© 2020, American Psychological Association. This paper is not the copy of record and may not exactly replicate the final, authoritative version of the article. Please do not copy or cite without authors' permission. The final article will be available, upon publication, via its DOI: 10.1037/xhp0000749

I haven't a clue! - Expectations based on repetitions and hints facilitate perceptual experience of ambiguous images

Uri Hertz *1, Colin Blakemore 2,3, Chris D Frith 2,4

*Corresponding Author uhertz@is.haifa.ac.il

1. Department of Cognitive Sciences, University of Haifa, 3498838, Haifa, Israel

2. Institute of Philosophy, School of Advanced Study, University of London, London WC1E 7HU, UK

3. Hong Kong Institute for Advanced Study, City University of Hong Kong, 83 Tat Chee Avenue, Kowloon Tong, Hong Kong

4. Wellcome Trust Centre for Neuroimaging, University College London, 12 Queen Square, London WC1N 3BG, UK

Public Importance (117 words)

Visual perception relies on bottom-up processes in which visual information is accumulated, and modified by top-down expectations, but it is not clear how different sources of expectations interact to affect perception.

We used multiple presentations of ambiguous two-tone images and verbal hints to study their different effects on perceptual experience.

We found that object detection and vividness ratings significantly increased when images were introduced with a verbal hint, and that vividness increase was gradual over repetitions, indicating a role for attention and recollection on object perception.

Our findings confirm that top-down verbal hints can benefit object perception, but also how this may be attenuated by the specific visual information available, i.e. bottom-up processes, and by memory capacity.

Abstract (231 words)

In recent years the role of top-down expectations on perception has been extensively researched within the framework of predictive coding. However, less attention has been given to the different sources of expectations, how they differ and how they interact. Here we examined the effects of informative hints on perceptual experience, and how these interact with repetition-based expectations to create a long-lasting effect. Over seven experiments, we used verbal hints and multiple presentations of ambiguous two-tone images. We found that vividness ratings increased from one presentation to the next, even after the object in the image had been identified. In addition, vividness ratings significantly increased when images were introduced with a hint, and this boost was greater for more detailed hints. However, the initial increase in vividness did not always carry over to the next presentation. When recognition of the image in the presentation was hard, due to memory load, inconsistent presentation, or noise level of the image, the initial advantage in vividness was attenuated. This was most apparent when participants were primed with greyscale version of the two-tone image. A computational model based on evidence accumulation was able to recover these patterns of perceptual experience, suggesting that the effect of hints is short lived if it cannot be encoded in memory for future presentations. This notion highlights the different contributions of attention, memory and their interactions on forming expectations for perception.

Keywords: Perception, Semantic Information, Top-Down Effects, Vision, Attention, Memory

Word Count: 9906

Introduction

Visual perception is considered an inferential process in which percepts are inferred from a unique combination of low-level features and high-level expectations. The classic theory of visual perception as a hierarchical bottom-up process (Van Essen & Maunsell, 1983) has been enriched in recent years by findings highlighting the effects of expectations, predictions and prior knowledge on visual perception (De Lange, Heilbron, & Kok, 2018; Gilbert & Li, 2013; Pearson & Brascamp, 2008; Summerfield & De Lange, 2014). At least two sources of top-down influences, which can shape, affect and tune visual perception, have been identified. The first is the effect of informative hints, in which expectations of visual inputs are set by using cues or primes which immediately promote object detection (Eimer, 2014; Summerfield & Egnér, 2009). Expectations can also be formed by previous experience with the visual stimuli at hand, which increases the vividness of the perceptual experience over repeated presentations, through recollection of the stimuli (Greve, Cooper, Tibon, & Henson, 2019; Pearson & Brascamp, 2008; Marlieke T.R. van Kesteren, Ruiter, Fernández, & Henson, 2012). As these two sources of top-down effects are usually tested separately, their combined contribution to perceptual experience is unknown. Here we set out to examine the top-down contribution of previous experience and hints on object detection and perceptual vividness of ambiguous visual stimuli.

It has been suggested that attention can create visual expectations by giving the observer a hint, or cue, just before performing the visual task (Eimer, 2014; Summerfield & Egnér, 2009). In this account object detection can be facilitated by directing attention towards the object. In a crowded scene with multiple objects, such as in visual search paradigms, providing a visual or semantic-verbal cue about the type of object that needs to be detected activates an attention template, which allows accelerated detection of the object (Nako, Smith, & Eimer, 2015; Sun, Shen, Shaw, Cant, & Ferber, 2015). Expectations about the type and likelihood of a stimulus can also facilitate object detection (Summerfield & De Lange, 2014). For example, people are more likely to interpret an image that contains a mixture of

face and house as a face when prompted to look for faces, but more likely to interpret the same image as house when prompted to detect houses (Summerfield et al., 2006). In addition, contextual visual features can provide expectations that enhance object detection. For example, detecting the orientation of Gabor patches is facilitated when these are aligned with contours in the surrounding image (Neri, 2014). In both cases, attentional control can either directly, in the case of visual primes, or through recruitment of previous knowledge, activate either an attentional template or a set of visual expectations to facilitate object detection. The effects of such cues were usually examined immediately after cue presentation, revealing a transient temporal profile of the effect of hints on perception (Summerfield & De Lange, 2014; Sun et al., 2015); the same visual input can be interpreted in different manners depending on the hint that preceded it.

