Cognition 217 (2021) 104903

Contents lists available at ScienceDirect

Cognition

journal homepage: www.elsevier.com/locate/cognit

Original Articles

What kind of empirical evidence is needed for probabilistic mental representations? An example from visual perception

Ömer Dağlar Tanrıkulu [a,*], Andrey Chetverikov [a,b,d], Sabrina Hansmann-Roth [a], Árni Kristjánsson [a,c]

[a] Faculty of Psychology, School of Health Sciences, University of Iceland, Reykjavik, Iceland
[b] Visual Computation Lab, Center for Cognitive Neuroimaging, Donders Institute for Brain, Cognition and Behavior, Radboud University, Nijmegen, the Netherlands
[c] School of Psychology, National Research University Higher School of Economics, Moscow, Russia
[d] Cognitive Research Lab, Russian Academy of National Economy and Public Administration, Moscow, Russia

ARTICLE INFO

Keywords:
Probability
Representation
Perception
Cognitive science
Vision
Bayesian

ABSTRACT

Recent accounts of perception and cognition propose that the brain represents information probabilistically. While this assumption is common, empirical support for such probabilistic representations in perception has recently been criticized. Here, we evaluate these criticisms and present an account based on a recently developed psychophysical methodology, Feature Distribution Learning (FDL), which provides promising evidence for probabilistic representations by avoiding these criticisms. The method uses priming and role-reversal effects in visual search. Observers' search times reveal the structure of perceptual representations, in which the probability distribution of distractor features is encoded.

We explain how FDL results provide evidence for a stronger notion of representation that relies on structural correspondence between stimulus uncertainty and perceptual representations, rather than a mere co-variation between the two. Moreover, such an account allows us to demonstrate what kind of empirical evidence is needed to support probabilistic representations as posited in current probabilistic Bayesian theories of perception.

1. Introduction

Within cognitive science the mind is considered to be an information processing system that makes inferences about the external states of the world using information from the senses or memory. However, this information is generally noisy and incomplete. One of the main challenges for cognition is therefore to make reliable inferences about the world in the face of uncertainty. This entails the idea that our brains perform probabilistic calculations involving uncertainty. With advances in computer science and mathematical modeling, probabilistic approaches to cognition involving Bayesian statistics have become a unifying framework for studying human cognition (for reviews see: Chater, Tenenbaum, & Yuille, 2006; Griffiths, Chater, Kemp, Perfors, & Tenenbaum, 2010).

These probabilistic models and approaches are particularly advanced and successful in the field of visual perception (e.g., Kersten & Yuille, 2003; Mamassian, Landy, & Maloney, 2002). A crucial assumption of these Bayesian theories is that the brain represents information probabilistically. For example, information is considered to be represented as a conditional probability density function of a set of hypotheses about a distal stimulus rather than as a single estimate of that stimulus. While a large amount of experimental results is consistent with probabilistic representations, experiments that directly investigate how information is represented by the brain are scarce (Knill & Pouget, 2004). Carefully designed experiments are needed to elevate this claim to more than simply an assumption in the field.

Recently, empirical support in favour of probabilistic representations in visual perception has been strongly criticized (Rahnev, 2017; Block, 2018; Yeon & Rahnev, 2020, see discussion in Rahnev, Block, Denison, & Jehee, 2021). In Section 2, we provide a short review of these recent criticisms, which involve the claim that proposed empirical evidence for probabilistic representations in the literature can be explained by positing non-probabilistic representations. While we agree with such criticisms, in Section 3 we present a recently developed psychophysical methodology, Feature Distribution Learning (FDL, Chetverikov, Campana, & Kristjánsson, 2016, 2017a), and argue in Section 4 that the FDL

* Corresponding author.
E-mail address: daglar83@gmail.com (Ö.D. Tanrıkulu).

https://doi.org/10.1016/j.cognition.2021.104903
Received 10 February 2021; Received in revised form 31 August 2021; Accepted 6 September 2021
Available online 14 September 2021

method enables the study of perceptual representations by avoiding the methodological criticisms of experimental studies that attempt to provide empirical evidence for probabilistic representations. We discuss the first criterion for providing evidence for probabilistic representations, which is to demonstrate that representations that involve probability distributions are not imposed on the task by the experimenter, but instead generated by the brain to be utilized later. In Section 5, we propose a second criterion to provide empirical evidence for probabilistic representations. According to our account, experimental results that demonstrate the utilization of correlations between the internal states of the visual system and stimulus uncertainty will not suffice as evidence for probabilistic representations as they are defined in the empirical literature, whereas demonstrating utilization of structural correspondence between the two would. Subsequently, we demonstrate how the results obtained with the FDL method constitute a prime example of such evidence. In Section 6, we present our overall conclusions and discuss unanswered questions that warrant further investigation.

Our main argument here is mainly built on Block (2018) and Rahnev's (2017) criticisms of the notion of probabilistic representations that incorporate probability distributions over possible estimates of a visual feature, which we take to be the notion posited in perceptual sciences. One can argue against Block's criticism by appealing to different notions of representation that incorporate probabilities in different ways (e.g., Gross, 2020; Shea, 2020; Shea & Frith, 2019). However, here, we mostly agree with Block's or Rahnev's criticisms; and we use them as steppingstones to elaborate on what satisfactory empirical evidence for probabilistic representations should look like, by presenting a particular psychophysical method (and the results obtained from this method) as a prime example of this.

2. Criticisms of experiments supporting probabilistic representations

Recent criticisms of empirical evidence for probabilistic representations can be grouped into two categories. The first is directed at the design of perceptual studies providing evidence for probabilistic representations, whereas the second category focuses on the interpretation of the results obtained from such studies.

2.1. Design of the experiments

Block (2018) argues that Bayesian theories of visual perception do not necessitate positing probabilistic representations. While he focuses on the connection between visual phenomenology and probabilistic representations, he also critically reviews influential experiments that argue for probabilistic representations. The common premise of those experiments is that while our visual system represents information in terms of probability distributions, our discrete responses or perceptual decisions are samples[1] taken from these distributions. Block argues that while these results provide an account of perceptual decisions, they do not say anything about perception itself.

Experiments on perceptual processes generally involve tasks where observers are asked to make judgments about the perceptual feature under scrutiny. Observers are generally required to focus on a visual stimulus (e.g., an object oriented in a specific way, or having a certain colour), and then answer a question based on that stimulus (e.g., is the object tilted right or left; is there a red target among orange distractors?). Block argues that while decisions in response to such queries can be described probabilistically, in ordinary perception, an object's feature, such as its orientation or colour, would appear, for example, vertical or blue to an observer without any task or query. There might be competing representations, and the result of this competition can be described probabilistically (e.g., observers answered "tilted right" 60% of the time), but this does not mean that the competing representations are themselves probabilistic.

Moreover, Block (2018) argues that such perceptual decision tasks require observers to impose cognitive categories on perceptual information. For example, if you see an animal outside of my house on a foggy night, your visual system might assign a high probability that the visual object is a four-legged animal. However, if you try to judge whether it is a cat or a small dog, then a different probability would be assigned to the result of my perceptual inference. In other words, the probabilities potentially assigned by the visual system cannot be assigned until the category that the perceptual decision is being made within is specified. In Bayesian terms, a hypotheses space must be specified before a probability distribution of that space can be defined, and this space determines the shape of the probability distribution of possible hypotheses (i.e., dog, cat, human, etc.) given the scene. In the previous example, this conditional probability distribution will differ by whether the space of hypotheses is composed of four-legged animals or animate objects in general. For perceptual tasks, these categories are imposed onto the task in advance by the experimenter. In other words, the space of hypotheses is generated for the observers beforehand, but this does not occur in the real world. Block therefore claims that the results of perceptual studies depend on the cognitive categories that researchers impose on the perceptual task.

Perceptual tasks in such experiments are designed to be challenging to obtain more informative results. This allows researchers to examine patterns of mistakes that the visual system makes. For example, the visual items could be presented very briefly, they can be presented in a crowded scene (either spatially or temporally) or in the peripheral visual field. However, these tasks can be so demanding from the observers' perspective that they can become subjective guessing tasks rather than perceptual tasks (Block, 2018). For example, try to discriminate the letters used in *this* word, while fixating on *this* word. You will probably perceive letters in your peripheral visual field, but your decision on the word will likely feel like a subjective guess. Moreover, these tasks may require very complex cognitive decisions. For example, a perceptual task may involve several response choices with different monetary reward tied to each response under varying time restrictions. Such experiments can force observers to adopt a certain cognitive strategy when performing the task, which, in turn, may mean that the results say more about cognitive decisions than perception. Block claims that these issues create a gap between *studies* of perception and what normally happens in perception.

