Sciences*, **41**. doi:10.1017/S0140525X18000936

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.

Rushworth, M. F. S., & Behrens, T. E. J. (2008). Choice, uncertainty and value in prefrontal and cingulate cortex. *Nature Neuroscience*, **11**(4), 389–397.

Sahani, M., & Dayan, P. (2003). Doubly Distributional Population Codes: Simultaneous Representation of Uncertainty and Multiplicity. *Neural Computation*, **15**(10), 2255–2279.

Schmack, K., Bosc, M., Ott, T., Sturgill, J. F., & Kepecs, A. (2021). Striatal dopamine mediates hallucination-like perception in mice. *Science*, **372**(6537). doi:10.1126/science.abf4740

Sedley, W., Gander, P. E., Kumar, S., … Griffiths, T. D. (2016). Neural signatures of perceptual inference. *ELife*, **5**, e11476.

Shannon, C. (1948). A Mathematical Theory of Communication. *The Bell System Technical Journal*, **27**(3), 379–423.

So, N., & Stuphorn, V. (2016). Supplementary Eye Field Encodes Confidence in Decisions Under Risk. *Cerebral Cortex*, **26**(2), 764–782.

Sohn, H., & Narain, D. (2021). Neural implementations of Bayesian inference. *Current Opinion in Neurobiology*, **70**, 121–129.

Stern, E. R., Gonzalez, R., Welsh, R. C., & Taylor, S. F. (2010). Updating Beliefs for a Decision: Neural Correlates of Uncertainty and Underconfidence. *Journal of Neuroscience*, **30**(23), 8032–8041.

Strange, B. A., Duggins, A., Penny, W., Dolan, R. J., & Friston, K. J. (2005). Information theory, novelty and hippocampal responses: unpredicted or unpredictable? *Neural Networks*, **18**(3), 225–230.

Tan, H., Wade, C., & Brown, P. (2016). Post-Movement Beta Activity in Sensorimotor Cortex Indexes Confidence in the Estimations from Internal Models. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, **36**(5), 1516–1528.

Ting, C.-C., Yu, C.-C., Maloney, L. T., & Wu, S.-W. (2015). Neural mechanisms for integrating prior knowledge and likelihood in value-based probabilistic inference. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, **35**(4), 1792–1805.

Todorov, E. (2004). Optimality principles in sensorimotor control. *Nature Neuroscience*, **7**(9), 907–915.

Tomov, M. S., Truong, V. Q., Hundia, R. A., & Gershman, S. J. (2020). Dissociable neural correlates of uncertainty underlie different exploration strategies. *Nature Communications*, **11**(1), 2371.

Trommershäuser, J., Maloney, L. T., & Landy, M. S. (2008). Decision making, movement planning and statistical decision theory. *Trends in Cognitive Sciences*, **12**(8), 291–297.

Trudel, N., Scholl, J., Klein-Flügge, M. C., … Rushworth, M. F. S. (2021). Polarity of uncertainty representation during exploration and exploitation in ventromedial prefrontal cortex. *Nature Human Behaviour*, **5**(1), 83–98.

Tzagarakis, C., Ince, N. F., Leuthold, A. C., & Pellizzer, G. (2010). Beta-band activity during motor planning reflects response uncertainty. *The Journal of Neuroscience: The Official Journal of the Society for Neuroscience*, **30**(34), 11270–11277.

van Bergen, R. S., & Jehee, J. F. M. (2019). Probabilistic Representation in Human Visual Cortex Reflects Uncertainty in Serial Decisions. *Journal of Neuroscience*, **39**(41), 8164–8176.

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.

Example of process approach that uses a probabilistic population code estimated in a data-driven manner by means of a generalized linear model. The uncertainty derived from fMRI activity correlates with behavioral variability.

Vilares, I., Howard, J. D., Fernandes, H. L., Gottfried, J. A., & Kording, K. P. (2012). Differential Representations of Prior and Likelihood Uncertainty in the Human Brain. *Current Biology*, 22(18), 1641–1648.

Example of descriptive approach that used specific features of the input (scatter) as a proxy to uncertainty (about the location of a cloud of dots). The fMRI correlates of this uncertainty are distinct from prior uncertainty.

Walker, E. Y., Cotton, R. J., Ma, W. J., & Tolias, A. S. (2020). A neural basis of probabilistic computation in visual cortex. *Nature Neuroscience*, 23(1), 122–129.

Example of the process approach that uses a probabilistic population code estimated in a data-driven manner by means of an artificial neural network. The uncertainty derived from multi-unit recordings accounts for the monkey choices.

Yeon, J., & Rahnev, D. (2020). The suboptimality of perceptual decision making with multiple alternatives. *Nature Communications*, **11**(1), 3857.

Zhou, Y., Acerbi, L., & Ma, W. J. (2020). The role of sensory uncertainty in simple contour integration. *PLOS Computational Biology*, **16**(11), e1006308.

Zylberberg, A., Fetsch, C. R., & Shadlen, M. N. (2016). The influence of evidence volatility on choice, reaction time and confidence in a perceptual decision. *ELife*, **5**, e17688.

Zylberberg, A., Roelfsema, P. R., & Sigman, M. (2014). Variance misperception explains illusions of confidence in simple perceptual decisions. *Consciousness and Cognition*, **27**, 246–253.