Another source of top-down expectation is prior experience, which allows previous encounters with the stimuli to facilitate perception by partial recollection of the stimulus to guide attention to the most relevant features (Pearson & Brascamp, 2008). This is in line with the notion of memory schema (Marlieke T.R. van Kesteren et al., 2012), in which information that can be associated with prior knowledge (the memory schema) is more likely to be consolidated, updating this schema (Tse et al., 2007). Such prior knowledge effects can facilitate understanding of movies and stories (M. T. R. van Kesteren, Fernandez, Norris, & Hermans, 2010), face recognition (Liu, Grady, & Moscovitch, 2018) and associative representations (Greve et al., 2019; Ozubko, Moscovitch, & Winocur, 2017). Using a series of stimulus presentations, these studies show that prior knowledge, acquired during a training phase, affects performance in the following perceptual task. An interesting observation is that prior knowledge may accumulate over time, making perception more vivid with each encounter (Godbole, Delaney, & Verkoeijen, 2014; Liberman, Fischer, & Whitney, 2014; Pearson & Brascamp, 2008). This increased vividness occurs in parallel with the effect of visual imagery on perception. In the case of imagery no visual input is presented, but memory of the previously encountered event can be used by the participant to perceive

a certain stimulus in the mind's eye (N. Dijkstra, Zeidman, Ondobaka, Van Gerven, & Friston, 2017). Self-reported vividness ratings during imagery have been shown to predict future performance in perceptual tasks (Pearson, Rademaker, & Tong, 2011), and neural activity in primary visual cortex (Fulford et al., 2018) and even pupillary dilation (Zokaei, Board, Manohar, & Nobre, 2019). These findings suggest a long time course of top-down effects, as memory can affect perception over multiple presentations of the stimulus, through encoding and decoding cycles. In addition, the perceptual experience may become more vivid over repetitions as the memory schema evolves and becomes more detailed.

Several of studies have used two-tone images to explore the effects of expectation on perception (Moore & Cavanagh, 1998). Two-tone images (also known as 'Mooney figures') are images which are reduced to black and white presentations, usually by adjusting the contrast of a grey-level image (see an example in Figure 1). When first encountering these images, it is often very hard to detect the object in the image, but after observing the greyscale image or after looking at the image for a long time it is possible to detect the object. A number of studies utilised the fact that the two-tone images remain the same before and after observing the greyscale image, i.e. their low-level visual features are unchanged. It was therefore possible to use these images to show how activation patterns in early visual cortices are affected the meaning of the visual stimuli (Hsieh, Vul, & Kanwisher, 2010), and to detect the neural basis of top-down effects on perception (Chang, Baria, Flounders, & He, 2016; Dolan et al., 1997; Hegde & Kersten, 2010; Samaha, Boutonnet, Postle, & Lupyan, 2018). Such studies usually examined object detection immediately after the informational cue (greyscale version of the two-tone image) was presented. In addition to the study of object recognition, two-tone images have been used to examine the effects of prior exposure to the object of the two-tone image, usually via a greyscale image (Chang et al., 2016; Dolan et al., 1997; Teufel, Dakin, & Fletcher, 2018), or through a verbal hint (Samaha et al., 2018) on performance in perceptual tasks, such as edge detection or similarity judgements. These studies demonstrated a top-down effect of cues that resolved

the ambiguity of the two-tone images either in immediate object detection accuracy or in consequent performance in perceptual tasks. However, the source of the effect, i.e. previous experience or informative hint, was not directly considered.

Here we suggest that it is possible to illuminate the specific contribution of hints and repetition-based top-down expectation effects through their consequences for perceptual experience over time. We used multiple presentations of two-tone images which included hard to see objects presented with and without verbal hints (such as “An elephant with its trunk in its mouth”, see Figure 1). Unlike previous studies, participants did not observe the greyscale images, only the two-tone images (with the exception of experiment 5). In addition, in our study participants observed the stimuli over repeated presentations, allowing us to observe their perceptual experience (through vividness ratings) in a continuous manner, going beyond a one-shot binary object detection.