2.2. Interpretation of experimental results

Empirical evidence for probabilistic representations in perception comes from studies indicating that the visual system represents its own sensory uncertainty, caused either by stimulus properties and/or by the noise in the internal mechanism of the system. This idea can be seen in perceptual cue combination studies (Ernst & Banks, 2002; Körding & Wolpert, 2006) or studies focusing on Bayesian priors and their effect on perceptual decisions (e.g., Weiss, Simoncelli, & Adelson, 2002). However, Rahnev (2017) argues that no empirical evidence has so far shown that uncertainty is represented or used as a full probability distribution in perceptual decisions. He proposes alternative representational schemes that can account for the experimental findings without necessitating that information about the whole probability distribution is represented. For example, viewing a moving bar would create a population code in motion-sensitive neurons where each neural structure is tuned to a different direction (Fig. 1A). This could be characterized as a sensory distribution over possible orientations, where the height of the curve represents the level of neural activity. Rahnev argues that this information is not available for perceptual decisions. Instead, only the

[1] Rather than being completely random, these samples are mostly considered to be outputs of a cost minimization computation given the probability distribution and a cost function applied to it.

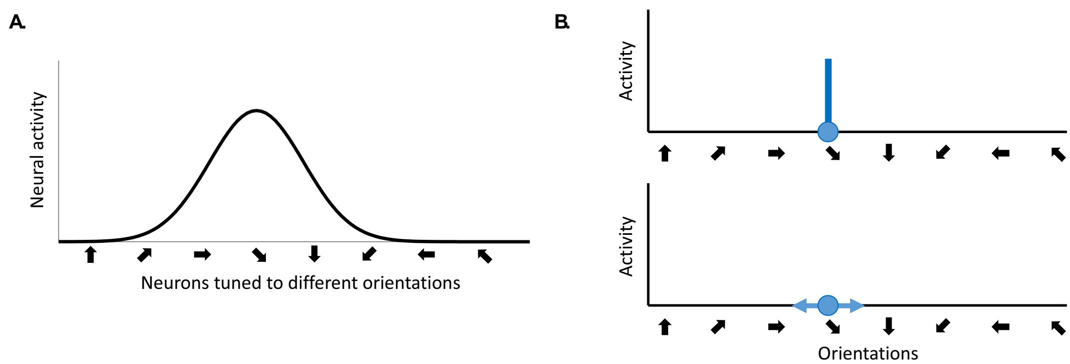

Fig. 1. *A.* Example response pattern of motion sensitive neurons in area MT evoked by an image of a bar moving diagonally in the down-right direction. *B.* Two possible ways of representing this information are shown. At the top: Only a point estimate and a strength-of-evidence value, which corresponds to the height of the neural activity distribution, are extracted from the full distribution. At the bottom: The mean and the variance of the neural activity distribution are extracted. In both cases, the perceptual representation includes a point estimate and a variable that correlates with the uncertainty in the sensory evidence (Rahnev, 2017).

summary statistics of this distribution (e.g., mean) can be accessed and used in perceptual decisions. Moreover, these summary statistical features can be represented in different ways. For example, only the hypothesis with the highest activity accompanied by the strength of this activity (Fig. 1B, top row), or only the mean and the variance of the distribution might be represented (Fig. 1B, bottom row). Experimental results assumed to support representations containing probability distributions can also be explained by these alternative summary statistical schemes.

Rahnev (2017) also proposes possible methodological approaches to empirically distinguish whether perceptual representations include full distributions or only their summaries. One critical approach highlighted by Rahnev (2017) is to use visual stimuli that would produce non-Gaussian sensory evidence. When the stimulus variability is Gaussian, distinguishing a full representation and a summary representation becomes almost impossible. Encoding the summary statistics of a Gaussian distribution (e.g., mean and variance) is equivalent to having information about the full Gaussian distribution. However, if multi-modal (e.g., bimodal) distributions of sensory evidence are used, a summary statistical scheme would only represent a single mode of that distribution, whereas full distribution representations would carry information about other peaks as well. Therefore, such experiments could provide more accurate insights about the format of perceptual representations.

Block (2018) also claims that the visual system does not represent uncertainty in the way that is often assumed. While structures in the visual system are *sensitive* to uncertainty, this does not mean that these structures *represent* this uncertainty. Therefore, experiments that putatively support probabilistic representations can only be used to claim that the visual system is sensitive to uncertainty. Such sensitivity to uncertainty can give rise to complex adaptive responses, where seemingly "representational" processes are embedded in the physiological architecture of the visual system. These processes can be complex for researchers to understand (and make them posit probabilistic representations to understand the system), but they are not complex from the perspective of the system itself.

While there are differences between the standpoints of Block and Rahnev (e.g., whether perception or perceptual decision is chosen as the locus of criticism), their views on the perceptual representation of uncertainty parallel one another. Rahnev argues that visual representations consist of a single estimate of the stimulus feature, coupled with an internal variable (e.g., strength of evidence or variance) that correlates with sensory uncertainty, instead of a probabilistic representation of that uncertainty. In other words, the sensory uncertainty is not fully represented but there are internal states that co-vary with it which make these visual processes sensitive to uncertainty. Block reaches a similar conclusion (see Section 5 for details), although Rahnev does not make an explicit distinction between sensitivity and representation. Overall, they both agree that the current empirical evidence indicates that there are internal visual states that co-vary with stimulus uncertainty, but this is not enough to claim that our visual system uses probabilistic representations that incorporate probability distributions over possible estimates of a visual feature.

3. The feature distribution learning (FDL) method

Recently, Chetverikov et al. (2016, 2017a); Chetverikov, Campana, and Kristjánsson (2017c) introduced a new method to assess representations of visual feature distributions. Their results indicate that observers encode not only the summary statistics of visual feature distributions, but also the distributions themselves.

In classical visual search tasks, observers see a display containing several items composed of distractors and a target. Sometimes observers are told what feature they should search for (e.g. search for '/' in Fig. 2B), while in others, observers are not informed about the features of the target in advance, but are told to search for the odd-one-out item (e.g. find the oddly oriented line in Fig. 2B, without knowing in advance that it is '/'). In the former, as soon as there is a match between the feature value of the target and an item in the display, the search ends. However, this is not the case for odd-one-out search, where observers have to visually process all the items. Whether an item is a target depends on whether all the other items can be grouped into a category of distractors, to which the target is unlikely to belong. Odd-one-out search forces observers to process distractor orientations so that the target orientation can be defined. However, this search can become more efficient when priming effects occur (Kristjánsson & Driver, 2008; Maljkovic & Nakayama, 1994). If observers perform a series of odd-one-out searches where the distribution of distractor features is repeated and observers encode the features of the target and the distractors, the search becomes increasingly similar to a search where the target and distractor features are known beforehand. As a result, search times decrease as observers perform repeated odd-one-out search with roughly constant distractor and target features. This is what happens during the learning trials of the FDL method (Fig. 2A).

Once observers encode the repeated target and distractor features, role-reversal effects are induced when target and distractors features are swapped, resulting in increased search times (e.g., the features of the target and distractors in Fig. 2A are swapped in the search display in Fig. 2B). Increased search times on the test trials reflect the attentional suppression applied to the distractor orientations repeated during the learning trials (see Geng, Won, & Carlisle, 2019, for a review of how the visual system ignores distractors). Due to the role reversal, the target orientation on the test trial has the feature that the visual system has learned to ignore. A crucial point, however, is that this attentional suppression is not applied in a constant or arbitrary manner, but