This design allowed us to evaluate different hypotheses regarding top-down effects on perception, derived from informative hints or previous experience. First, the use of verbal hints has been shown to trigger attentional templates that facilitated visual search and detection (Nako et al., 2015; Samaha et al., 2018; Summerfield et al., 2006). Our design allows us to test the hypothesis that different types of verbal hints can facilitate object detection in two-tone images, in line with the attentional template account, and compare their effect with that of visual greyscale cues. Second, our design permits examination of the hypothesis that, after the object has been detected, object vividness will gradually increase over multiple repetitions, as a schema of the ambiguous image is formed and updated. Finally, our design allowed us to test whether the temporal effect of hints are transient effects, as suggested by the attentional template account, or sustained effects of repetition carried from one presentation to the next.

Experiment 1

Our first experiment was designed to examine the effect of verbal hints on object detection and perceptual experience over two stimuli presentations. We expected verbal hints to facilitate object detection and decrease the amount of time it takes participants to detect the object in the two-tone image. We were also interested in the subjective perceptual experience of participants, and to examine how it evolves over time – whether similar vividness is experienced in all conditions after the object is detected, or whether the experience is affected by the detail level of expectations, for example when a semantic hint is more detailed or when observing the image for a second time and utilising the memory of the previous presentation.

Methods

We recruited participants through the Amazon Mechanical Turk online platform. All participants provided informed consent. All experiments were approved by UCL Research Ethics Committee (project ID 5375/001). Participants earned a fixed monetary compensation of 4\$ for their participation, which lasted 20 minutes on average. We tested 45 participants in experiment 1 (28 male participants, age (mean $\pm$ std): 34.8 ± 7.37). This sample size was based on an exploratory expected effect size of $d_z = 0.5$, alpha of 0.05 and power of 0.9 in a mixed-effects design using the power analysis calculator provided by Judd et al. (Judd, Westfall, & Kenny, 2017).

We used natural pictures of animals or humans as stimuli in our experiments. All images were obtained through Google and Flickr searches for public domain images (CC0) of animals and humans. We then resized the images to have a similar height of 300pixels, and binarized them according to pixel intensity using Matlab's image processing toolbox (`im2bw` function, Matlab R2016a, MathWorks, Natick, MA), using a threshold of 0.25, i.e. pixels with greyscale value of 0.25 or higher were set to white and pixels below the threshold were set to black. The images were screened by independent observers and images which contained no information (for example, some images were completely bleached by the transformation), and images in which the objects were immediately identified were removed from the set. 30

images were presented in random order, 10 with no hint, 10 with a general verbal hint (general hint), and 10 with a detailed verbal hint (detailed hint). The general hints indicated the category of the animal, e.g. 'farm animal'. The detailed hints indicated the specific animal in the image and its posture, e.g. 'an elephant with its trunk in its mouth'). All images and hints are available at: <https://figshare.com/s/bb2b0059a6bf86ec1bcd> .

Participants were presented with 30 two-tone images. Each image was presented twice. Image presentation was self-paced. The participants were instructed to look at the images until they identified the object in the image, and press a Got-It button to move on to the rating stage. After 10 seconds a Give-Up button appeared, allowing participants to move on even if they failed to detect the object in the image. Participants then rated the vividness of their perception of the object on a six-point scale (see figure 1A), unless they had pressed the Give-Up button. Specifically, participants were told that the lowest rating meant 'I did not see the object at all', while the highest rating meant 'The object was very vivid!'. In the case of giving up the vividness rating was automatically set to zero. All the images were presented for a second time in a different order, but with no hints.

To further motivate participant to correctly identify the objects in the images, participants were told that four images would be randomly selected to form an identification quiz at the end of the experiment. Correct answers to the quiz would gain them a bonus payment of 1\$. All the experiments reported here ended with the 4-item quiz.

A. Self Pace Experiments

B. Intervals Experiments

Figure 1 – Presentation of images

(A) In Experiments 1 and 2, the images were accompanied with either a verbal detailed hint in which detailed information about the object in the image was given, a verbal general hint in which a general category information was given, or no hint. Participants were asked to press the Got-It button when they recognize the object in the image, or the Give-Up button that appeared 10 seconds after the image presentation. They were then asked to give a vividness rating on a six-point scale.

(B) In experiments 3 and 4 images were presented for 2.5 seconds and were subsequently rated for vividness using a sliding scale (the cursor appeared with the first mouse click on the scale).

Results

Our experiment used multiple presentation of each image, randomly assigned to hint conditions. To account for the random effects of both image identity and participants, we used a mixed effect linear model with two random effects; subjects and images; and hint level (3 nominal levels), repetition and their interaction as fixed effects (Baayen, Davidson, & Bates, 2008; Judd et al., 2017). We fitted these models using Matlab's `fitlme` function.

We initially fitted a model to the time spent with the images, i.e. how long before the participants pressed the Got-It or Give-Up button (Figure 2A). We found a significant decrease in time spent with image for general hints (estimate $\pm$ se: -2.1 ± 1 seconds, $t(2694) = -2.07$, $p = 0.04$) and detailed hints compared with no-hint (-3.5 ± 1 seconds, $t(2694) = -3.45$, $p = 0.0006$), and a significant repetition effect (-3.1 ± 0.45 seconds, $t(2694) = -6.85$, $p < 10^{-5}$) as participants spent less time with the images when they observed them for the second time. The interaction effect was not significant.

We then fitted a model to the vividness ratings made by the participants. We did not include the 'give-up' trials as participants did not provide ratings in these trials (Figure 2B). We found that vividness ratings increased from the first presentation to the second (0.52 ± 0.08 , $t(2402) = 6.12$, $p < 10^{-8}$), and that detailed hints lead to significantly higher vividness ratings than no-hints (0.7 ± 0.19 , $t(2402) = 3.78$, $p = 0.00016$). Interaction effects were not significant.

Finally, we examined the likelihood of Give-Up responses in each experimental condition using a logistic mixed effect linear regression (Figure 2C). Compared with no-hint images, participants had less Give-Up trials when they had a general hint (-0.95 ± 0.27 , $t(2694) = -3.45$, $p = 0.0006$) and when they had a detailed hint (-1.49 ± 0.35 , $t(2694) = -4.25$, $p < 0.0001$). They were also less likely to miss an image on the second presentation compared with the first presentation (-0.31 ± 0.11 , $t(2694) = -2.84$, $p = 0.005$). However, repetition did not benefit participants with general hints, as captured by a significant interaction effect for repetition and general hint (0.34 ± 0.17 , $t(2694) = 1.97$, $p = 0.05$). The interaction with detailed hints was not significant.

Discussion

The effect of hint level suggests that top-down expectations increased participants' perceptual experience (Carlisle, Arita, Pardo, & Woodman, 2011). We observed an effect of hints and repetition on object detection, captured by the reduction in Give-Up trials. In addition, a more gradual effect on visual experience, beyond object identification, was observed as time spent with images reduced and vividness ratings increased over hint types and repetition. The repetition effect suggests a gradual schema-driven process in which objects become more vivid with repeated presentations. The hint level effect suggests an attentional-template based effect, in which semantic information is used to set expectation immediately after the cue. Interestingly, the attentional template effect was carried to the second presentation and affected vividness and time with image on the subsequent presentation, even though the hint was not presented prior to the second presentation.

Experiment 2

In this experiment we wanted to further differentiate between hint effect and repetition effect by making the images less identifiable in the second presentation. For the effect of the hint given on the first presentation to carry over to the second presentation, when no hint was given, participants had to recognise, at some level, that this was the image that they had seen before and remember the hint that went with it. This requires integration of the low-level visual features of the image with high-level process associated with the hint, in line with the memory schema hypothesis (Marlieke T.R. van Kesteren et al., 2012). We hypothesized that, if object identification is enhanced by repeated examination of the image, then inconsistency in the way the image is presented should hinder this process, as it will make it harder to recognise the image and activate the appropriate memory schema. To explore this hypothesis, we adapted the experimental design by creating inconsistency in the presentation of the low-level features.

Methods

Experiment 2 was identical to experiment 1, except that the images in the second presentation were L-R mirror reversed (i.e. flipped around the vertical axis). We recruited 45 online participants for this experiment to match the sample size of experiment 1. We analysed data from 41 participants in experiment 2 (26 male participants, age: 37.3 ± 9.6), as 4 participants did not finish the experiment.

Results

The data were analysed using a mixed effect linear model as in experiment 1, with repetition, hint type, and their interactions as fixed effects and participant and image as random variables. Analysing the time spent with images (Figure 2D) we found a significant repetition effect (-4 ± 0.62 sec, $t(2394) = -6.44$, $p < 10^{-9}$), and a significant reduction in time with the image with general hints (-6.5 ± 1.4 sec, $t(2394) = -4.66$, $p < 10^{-5}$) and with detailed hints (-7.65 ± 1.4 sec, $t(2394) = -5.47$, $p < 10^{-7}$). However, as we expected, there was an attenuation of the effect of hints in the second, horizontally flipped, presentation. A significant interaction between repetition and general hints 3.29 ± 0.88 sec, $t(2394) = 3.72$, $p = 0.0002$) and repetition and detailed hint (3.16 ± 0.88 sec, $t(2394) = 3.58$, $p = 0.00034$) showed that as hints became more detailed, and resulted in lower engagement time in the first presentation, the more they suffered from inconsistency in presentation in the second presentation.