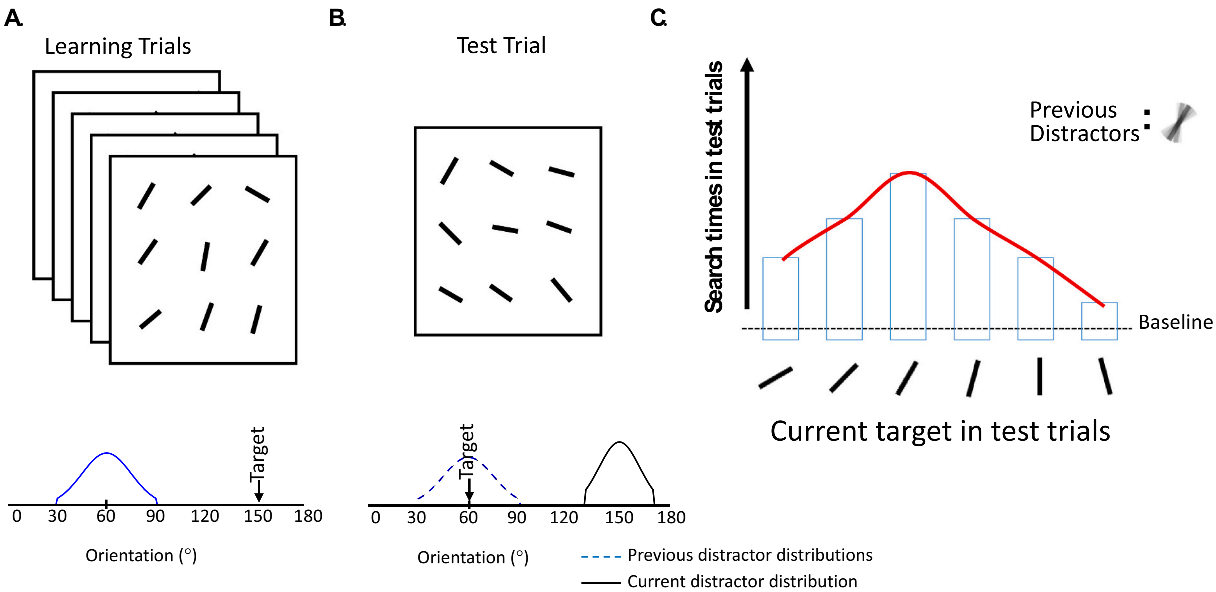

Fig. 2. The FDL method consists of blocks of trials, where each block includes a series of learning trials and then a test trial. *A*. On learning trials observers perform a series of odd-one-out searches where distractor orientations are sampled from the same distribution. In this example, it is a Gaussian distribution centered at 60°, shown with the blue curve at the bottom. *B*. On the test trial the features of the target and distractors are swapped and the target can now have the same features as the previous distractor distribution, shown with the dashed blue curve at the bottom. *C*. Hypothetical search times on test trials as a function of the target orientation from test trials. Search times following role-reversals depend on the similarity between the current target and the previous distractors as encoded by the visual system. Manipulating this similarity and observing its effects on search times reveals observers' representations of previous distractors (the red curve). (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

depends instead on the shape of the distractor distribution encoded during the learning trials. In Fig. 2, orientations around 60° will be suppressed the most since distractors on search trials preceding the role reversal were mostly oriented around 60°. However, this suppression will decrease for other orientations, in proportion to the probability of being a distractor as encoded by the visual system. The search times in Fig. 2C reflect this, and hence, allow assessment of how observers encoded the orientation distribution of the distractors from the learning trials.

To summarize the rationale behind FDL: Search times on test trials depend on the similarity between the target on the test trial and the distractors from the preceding learning trials as represented by the visual system. The structure of this representation is revealed by probing it at different locations in feature space, which is achieved by manipulating the similarity between the current target and the previous distractors. In FDL, this similarity is quantified by the parameter "CT-PD", which is defined as the distance (in feature space) between the feature value of the current target ("CT") and the mean feature value of the previous distractor distribution ("PD", see examples in Fig. 3, middle column). When search times from the test trials are plotted as a function of CT-PD distances (which will be referred as CT-PD curves from now on), this reveals observers' representations of the distractor distribution from the learning trials.

Chetverikov et al. (2016, 2017a); Chetverikov, Campana, and Kristjánsson (2017b, 2020) showed that observers' search times as a function of the CT-PD distance follow the shape of the distractor feature distribution on the preceding learning trials. This even occurred when Gaussian and uniform distributions with the same mean and range were contrasted (Fig. 3). This demonstrates that observers build a model of the distractor features in terms of a probability distribution. Similar results have been observed with asymmetric skewed distributions (Chetverikov et al., 2016), with bi-modal distributions (Fig. 4B; Chetverikov et al., 2017b, 2020), with hue distributions of isoluminant colored items (Chetverikov et al., 2017a; Hansmann-Roth, Chetverikov, & Kristjánsson, 2019), and when the search array appeared in the peripheral visual field (Tanrıkulu, Chetverikov, & Kristjánsson, 2020; for a review, see Chetverikov, Hansmann-Roth, Tanrıkulu, & Kristjánsson, 2019).

4. How does the FDL method differ from methodologies of other experiments supporting probabilistic representations?

The overarching problem with providing convincing empirical evidence for probabilistic representations is that observers' responses vary, but its source is difficult to identify. When faced with a perceptual decision on colour, an observer might respond "red" *70%* of the time. This observation does not directly demonstrate that the observer has a representation assigning *0.7* probability to "red". This response variation may also reflect artefacts from subjective guessing mentioned by Block (2018), from variability across experimental trials, or the process of sampling from the underlying distribution that the visual system performs when making perceptual decisions (e.g., Vul, Hanus, & Kanwisher, 2009; Vul & Pashler, 2008).

The crucial aspect of FDL is that the response feature differs from the visual feature being investigated. Therefore, FDL involves i) no query, ii) no perceptual decision about the distractor features that are learned, iii) no imposition of cognitive categories, iv) no sampling and v) no subjective guessing *about* the relevant visual feature. When the experimenter investigates how observers encode the orientation of a set of lines, the task only requires observers to respond to the location of the target in the search array. When colour encoding is investigated, the observer only judges the location of the cut-off on a diamond-shaped item. The task does not include any perceptual decisions about the visual search feature, which is merely processed for visual analysis of the stimulus. This visual process might include a step in which the visual system has to "decide" that a certain line (or a diamond) is the target before responding to its location (or its cut-off). However, this step neither requires a response nor any explicit analysis of the target features, because most of the time the target pops out among the distractors. Even when observers passively viewed the search displays during the learning trials (i.e., without making any response), response times on the test trials still demonstrated that observers were able to

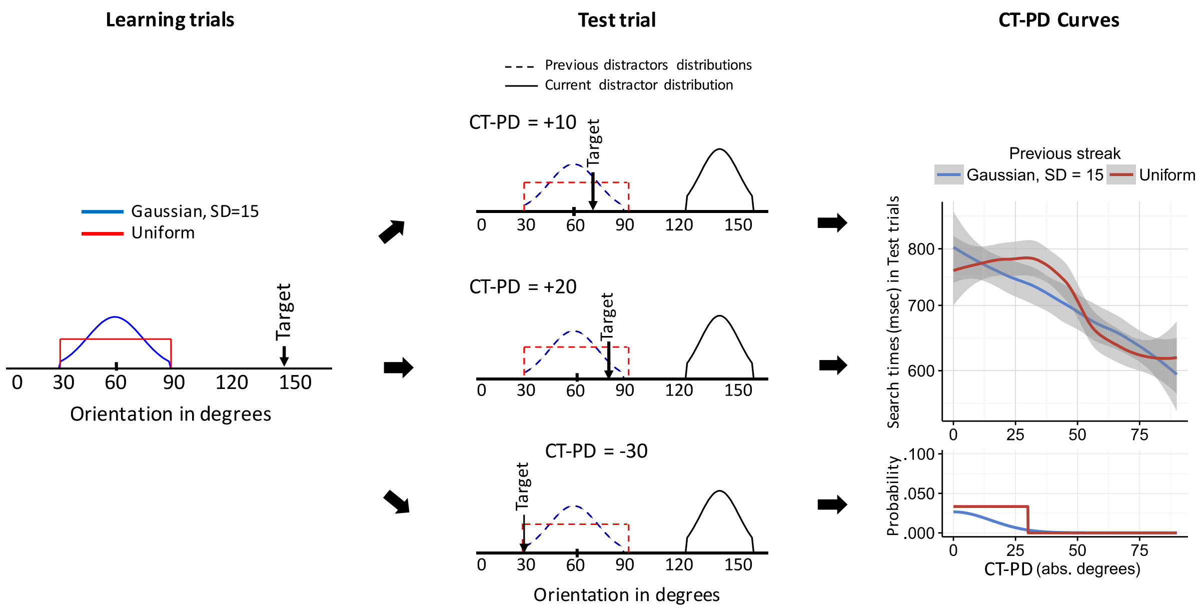

Fig. 3. An example of how Chetverikov et al. (2016) created CT-PD (Current Target – Previous Distractor distance) curves. On the learning trials, the distractor orientations were drawn from either Gaussian or uniform distributions (Left column). CT-PD distance was manipulated across blocks of trials throughout the experiment (Middle column). The search times from the test trials were plotted as a function of CT-PD distances which correspond to observers' representations of the distractor distributions. The physical distribution of the distractors is shown below the CT-PD curves.