We examined the vividness ratings (Figure 2E), excluding trials in which participants used the Give-Up button, using a mixed effect linear model. We found a significant effect of repetition (0.37 ± 0.086 , $t(2048) = 4.34$, $p < 0.0001$), as participants rated the images as more vivid in the second presentation. A significant general-hint effect was observed (0.41 ± 0.19 , $t(2048) = 2.15$, $p = 0.031$) as well as a significant detailed-hint effect (0.57 ± 0.18 , $t(2048) = 3.06$, $p = 0.0022$). We also observed a significant repetition and general-hint interaction (-0.24 ± 0.12 , $t(2048) = -2$, $p = 0.045$) as ratings for the general hint did not increase with repetition. The repetition and detailed hints interaction was not significant.

We examined the rate of Give-Up trials using a mixed effect logistic regression (Figure 2F). Participants pressed the Give-Up button less when they had a general hint (-0.7 ± 0.28 , $t(2394) = -2.54$, $p = 0.012$) and when they had a detailed hint (-1.98 ± 0.32 , $t(2394) = -5.87$, $p < 10^{-8}$) compared with no hints. They were also less likely to miss an image on the second presentation compared with the first presentation (-0.28 ± 0.12 , $t(2394) = -2.31$, $p = 0.021$). However, we observed a marginal general-hint and repetition effect (0.32 ± 0.18 , $t(2394) = 1.82$, $p = 0.07$) and a significant detailed-hint and repetition effect (0.81 ± 0.2 , $t(2394) = 3.94$, $p = 0.0001$). As demonstrated in figure 2F, Give-Up rates increased from the first to the second presentation of images with detailed-hints, an effect demonstrated in the interaction effect.

Figure 2- Effect of hints on object identification

Time spent with images (A), vividness ratings (B) and Give-Up rates (C) in experiment 1. Participants spent less time with images, gave generally higher ratings, and were more likely to identify them when the images were accompanied with hints in the first presentation, and when viewing an image for the second time without hints. In experiment 2, in which images were horizontally flipped in the second presentation, the same main effects of time (D), ratings (E), and Give-Up rates (F) were obtained but with interactions with hinted images, as presentation inconsistency affected hinted images more than non-hinted images. This is most apparent in the increase in Give-Up rates for detailed-hint images when images were flipped. This effect was verified by directly comparing Experiment 1 and 2 results. These effects highlight the dependency of repetition effects on image recognition, as presentation inconsistency attenuated the repetition effects. This attenuation was apparent when images were initially presented with hints, indicating that attentional-hint effects were most affected by the presentation inconsistency. Error bars represent standard error of the means (SEM).

To directly evaluate the detrimental effect of presentation inconsistency on identification visual experience we included participants from both experiment 1 and 2 in one model, which included an experiment fixed effect in addition to the repetition and hints effects (participants random effect was nested in the inconsistency effect), as inconsistency was expected to generate a triple interaction between hint, repetition and experiment. Here we report only the effects associated with experiment effect, as hints and repetition effects were consistent across experiments and were significant here as well.

Comparing time with images across experiments we found a significant triple interaction between general hints, repetition and experiment (2.72 ± 1.07 sec, $t(5088) = 2.52$, $p = 0.011$), and a significant triple interaction between detailed hints, repetition and experiment (2.62 ± 1.07 sec, $t(5088) = 2.44$, $p = 0.014$), showing that inconsistency in the second presentation increased the time spent with hinted images. These effects were accompanied with a marginal overall experiment effect (2.49 ± 1.37 sec, $t(5088) = 1.83$, $p = 0.07$), indicating an overall increase in time spent with images in the second experiment, and significant general-hint and experiment 2 interaction (-4.3 ± 1.7 sec, $t(5088) = -2.56$, $p = 0.01$) and a significant detailed-hint and experiment 2 interaction (-4.13 ± 1.7 sec, $t(5088) =$

-2.42, $p = 0.015$), indicating that the overall increase in time in experiment 2 was mostly driven by the non-hinted images.

Analysis of vividness ratings (excluding Give-Up trials) did not reveal any significant effect of presentation inconsistency, beyond the shared effects of repetition and hints.

Analysis of Give-Up trials using a mixed effect logistic regression revealed a significant triple interaction between detailed hints, repetition and experiment (0.65 ± 0.3 sec, $t(5088) = 2.13$, $p = 0.03$), indicating an increase in Give-Up rates when images were originally presented with a detailed hint and then horizontally flipped. This direct comparison indicates that presentation inconsistency, which led to attenuation in object recognition in the second presentation, limited the beneficial effect of attentional expectations.

Discussion

Inconsistency in the presentation of low-level features hindered the improvement in object identification and time with images, but only when verbal hints had been given on the first presentation. Indeed, in the conditions with hints, it was as if participants were seeing the image for the first time (though not down to the level of seeing it with no hints). This was most apparent in the Give-Up rates of the detailed-hint images, which increased from the first presentation to the next when images were flipped, but not when presentation was consistent. One explanation might be that, in these conditions, participants had to exert more time and effort to use a memory schema to link the first presentation with the second presentation properly, which attenuated their increase in visual experience. Importantly, the increase in Give-Up rates on the detailed-hint condition in the second experiment may indicate the temporal limit of hints, as their beneficial effect was not carried over to the second presentation.

Experiment 3

Experiments 1 and 2 revealed a gradual increase in perceptual experience, which continued beyond object recognition. However, the time spent with images was affected by the hint

level, and so were the vividness ratings. To better characterise the gradual increase in vividness ratings and perceptual experience, and to better differentiate between the transient hints effect in the first presentation and the effect of repetition on consequent presentations, we changed our experimental design in Experiment 3. Here we used multiple presentations of images for fixed duration of time, and asked participants to report how vivid was the image. We used only two levels of hints, no-hints and detailed hints, as they seem to capture the main effects of interest.

Methods

We presented 20 two-tone images (only animals) in 6 blocks, so each image was presented 6 times, with each presentation being fixed at 2.5 seconds. The order of image presentation in each repetition was random. The images were first presented with either a detailed hint or no hint. Images were not accompanied with hints after their first presentation. Each presentation was immediately followed by a vividness rating. Ratings were made using a continuous sliding scale, ranging between 0 and 50 (Figure 1B). The scale marker appeared only after the first mouse click on the scale, to prevent anchoring of answers. We used the effect size from experiments 1 and 2 to set the number of participants in the following experiments. We calculated an effect size of 0.6 for the hint effect, and adjusted our sample size to be $n=22$ (Judd et al., 2017). We recruited 27 online participants, and analysed data from 22 participants (14 male participants, age: 37.1 ± 11.5), as 5 participants did not finish the experiment online due to technical issues.

As we used multiple presentations, it was possible to examine the shape trajectory followed by the vividness ratings. A gradual increase in average vividness ratings could stem from gradual increases in the individual level, i.e. everyone shows the same gradual increase, or could be an artefact of averaging over many participants. Participants could display a step-function in their vividness response – giving low vividness ratings until they recognize the object and then switching to high vividness ratings. Averaging across participants with different points of switch will result in smooth increase in the group average (Chen &

Krajbich, 2017). To evaluate these possibilities, we fitted individual participants' vividness ratings trajectories to a step function:

$$Ratings = a \cdot \begin{cases} 0 & Interval < b \\ 1 & Interval \geq b \end{cases}$$

Where a and b are free parameters, b indicates the interval in which detection was made, and a is a scaling parameter, indicating the final level of ratings reached. And to a linear function:

$$Ratings = c + d \cdot Interval$$

Where c is the intercept and d is the slope of the trend. A comparison of model fits would indicate whether individual ratings were better described as step function or gradual linear increase.

Fitting procedure was carried for each participant and for each image independently using Matlab's `fminsearch` function.

Results

The vividness ratings were analysed with a mixed-effect linear model, with repetition number and hint as fixed effects, and participant and image as random effects (Figure 3A). We found a significant repetition effect, as vividness ratings increased with image repetition (2.13 ± 0.21 , $t(2636) = 10.13$, $p < 10^{-10}$). We also found a significant effect of hints, with hinted images getting higher vividness ratings (7.65 ± 1.16 , $t(2636) = 6.55$, $p < 10^{-10}$). In this experiment we also observed a significant interaction effect, as the increase in vividness ratings was more pronounced in the 'no-hint' condition (-1 ± 0.3 , $t(2636) = -3.36$, $p = 0.0007$, see Figure 3A). Across the 6 presentations, we observed smooth increase in group vividness rating, which is consistent with our hypothesis that object perception emerges gradually. To make sure this was not an artefact of averaging over many participants, we fitted step-function and linear function to participants' ratings. We found that the mean

squared error of all participants was higher for the step function, indicating a better fit of the linear model for all participants trajectories (paired t-test: $t(21) = 7.15$, $p < 0.00001$).

Discussion

In Experiment 3 we observed a similar pattern of results to those obtained in Experiment 1. First, hints increased the vividness ratings, i.e. the perceptual experience of the participants, in the first presentation. This increased vividness persisted in follow up presentations, indicating that the benefit from hint effect was also encoded in memory for future encounters. As the experiment lasted for multiple presentations, we were able to reveal an interaction effect, as vividness ratings seem to slow down as the image becomes more and more vivid. These results served as a basis for our next investigations, in which the effect of hints was further investigated.