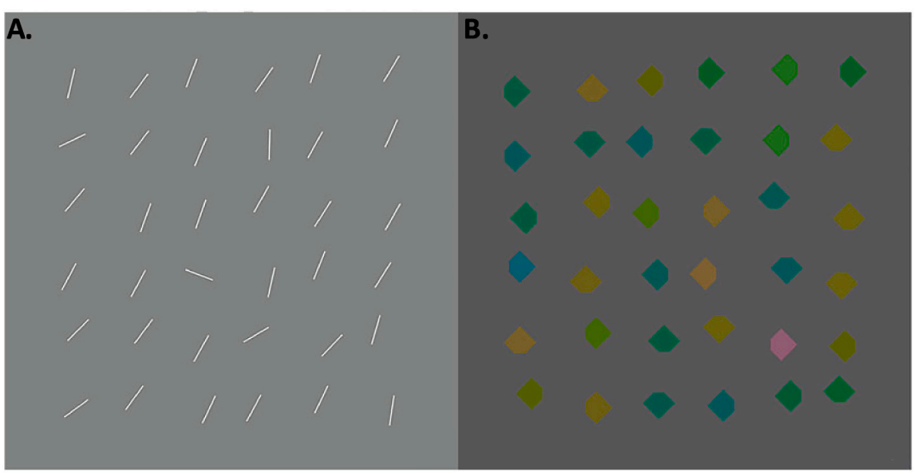

Fig. 4. *A.* An example search display for orientation. Observers search for the oddly oriented line and then indicate whether that line is located in the upper or lower half of the search display. In this example, the correct response would be "bottom". *B.* An example search display for colour. Observers search for the oddly colored diamond and indicate the location of the cut-off on that diamond. In this example, the correct response is "left".

encode information about the distractor orientation distribution (Kristjánsson, Ceylan, & Pascucci, 2021). This indicates that even passive exposure to the search array makes the visual system automatically distinguish the target from the distractors. The crucial point here is that the probability distributions revealed by the FDL method are not imposed onto the task by the experimenter via probabilistic description of observer's responses, because the method does not require a response concerning the visual feature whose distribution is being assessed. Instead, these probability distributions originate from the visual process necessary to perform the search. FDL therefore avoids the methodological criticisms of Block (2018) and Rahnev (2017) summarized in Section 2.1.

Moreover, learning effects observed with FDL cannot be explained by low-level neural mechanisms (e.g., adaptation) that occur before the visual input reaches the primary visual cortex (Maljkovic & Nakayama, 1994). These effects also cannot be attributed to post-perceptual or cognitive processes. FDL relies on role-reversal effects, which Kristjánsson and Driver (2008) describe as a form of "negative priming" (Tipper, 1985, 1992), where the internal representation of the to-be-ignored stimulus (i.e., distractors) is associated with suppression. This representation is considered to reflect a primitive memory system that operates at different levels of the visual processing hierarchy (Kristjánsson & Ásgeirsson, 2019; Kristjánsson & Nakayama, 2003; Nakayama, Maljkovic, & Kristjánsson, 2004). This system has strong parallels with the "perceptual representation system" (Magnussen & Greenlee, 1999; Tulving & Schacter, 1990) which is defined as a pre-semantic perceptual memory operating independently of other memory systems (for a review, see Kristjánsson, 2006; Kristjánsson & Campana, 2010). Such priming effects have been observed independently of observers' expectancies, perceptual learning, anticipatory strategies (Becker, 2008; Kristjánsson & Driver, 2008; Maljkovic & Nakayama, 1994, 1996; Shurygina, Kristjánsson, Tudge, & Chetverikov, 2019;

Sigurdardottir, Kristjánsson, & Driver, 2008), or priming of responses (Goolsby & Suzuki, 2001; Sigurdardottir et al., 2008). The strength of inter-trial priming in visual search depends on the level of uncertainty in the task (Meeter & Olivers, 2006; Olivers & Meeter, 2006), suggesting that priming helps the visual system deal with uncertainty. All this means that investigations of role-reversal effects allow us to examine how information is encoded by a representation system that is primarily *perceptual*.

In sum, the characteristics of FDL summarized in this section allow this method to meet the first criterion of providing evidence for probabilistic representation. They ensure that the probability distribution revealed by this method is not imposed upon the task by the experimenter. Instead, it is generated as a result of the search process for the visual system itself, rather than being a probabilistic way of describing the output of the search process from the perspective of the experimenter. The method achieves this by making sure that the representations studied are primarily perceptual and are not contaminated by any post-perceptual judgments or categories. The next section describes how the FDL method meets the second criterion for probabilistic representations, which is to provide evidence for a structural correspondence between the probabilistic structure of the world states and the inner states. The section also describes evidence showing that the brain actually exploits such structural correspondences.

5. How do FDL results differ from results of other experiments supporting probabilistic representations?

5.1. Summary of key results obtained with FDL

In Chetverikov et al. (2016), the distractor distribution used on learning trials was either Gaussian or uniform (Fig. 3). They found remarkable correspondence between the shape of the underlying representation of the distractor distribution (i.e., CT-PD curves) and the shape of the physical distractor distribution used as the stimulus. This correspondence is especially important given that the two different distractor distributions had the same mean and range, which are summary statistical variables crucial for outlier detection within sets of visual items (Hochstein, Pavlovskaya, Bonneh, & Soroker, 2018). Therefore, the difference observed in the CT-PD curves (Fig. 3) cannot be attributed to summary statistical representations. Notably, the results were replicated when Gaussian and uniform distribution were equated on variance rather than range or when the target feature is kept constant or varied randomly (within a certain range) during the learning trials (Chetverikov et al., 2016, Experiments 3A-3C). This demonstrates that encoding of the target features alone cannot account for the correspondence between the CT-PD curves and the physical distractor distribution.

These results cannot be explained by representational schemes that rely on summary statistics (e.g., representing only the mean and variance). Instead, they can only be accounted for by a representation that includes information about the full probability distribution of the distractors.[2] For example, the search times obtained on the test trials when the orientations of the previous distractors were distributed uniformly were fairly similar inside the range of the uniform distribution (Fig. 3, rightmost column, the part of the red line corresponding to orientations smaller than *30*°). If the visual system specifically represents the mean of this distribution rather than the distribution itself, the search times should have decreased as the CT-PD distance increased (which is what happens when the previous distractor distribution was Gaussian). However, the fact that the search times are constant in the range of the uniform distribution indicates that the distribution (rather than just its mean) is represented.

[2] At a minimum, only a representation including information beyond and above distribution summary statistics can account for these results.

In a later study, Chetverikov et al. (2017b) used a bimodal distractor distribution composed of two separate uniform segments on learning trials. Following two learning trials, CT-PD curves revealed a unimodal representation of the distribution with a peak between the two modes. However, following more learning trials, the shape of the CT-PD curves took the form of the bimodal distribution. This is strong evidence that observers initially assumed a unimodal distribution but updated their representation to a bimodal distribution as more sensory evidence accumulated. Moreover, Chetverikov et al. (2020) presented two targets simultaneously within a single search trial, where distractor orientations were drawn from a bimodal distribution (Fig. 5A). Search times for targets between the peaks of the bimodal distribution were lower than if the target appeared on one of the peaks. They also ran simulations with three different models (the three models in Fig. 5B) to distinguish between a probabilistic representation and two other possible summary statistical representation models. The probabilistic model predicted observers' search times (and the order in which the targets were reported) better than the other two, which indicates that observers' representations approximated the physical distribution. Even though the gap between the two peaks was partly filled, the bimodality of the distribution was still visible in the search times. As also suggested by Rahnev (2017), using bimodally distributed sensory evidence provides clearer evidence for distinguishing between full-distribution and summary statistical schemes.

5.2. What kind of empirical evidence is needed for probabilistic representations?

The strongest empirical evidence for probabilistic perceptual representations would involve demonstrating that the visual system actually uses the probabilistic information about sensory uncertainty to guide behavior. While some have attempted to show this (for a review see Ma & Jazayeri, 2014) it is still unclear whether this evidence for representing sensory uncertainty is linked to perceptual or post-perceptual processes (Gross, 2018), which is very difficult to distinguish empirically as mentioned before. Positive evidence for links to perceptual processes comes from an fMRI study by Van Bergen, Ma, Pratte, and Jehee (2015); also see Walker, Cotton, Ma, and Tolias (2019). Observers were shown an oriented grating and then asked to report the orientation by rotating a bar. Similar to previous studies, observers' orientation judgments were biased away from cardinal axes. This bias increased when there was high sensory uncertainty indicating that observers used that uncertainty in their perceptual decisions. However, the novel part of their study was that the uncertainty was not induced by any external sources (such as adding noise to the visual stimulus), but was decoded from the visual cortex with fMRI. Therefore, even when the physical features of the visual stimulus were kept constant, the degree of bias in observers' judgments was highly correlated with uncertainty in the visual cortex. However, Block (2018) argued that what may look like a *representation* of uncertainty in that study could instead reflect mere *sensitivity* to sensory uncertainty (or in Block's terms: sensitivity to the degree of competition between non-probabilistic representations). Gross (2020) refers to this criticism by Block as the "mere sensitivity challenge", which boils down to the question of whether sensitivity to sensory uncertainty qualifies as a probabilistic representation of that uncertainty.