Experiment 4

In Experiment 4 we aimed to explore the role of previous encounters in the gradual increase in perceptual experience across repeated presentations. We increased the number of images from 20 to 30. This meant that there were more images to track over time, and the intervals between presentations were, on average, longer. We hypothesised that this increased memory load would cause the vividness of the percept to emerge more slowly. In addition, we wanted to evaluate the attentional effect. As beneficial hints were shown to have long lasting effect on vividness ratings, we examined whether wrong hints, i.e. hints that contain wrong information about the image, may have long lasting detrimental effects.

Methods

In experiment 4 thirty images were used, 10 of which were not accompanied with hints. In this experiment hints were introduced as questions before observing the image, e.g. 'is this a picture of a jumping kangaroo?', to which the answer was 'no' for half of the presentations. Participants were told that the hints were either correct or wrong (10 wrong and 10 correct hints). After the first presentation all images were presented for an additional four times

without hints, and vividness ratings were collected after each presentation. Vividness was rated as in experiment 3. We recruited 24 participants in experiment 4 (9 male participants, age: 33.95 ± 8.44), to match the sample size of experiment 3.

Results

We first confirmed that individuals' ratings of each image were better described by linear model rather than a step model, using the same fitting procedure as in Experiment 3. We again observed that the mean squared error of all participants was higher for the step function, indicating a better fit of the linear model to all participants trajectories (paired t-test: $t(23) = 8$, $p < 0.00001$).

While initial vividness ratings were very similar to the ratings from experiment 3, the rate of increase in ratings was slower – with no increase in the 'correct hints' condition (Figure 3B). Using a mixed-effects linear model with three levels of hints (none, wrong and detailed) and repetitions as fixed effects, we found a significant effect of repetitions (1.19 ± 0.23 , $t(3594) = 5.21$, $p < 10^{-6}$). Wrong hints had a significant negative effect on vividness ratings (-3.38 ± 1.08 , $t(3594) = -3.13$, $p = 0.0017$), while correct hints had a significant positive effect on ratings (8.1 ± 1.08 , $t(3594) = 7.49$, $p < 10^{-13}$). In addition, we found a significant interaction between wrong hints and repetitions (0.67 ± 0.32 , $t(3594) = 2.05$, $p = 0.04$), indicating a more steep increase in ratings over repetition compared with no-hints. A significant interaction between correct hints and repetitions (-0.9 ± 0.32 , $t(3594) = -2.76$, $p = 0.006$) indicated that correct hints led to more gradual increase in vividness ratings compared with no-hints.

To directly assess the effect of increase of number of images and memory load on vividness ratings, we fitted the ratings from the first 5 presentations of experiment 3 and the correct-hint and no-hint images ratings from experiment 4 with one mixed effect linear model, which included repetition, hint type, experiment and their interactions as fixed effects, and image number and participants as random effects (participants were nested within experiments).

We found a significant hint effect (9.07 ± 1.29 , $t(4592) = 6.98$, $p < 10^{-11}$), a significant repetition effect (2.5 ± 0.27 , $t(4592) = 9.06$, $p < 10^{-18}$), and a significant negative interaction between them (-1.39 ± 0.39 , $t(4592) = -3.55$, $p = 0.0004$) as hinted images showed slower increase in vividness compared with non-hinted images. We did not observe an overall experiment effect, but observed a significant negative repetition and experiment interaction (-1.31 ± 0.38 , $t(4592) = -3.42$, $p = 0.0007$). This effect indicates that the increase in vividness ratings over repetitions was slower in experiment 4, which included more images than experiment 3 and therefore increased memory load.

Discussion

The findings from experiment 4 reveal the effects of hints over time, and how memory constraints affect perceptual experience. Expectations based on hints showed a marked effect in the first presentation, with correct hints leading to higher vividness ratings and incorrect hints leading to lower vividness ratings. However, these initial effects did not transfer completely to future presentations, indicating a transient profile of hint-based expectations. Images that were hinted with correct hints did not increase in vividness so much, and even suffered a slight decrease in vividness in the second presentation. In addition, wrong hints did not hinder object identification for very long, converging with the ratings of the 'no hint' condition after 3 repetitions. We increased the number of images in the experiment from experiment 3, which meant that more images had to be maintained in memory, and that the average time between presentations of images was longer compared with experiment 3. The increased memory load led to a slower increase in the level of vividness, demonstrated by the interaction between repetition and experiment when experiments 3 and 4 were compared directly. However, in all cases after the initial effects of hints in the first presentation a gradual increase in vividness ratings was observed, indicating the contribution of memory-based expectations to perceptual experience.