Some proponents of probabilistic theories have also raised similar concerns. For example, Knill and Pouget (2004) acknowledge that empirical results in the literature can also be accounted for by a covariation relationship between internal structures in the visual system and uncertainty in the visual stimulus, rather than reflecting probabilistic representations of that uncertainty. Chater et al. (2006) agree that a set of heuristic tricks implemented by the visual system can account for empirical results suggesting computations of probabilistic representations. However, contrary to Block, both Knill and Pouget (2004) and Chater et al. (2006) predict that a system whose operation depends on

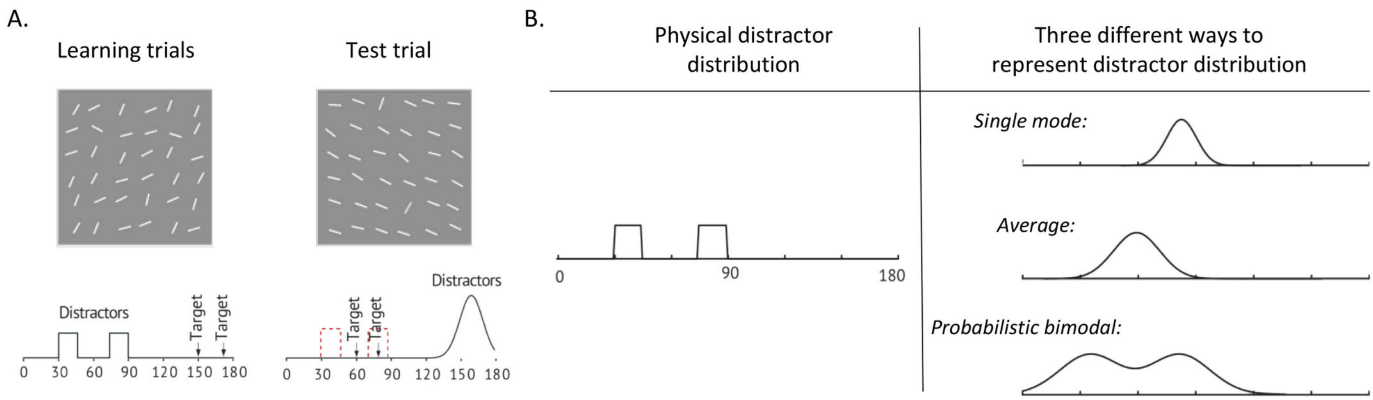

Fig. 5. *A*. Example learning and test trials used in Chetverikov et al. (2020). The plots at the bottom show the features of targets and distractors used in this example trial. The solid lines show the distractor distribution used on the current trial, while the dashed lines show the distractor distribution on the previous trial. *B*. The left column shows the physical distribution of distractor orientations. The right column shows possible different schemes for how this information might be represented. *Top row*: Only one of the modes of the bimodal distribution is chosen and represented. *Middle row:* The distribution can be represented with summary statistics (mean and variance). *Bottom row*: Both modes of the distribution are represented, which approximately equals the full distribution (see Chetverikov et al., 2020 for more details).

such heuristics cannot explain the high-level flexibility and generality of human perception and cognition.

Gross (2020) states that Block's conception of mere sensitivity is well captured with Shea's notion of representations that is based on the exploitation of correlational information carried by a range of states (Shea, 2018, p.78). An example of this notion would be the correlation between the honeybee nectar dance and the location of a nectar source, where the angle of the performed dance and the number of waggles in it provide information about the direction and the distance of the nectar source from the hive. Bees are not only *sensitive* to the correlation between the features of the dance and the range of different directions and distances for the food source, but they can exploit this correlation to guide their behavior and forage successfully. Similarly, the visual system's exploitation of the correlation between its inner states and sensory uncertainty caused by a distal stimulus (as shown by Van Bergen et al., 2015) involves more than mere sensitivity, so it qualifies as a representation of that uncertainty (see also Shea, 2020).

However, even though such a correlational notion of representation could overcome the mere sensitivity challenge, it would be in line with Rahnev's (2017) alternative summary representational schemes (Fig. 1B). This correlational representation of sensory uncertainty includes a single estimate of the distal visual feature coupled with a parameter indicating its level of uncertainty. However, as Rahnev (2017) rightly states, assumptions in current probabilistic accounts of perception do not involve such summary representational schemes, but instead they involve perceptual representations that include probability distributions of possible estimates of the visual feature. The latter, but not the former, reflects the type of representation schemes both Rahnev (2017) and Block (2018) criticize. Therefore, a correlational notion of representation cannot avoid Block's and Rahnev's criticism of probabilistic representations. This correlational notion of representation parallels the account based on co-variation relations that Knill and Pouget (2004) — or the heuristic account that Chater et al. (2006) — contrasted with truly probabilistic representations assumed in the empirical literature. Therefore, even if empirical results indicating an exploited correlation between sensory uncertainty and internal states of the visual system (such as those obtained by Van Bergen et al., 2015; or by Walker et al., 2019) can be considered representations of sensory uncertainty, they still cannot be considered empirical evidence for probabilistic representations of sensory uncertainty as posited in probabilistic Bayesian accounts of perception, because the criticisms of Block and Rahnev would still hold.

Rather than a notion of representation based on exploited correlational information, a representational scheme based on an exploited *structural correspondence* between internal states of the brain and the external states of the world would provide empirical evidence for probabilistic representations. The structural correspondence in these types of representations emerges when "a collection of representations in which a relation on representational vehicles represents a relation on the entities they represent." (Shea, 2018, p.118). A simple example of this structural notion of representation would be a topographical map of a landscape. The reason a hiker can use such a map as a representation of the actual landscape is the structural correspondence between the map and the landscape. Exploitation of this correspondence (i.e., considering the map as a representation) would be crucial for explaining the hiker's behavior (see also Ramsey, 2007, p.77–80).

Representations with structural correspondence better capture the notion of probabilistic representation posited in probabilistic Bayesian accounts of perception. In probabilistic representations, not only is the correlational information about the probability of an estimate of a distal visual feature being exploited, but the *relations* between different correlational representations are also exploited. To clarify this distinction, let's first consider again the correlational relation between the honeybee nectar dance and the location of the nectar source with respect to the location of the beehive. This relation does indeed include a structural relation since the difference between the number of waggles in two different dances corresponds to the distance difference between the two nectar sources. However, as Shea (2014, 2018) states, this structural relation is not exploited because the relation between the dances is not used in guiding the behavior of the bees. For example, the fact that the 6-waggle dance is three waggles more than the 3-waggle dance is not used by the bee for nectar foraging. Therefore, the honeybee nectar dance case does not qualify as an *exploited* structural representation. Contrary to this, probabilistic representations in terms of probability distributions include information about the relations between represented probabilities of possible estimates of a visual feature. For example, a probability distribution over different orientations would provide information about the relation between the probabilities assigned to different orientation estimates.

An account of probabilistic representations requiring structural correspondence is also consistent with Rahnev's (2017) view that the alternative representational schemes that he proposed (Fig. 1B) are not genuine probabilistic representations, but instead they provide correlational information about a single estimate. These alternative representational schemes can of course carry information about a range of different orientations, or there can be multiple representations each of which is correlated with a different orientation. However, these representational schemes do not include information about the probabilistic relations between different orientations, as probability distributions do. Hence, they do not qualify as an exploited structural representation. However, the probabilistic representations that consist of probability distributions over feature values, should be considered structural representations because the relations between probabilities for different feature values are assumed to be computed and used in guiding the behavior. This is the type of probabilistic representation posited in the perception literature, and both Rahnev (2017) and Block (2018) rightly claim that there is no empirical evidence for such probabilistic representations.

Given this, Rahnev's (2017) and Block's (2018) perspectives can be recast as conveying that there is no empirical evidence for probabilistic representation as assumed in the perceptual science literature, because the current empirical evidence is more in line with a notion of representation that carries correlational information, as opposed to relying on structural correspondence. This is the crucial difference between the experimental results Block and Rahnev discuss and the results obtained with the FDL method, which demonstrate an example of a structural correspondence between the inner states of the visual system (inferred from CT-PD curves) and the underlying probability distribution of the stimulus in the physical world. The difference in the amount of suppression applied to any two distractor orientations (or colors) corresponds to the difference between the probabilities that they are distractors.

5.3. How does FDL provide such empirical evidence?

To qualify as a structural representation, the structural relation between the probabilities tied to different orientations should also be exploited by downstream processes of the visual system. Shea (2018, 2020) proposes that if the degree of match (or mis-match) between the represented and the actual probabilistic structure explains the success (or failure) of behavior, this would indicate that the probabilities are included in the representational content, as opposed to being a part of the mode or the manner that non-probabilistic information is represented in, as discussed by Gross (2020). The results from Chetverikov et al. (2020) provide an example of how the probabilistic structure demonstrated by the FDL method is exploited by the visual system, and also show how the degree of mismatch between the CT-PD curves and the physical distribution of orientations in the stimulus explains observers' behavior.

Chetverikov et al. (2020) included two targets on their odd-one-out search trials (Fig. 5A). Observers were instructed to respond as soon as they found a target. Apart from observers' search times, they also looked at which of the two targets was more likely to be found first on test trials as a function of their orientation with respect to the distractor distribution on the preceding learning trials (Fig. 6). Targets whose orientation fell between the two modes of the previous distractor distribution were more likely to be found first when the other target was on one of the modes of that distribution. Such a result was expected, since the orientation of "Target B" depicted in Fig. 6 (middle row) would be suppressed more than the orientation of "Target A", given that the orientation of "Target B" is selected from inside one of the modes of the previous distractor distribution. An interesting condition appears when this comparison is done between a target orientation outside of the range of the bimodal distribution (Target C in Fig. 6) and a target in between the two modes of the previous distractor distribution (Target A in Fig. 6). The suppression of both these target orientations should be equal, because their probability of being a distractor is zero given the physical distractor distribution on the learning trials. However, Chetverikov et al. (2020) found that in this condition the target outside of the range of the previous bimodal distractor distribution (Target C) was more likely to be found first than the one in between the modes of that bimodal distribution (Target A). This observation can only be explained by the mismatch between how the observers represented the bimodal distractor distribution and the physical bimodal distractor distribution used in the experiment. Search times from the test trials revealed that the visual system approximates the distractor distribution with a probabilistic bimodal representation similar to the one shown in Fig. 5B (bottom of rightmost column). In other words, the gap between the two modes of the distribution was partly filled in by the representation of that bimodality, which explains why "Target C" was more likely to be found first than "Target A". The probability of being a distractor was higher for "Target C" than "Target A" given the *representation* of the

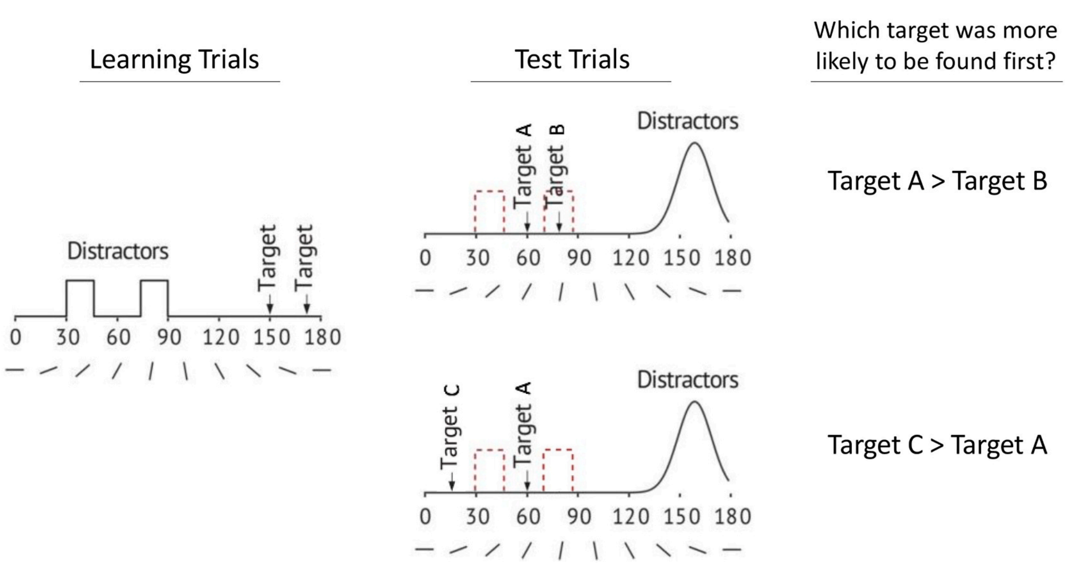

Fig. 6. An example of the target and distractor orientations used on the learning and test trials of Chetverikov et al. (2020). The red dashed lines in the "Test Trials" column show the distractor distribution used on the learning trials. The rightmost column shows which of the two targets was more likely to be picked by the observer given its orientation in relation to the distractor distribution on the learning trials. The exact locations of the targets in orientation space shown for the test trials are only shown as demonstrations. There are three types of targets categorized with respect to their location in orientation space. Target A is for target orientations that fall in between the two modes of the preceding distractor distribution. Target B corresponds to target orientations that fall inside one of the two modes of the preceding distractor distribution. Target C corresponds to target orientations that fall outside of the range of the preceding bimodal distractor distribution. See text for more details. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

distractor distribution, but not according to the actual physical distractor distribution. In other words, the order of the two targets found by the observers followed the relational structure of this internal representation, which clearly stands as a proxy for the physical distribution.

To summarize, we argue that existing evidence for probabilistic representations demonstrates that the visual system can exploit correlations between its inner states and the sensory uncertainty caused by the distal stimulus. However, such results only reveal sensitivity to sensory uncertainty, which is not enough to posit probabilistic representations as assumed in the perception literature (Block, 2018). Similarly, Rahnev's (2017) view is in line with the idea that representations based on exploited correlational information can account for the current empirical evidence presented in favour of probabilistic representations. Given these criticisms, we argue that empirical evidence demonstrating an exploited structural correspondence between the inner perceptual states and external stimulus conditions can be used as evidence for probabilistic representations. Such a notion of probabilistic representation fits better with the one assumed in current probabilistic accounts of perception. Importantly, FDL results not only provide evidence that is based on probabilistic structural correspondence, but also demonstrate how the degree of mismatch between the represented and the actual probabilistic structure of the feature distributions can explain observers' behavior.

5.4. Representing groups versus individual items

Both humans and animals seem to have an ability to perceive the underlying probabilities of statistical regularities in their environment (Gallistel, Krishan, Liu, Miller, & Latham, 2014). However, the main argument here concerns probabilistic representations, not representations of probabilities.

The CT-PD curves from FDL experiments reveal the representations of orientations (or colors) of a group of items. The represented entity in the FDL task could be described as an ensemble of visual items. In fact, FDL studies often refer to ensemble perception, which is the ability to represent multiple items as an ensemble instead of representing each individual item separately (for a review see Alvarez, 2011; Whitney & Yamanashi Leib, 2018). Therefore, it might look like there is a fundamental difference between processes used in FDL tasks, and simply representing the orientation or colour of a single item. However, this difference does not change or undermine the main conclusions inferred from FDL studies. Firstly, the fact that the search display in FDL studies includes multiple items is a trivial property of the task. For example, the space between the diamonds used in the search array of (Chetverikov et al., 2017a; see Fig. 4B) can be eliminated by positioning each colored item adjacent to each other, resulting in a single object that varies in colour. The same FDL task can be performed on this object by asking observers to locate whether the odd-colored patch is in the upper or lower half of this surface. Individuation of visual objects depends heavily on how the visual scene is interpreted, rather than being a direct reflection of how the external world is physically structured (Feldman, 2003). This interpretation depends mainly on how scene features are perceptually grouped by the visual system. Secondly, FDL tasks do not require observers to encode the underlying distribution of the distractor orientations. Not only that, even at a computational level, encoding the full distribution of the distractor orientation is not required for performing the search task. For example, encoding only a few summary statistical properties of the distractors (e.g. their range) would, in principle, be enough to detect the outlier in the search task. Therefore, if there are traces of ensemble processing or perceptual grouping in FDL studies, this should be attributed to how vision works, rather than to specifics of FDL.

The FDL method also includes temporal integration of visual information across learning trials (Chetverikov et al., 2017b). Therefore, results obtained by this method do not directly show that orientation of a single item on an individual trial is encoded probabilistically. Yeon and Rahnev (2020) argued that even though perceptual priors constructed over time could be considered as probabilistic, this does not indicate that the sensory representation of a single stimulus on a single trial is encoded probabilistically. The account presented here does not necessitate probabilistic representation of a single stimulus on a single trial, but it claims that probabilistic representations are best defined as structural representations in which relations between different correlational representations are exploited. These individual building blocks (e.g., correlational sensory representations) of perceptual representations are not what makes them structural (probabilistic) representations, but the fact that the relations between such individual building blocks are being exploited makes them structural and probabilistic.

Our visual system evolved to operate in complex environments. Visual items are never present in isolation in the real world. But instead they are always embedded within a spatial and temporal context. Representation of a visual feature of a single item on a single trial would still incorporate relational information with respect to its spatial (Utochkin & Brady, 2020) and temporal (Bae & Luck, 2017; Fischer & Whitney, 2014) surroundings. Therefore, we argue that an account of probabilistic representation based on structural correspondence that incorporates relational information would better serve empirical theories of perception.

6. Conclusions

Probabilistic approaches to perception and cognition have had great success, especially in building computational models of perceptual processes. This has led researchers to propose that the brain represents information probabilistically. Highly influential studies strongly suggest that probabilistic representations are used in visual perception. However, the methodology of such studies prevents them from providing clear evidence in favour of probabilistic representations (Block, 2018). Moreover, the results of these studies are more in line with summary statistical representations than representations in terms of probability distributions (Rahnev, 2017; Yeon & Rahnev, 2020). We have presented an experimental methodology that provides evidence for probabilistic representations in perception using a method that acknowledges these criticisms but crucially, also avoids them. Firstly, FDL is an ideal tool for examining perceptual representations because it does not require explicit judgments of the relevant visual feature being investigated. Secondly, and more importantly, the results obtained with FDL demonstrate an exploited structural correspondence between the internal states of the visual system and the probabilistic structure of the distal visual stimulus.

We argued that empirical evidence for probabilistic representations should demonstrate an exploited structural correspondence between the inner states of the brain and the external stimulus conditions. The notion of probabilistic representation posited by probabilistic Bayesian accounts calls for representations that depend on a structural correspondence. Moreover, representations based on structural correspondence could serve better for building theories with explanatory value in cognitive psychology (Gallistel, 1990; Gallistel, 2020; Ramsey, 2007).

Even though the methods and results in FDL paradigm potentially provide a prime example of what kind of empirical evidence is needed for probabilistic representations posited in probabilistic Bayesian accounts of perception, there are nevertheless open questions. For example, probabilistic representations are defined at Marr's (1982) computational level, and might not be needed at algorithmic or representational levels of explanation (Block, 2018). A common view within Bayesian approaches is that the processes at the algorithmic level approximate the Bayesian processes described at the computational level (Griffiths, Vul, & Sanborn, 2012). Block (2018) rephrases this statement by claiming that the visual system behaves *as if* it implements Bayesian inference. However, there is no consensus on whether there is a fundamental difference between these two (approximation vs. as if) claims (Rescorla, 2015). On the other hand, Sanborn and Chater (2016) argue that the brain does not represent probabilities, but instead functions as a Bayesian sampler (see also Griffiths et al., 2012). However, Icard (2016) argues that the sampling propensities are probabilities and are meaningfully represented by the brain. Investigating how probabilistic representations are defined at different levels of description can significantly contribute to the discussions raised in this paper.

The visual processes that are measured with FDL potentially build a better connection between studies in perceptual sciences and what actually happens during perception in real life. Perceptual grouping and ensemble perception are fundamental parts of visual processing. Attempting to isolate visual perception from such processes will prevent researchers from studying – in Block's (2018) terms – "genuine perception", rather than enabling such investigations. We also acknowledge the limitations of FDL. Firstly, FDL results have so far only provided evidence for probabilistic representations in low-level visual features on continuous scales (e.g. orientation, colour). How these findings will apply to higher-level visual representations or visual tasks that require representations of discrete quantities is still unknown. We believe, however, that much of our perceptual experience depends on the representation of these continuous feature spaces, such as colour, illumination, orientation, distance, etc. Secondly, the actual mechanism behind the role-reversal effects that FDL methods rely on is still not fully known. However, as we have argued, many studies on such priming effects strongly indicate that these effects emerge from manipulation of early perceptual representations. Further research is needed to assess the scope of the conclusions that can be drawn from FDL results. However, rather than being limitations, we consider all these questions as useful for making the problems about probabilistic perception more explicit and apparent. This would encourage constructive discussions for guiding future research to get to the crux of issues regarding probabilistic perception.

The notion of probabilistic representations has been a central theoretical concept within probabilistic accounts of perception. Regardless of its importance, such a notion still requires a bridge between its theoretical merits and empirical underpinnings. This is why we believe that the criticisms summarized here against the empirical evidence for probabilistic representations have to be acknowledged. But we have also presented an experimental method that sidesteps such criticisms, and provides an intriguing example of what kind of empirical evidence is needed for probabilistic representations in perception.

Acknowledgments

We thank for the valuable feedback from the attendees of the "Representation in Cognitive Science" workshop organized by Tobias Schlicht and Krzysztof Dolega at Ruhr-Universitat Bochum, where an earlier version of this paper was presented. ODT, SHR and AK are supported by grant IRF #173947-052 from the Icelandic Research Fund. AC is supported by Radboud Excellence Fellowship.

References

Alvarez, G. A. (2011). Representing multiple objects as an ensemble enhances visual cognition. *Trends in Cognitive Sciences, 15*(3), 122–131.

Bae, G. Y., & Luck, S. J. (2017). Interactions between visual working memory representations. *Attention, Perception, & Psychophysics, 79*(8), 2376–2395. https://doi.org/10.3758/s13414-017-1404-8.

Becker, S. I. (2008). The stage of priming: Are intertrial repetition effects attentional or decisional? *Vision Research, 48*, 664–684.

Block, N. (2018). If perception is probabilistic, why doesn't it seem probabilistic? *Philosophical Transactions of the Royal Society B, 373*, 1755.

Chater, N., Tenenbaum, J. B., & Yuille, A. (2006). Probabilistic models of cognition: Conceptual foundations. *Trends in Cognitive Sciences, 10*(7), 287–291.

Chetverikov, A., Campana, G., & Kristjánsson, Á. (2016). Building ensemble representations: How the shape of preceding distractor distributions affects visual search. *Cognition, 153*, 196–210.

Chetverikov, A., Campana, G., & Kristjánsson, Á. (2017a). Representing color ensembles. *Psychological Science, 28*(10), 1–8.

Chetverikov, A., Campana, G., & Kristjánsson, Á. (2017b). Rapid learning of visual ensembles. *Journal of Vision, 17*(2), 21, 1–15.

Chetverikov, A., Campana, G., & Kristjánsson, Á. (2017c). Learning features in a complex and changing environment: A distribution-based framework for visual attention and vision in general. In *, 236. Progress in brain research* (pp. 97–120). Elsevier.

Chetverikov, A., Campana, G., & Kristjánsson, Á. (2020). Probabilistic rejection templates in visual working memory. *Cognition, 196*, 104075.

Chetverikov, A., Hansmann-Roth, S., Tanrıkulu, Ö. D., & Kristjánsson, Á. (2019). Feature distribution learning (FDL): A new method for studying visual ensembles perception with priming of attention shifts. In *Neuromethods*. Humana Press.

Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature, 415*, 429–433.

Feldman, J. (2003). What is a visual object? *Trends in Cognitive Sciences, 7*(6), 252–256.

Fischer, J., & Whitney, D. (2014). Serial dependence in visual perception. *Nature Neuroscience, 17*(5), 738–743. https://doi.org/10.1038/nn.3689.

Gallistel, C. R. (1990). *The Organization of Learning*. London/Cambridge MA: MIT Press.

Gallistel, C. R. (2020). Where meanings arise and how: Building on Shannon's foundations. *Mind & Language, 35*, 390–401. https://doi.org/10.1111/mila.12289.

Gallistel, C. R., Krishan, M., Liu, Y., Miller, R., & Latham, P. E. (2014). The perception of probability. *Psychological Review, 121*(1), 96.

Geng, J. J., Won, B.-Y., & Carlisle, N. B. (2019). Distractor ignoring: Strategies, learning, and passive filtering. *Current Directions in Psychological Science*. https://doi.org/10.1177/0963721419867099.

Goolsby, B. A., & Suzuki, S. (2001). Understanding priming of color-singleton search: Roles of attention at encoding and "retrieval.". *Perception & Psychophysics, 63*, 929–944.

Griffiths, T. L., Chater, N., Kemp, C., Perfors, A., & Tenenbaum, J. B. (2010). Probabilistic models of cognition: Exploring representations and inductive biases. *Trends in Cognitive Sciences, 14*(8), 357–364. https://doi.org/10.1016/j.tics.2010.05.004.

Griffiths, T. L., Vul, E., & Sanborn, A. N. (2012). Bridging levels of analysis for probabilistic models of cognition. *Current Directions in Psychological Science, 21*(4), 263–268. https://doi.org/10.1177/0963721412447619.

Gross, S. (2018). Perceptual consciousness and cognitive access from the perspective of capacity-unlimited working memory. *Philosophical Transactions of the Royal Society, B: Biological Sciences*, 373. https://doi.org/10.1098/rstb.2017.0343.

Gross, S. (2020). Probabilistic representations in perception: Are there any, and what would they be? *Mind & Language, 35*, 377–389. https://doi.org/10.1111/mila.12280.

Hansmann-Roth, S., Chetverikov, A., & Kristjánsson, Á. (2019). Representing color and orientation ensembles: Can observers learn multiple feature distributions? *Journal of Vision, 19*(9), 1–17.

Hochstein, S., Pavlovskaya, M., Bonneh, Y. S., & Soroker, N. (2018). Comparing set summary statistics and outlier pop out in vision. *Journal of Vision, 18*(13), 12.

Icard, T. F. (2016). Subjective probability as sampling propensity. *The Review of Philosophy and Psychology, 7*(4), 863–903.

Kersten, D., & Yuille, A. (2003). Bayesian models of object perception. *Current Opinion in Neurobiology, 13*(2), 150–158.

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. https://doi.org/10.1016/j.tins.2004.10.007.

Körding, K. P., & Wolpert, D. M. (2006). Bayesian decision theory in sensorimotor control. *Trends in Cognitive Sciences, 10*(7), 319–326.

Kristjánsson, Á. (2006). Rapid learning in attention shifts - a review. *Visual Cognition, 13*, 324–362.

Kristjánsson, Á., & Ásgeirsson, Á. G. (2019). Attentional priming: Recent insights and current controversies. *Current Opinion in Psychology, 29*, 71–75.

Kristjánsson, Á., & Campana, G. (2010). Where perception meets memory: A review of repetition priming in visual search tasks. *Attention, Perception, & Psychophysics, 72*(1), 5–18.

Kristjánsson, Á., Ceylan, G., & Pascucci, D. (2021). Feature distribution learning by passive exposure. In *Poster Presented at the Annual Meeting of the Vision Science Society (V-VSS 2021)*.

Kristjánsson, Á., & Driver, J. (2008). Priming in visual search: Separating the effects of target repetition, distractor repetition and role-reversal. *Vision Research, 48*, 1217–1232.

Kristjánsson, Á., & Nakayama, K. (2003). A primitive memory system for the deployment of transient attention. *Perception & Psychophysics, 65*, 711–724.

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

Magnussen, S., & Greenlee, M. W. (1999). The psychophysics of perceptual memory. *Psychological Research, 62*, 81–92.

Maljkovic, V., & Nakayama, K. (1994). Priming of pop-out: I. Role of features. *Memory & Cognition, 22*, 657–672.

Maljkovic, V., & Nakayama, K. (1996). Priming of pop-out: II. The role of position. *Perception & Psychophysics, 58*, 977–991.

Mamassian, P., Landy, M. S., & Maloney, L. T. (2002). Bayesian modelling of visual perception. In R. RPN, B. A. Olshausen, & M. S. Lewicki (Eds.), *Probabilistic Models of the Brain: Perception and Neural Function* (pp. 13–36). Cambridge, MA: MIT Press.

Marr, D. (1982). *Vision*. San Francisco, CA: W.H. Freeman.

Meeter, M., & Olivers, C. N. L. (2006). Intertrial priming stemming from ambiguity: A new account of priming in visual search. *Visual Cognition, 13*, 1–28.

Nakayama, K., Maljkovic, V., & Kristjánsson, A. (2004). Short-term memory for the rapid deployment of visual attention. In M. Gazzaniga (Ed.), *Cognitive neurosciences III* (pp. 397–408). Cambridge, MA: MIT Press.

Olivers, C. N. L., & Meeter, M. (2006). On the dissociation between compound and present/absent tasks in visual search: Intertrial priming is ambiguity driven. *Visual Cognition, 13*, 1–28.

Rahnev, D. (2017). The case against full probability distributions in perceptual decision making. *bioRxiv*, 108944.
Rahnev, D, Block, N, Denison, R, & Jehee, J (2021). Is perception probabilistic? Clarifying the definitions. *PsyArXiv*. https://doi.org/10.31234/osf.io/f8v5r.
Ramsey, W. M. (2007). *Representation reconsidered*. Cambridge, MA: Cambridge University Press.
Rescorla, M. (2015). Bayesian perceptual psychology. In M. Matthen (Ed.), *The Oxford handbook of the philosophy of perception*. Oxford: Oxford University Press.
Sanborn, A. N., & Chater, N. (2016). Bayesian brains without probabilities. *Trends in Cognitive Sciences, 20*(12), 883–893. https://doi.org/10.1016/j.tics.2016.10.003.
Shea, N. (2014). VI—Exploitable isomorphism and structural representation. *Proceedings of the Aristotelian Society, 114*(2), 123–144.
Shea, N. (2018). *Representation in cognitive science*. Oxford: Oxford University Press.
Shea, N. (2020). Representation in cognitive science – Replies. *Mind & Language, 35*, 402–412. https://doi.org/10.1111/mila.12285.
Shea, N., & Frith, C. D. (2019). The global workspace needs metacognition. *Trends in Cognitive Sciences, 23*(7), 560–571.
Shurygina, O., Kristjánsson, Á., Tudge, L., & Chetverikov, A. (2019). Expectations and perceptual priming in a visual search task: Evidence from eye movements and behavior. *Journal of Experimental Psychology: Human Perception and Performance, 45*(4), 489.
Sigurdardottir, H. M., Kristjánsson, Á., & Driver, J. (2008). Repetition streaks increase perceptual sensitivity in brief visual search displays. *Visual Cognition, 16*, 643–658.
Tanrıkulu, Ö. D., Chetverikov, A., & Kristjánsson, A. (2020). Encoding perceptual ensembles during visual search in peripheral vision. *J. Vis., 20*(8). https://doi.org/10.1167/jov.20.8.20.
Tipper, S. P. (1985). The negative priming effect: Inhibitory priming by ignored objects. *Quarterly Journal of Experimental Psychology. A, Human Experimental Psychology, 37A*, 571–590.
Tipper, S. P. (1992). Selection for action: The role of inhibitory mechanisms. *Current Directions in Psychological Science, 1*, 105–109.
Tulving, E., & Schacter, D. L. (1990). Priming and human memory systems. *Science, 247*, 301–306.
Utochkin, I., & Brady, T. (2020). Individual representations in visual working memory inherit ensemble properties. *Journal of Experimental Psychology: General, 46*(5), 458–473.
Van Bergen, R. S., Ma, W. J., Pratte, M. S., & Jehee, J. F. (2015). Sensory uncertainty decoded from visual cortex predicts behavior. *Nature Neuroscience, 18*(12), 1728.
Vul, E., Hanus, D., & Kanwisher, N. (2009). Attention as inference: Selection is probabilistic; responses are all-or-none samples. *Journal of Experimental Psychology: General, 138*(4), 546–560. https://doi.org/10.1037/a0017352.
Vul, E., & Pashler, H. (2008). Measuring the crowd within: Probabilistic representations within individuals. *Psychological Science, 19*(7), 645–647.
Walker, E. Y., Cotton, R. J., Ma, W. J., & Tolias, A. S. (2019). A neural basis of probabilistic computation in visual cortex. *Nature Neuroscience, 23*(1), 122–129.
Weiss, Y., Simoncelli, E. P., & Adelson, E. H. (2002). Motion illusions as optimal percepts. *Nature Neuroscience, 5*(6), 598–604.
Whitney, D., & Yamanashi Leib, A. (2018). Ensemble perception. *Annual Review of Psychology, 69*, 105–129.
Yeon, J., & Rahnev, D. (2020). The suboptimality of perceptual decisions making with multiple alternatives. *Nature Communications, 11*, 3857. https://doi.org/10.1038/s41467-020-17661-z.