Figure 3 – The temporal dynamics of hints effect on identification

(A) The vividness ratings in Experiment 3 (20 images) started higher for the hinted images, and increased with repetition. The ratings in the non-hinted images initially increased at higher rate than for the hinted images, but were always lower than the hinted images.

(B) The vividness ratings in experiment 4 (30 images) showed a different pattern. Initially there was a large effect of hints, with correct hints being better than no hints, while incorrect were worse. There was no increase in vividness across repetitions in the correct hints condition and no difference between the incorrect and no hint conditions after the first presentation.

(C) The vividness ratings in experiment 5 (30 images) showed a similar effect in the hinted and non-hinted images. However, priming images with their greyscale version led to very high vividness ratings initially, which dropped dramatically in the following presentations and converged with the ratings of the hinted images.

(D) The vividness ratings in experiment 6 (30 images) showed a similar effect in the hinted and non-hinted images. However, very noisy and non-hinted images received very low vividness ratings, which remained low throughout the experiment, ruling out the effect of familiarity on ratings, and highlight the link between vividness ratings and object identification. Error bars represent standard error of the means (SEM).

Experiment 5

As noted in the introduction, previous studies used greyscale primes to evaluate the effect of hints and previous knowledge on detection of two-tone images (Neri, 2014; Teufel et al., 2018). In experiment 5 we wanted to evaluate how verbal hints compare with visual cues. We hypothesised that visual cues will behave as a 'super-detailed' hints, completely disambiguating the two-tone images, and therefore will lead to high vividness ratings immediately after presentation. However, examining how well this effect carries over to future presentations will indicate the temporal differences in hint-based and repetition-based expectations.

Methods

To compare the effect of verbal hints to greyscale versions of the two-tone images, we followed the same design as in experiment 4, with 30 images. Ten of these images were not accompanied with any hints, and ten were accompanied with hints as before. The other ten images were preceded with a greyscale version of the image. Hints were presented for 3 seconds before the presentation of the two-tone images, and greyscale images were presented for 1 second before the presentation of the two-tone images. After the first presentation all images were presented for an additional four times without hints, and vividness ratings were collected after each presentation. We recruited 25 participants in experiment 5 (11 male participants, age: 33.95 ± 8.44) to match the sample size of experiments 3 and 4.

Results

We first confirmed that individuals' ratings of each image were better described by linear model rather than a step model. As in experiments 3 and 4, we observed that the mean squared error of all participants was higher for the step function, indicating a better fit of the linear model to all participants trajectories (paired t-test: $t(24) = 8$, $p < 0.00001$).

We fitted a mixed-effects linear model to the vividness ratings, with hint (none, visual and verbal), repetition and their interactions as fixed effects, and images and participants as

random effects (Figure 3C). We found a significant effect of repetitions on ratings (1.06 ± 0.21 , $t(3744) = 5.05$, $p < 10^{-6}$), as vividness ratings increased over time. In addition we found a significant increase in vividness ratings which were initially accompanied with visual hints (10.19 ± 0.99 , $t(3744) = 10.26$, $p < 10^{-23}$), and verbal hints (5.01 ± 0.99 , $t(3744) = 5.05$, $p < 10^{-6}$). However, a negative interaction between visual-hints and repetitions (-1.72 ± 0.29 , $t(3744) = -5.78$, $p < 10^{-8}$) indicated that the initial increase in vividness did not persist, and was greatly attenuated on the second presentation. The interaction between verbal hints and repetition was not significant (-0.34 ± 0.3 , $t(3744) = -1.14$, $p = 0.25$).

Discussion

Experiment 5 replicated the results from the no-hint and verbal-hint conditions – there was an increase in vividness ratings over presentations, and hints led to higher vividness ratings than non-hinted images (although there was no interaction effect with verbal-hints).

Greyscale cues led to very high vividness ratings on the first presentation, but decreased dramatically on the second presentation, essentially converging with the ratings of the hinted images. After the second presentation ratings in the greyscale primed images continued to increase with repeated presentation, just as with the hinted images. These results indicate that the visual cues had a very strong and transient attentional expectation effect, which did not carry across presentations. However, the vividness ratings of visually hinted images did not decrease to the level of non-hinted images but converged with those of the verbally hinted images. This suggests that at least some of the perceptual gains made by the greyscale visual expectations were stored in memory and affected future interactions with the image. Importantly, memory-based expectations led to increased vividness ratings from the second presentation onwards. This recovery in vividness ratings suggests that through repeated exposure and recollection, memory-based expectations were formed.

Finally, initial vividness ratings were somewhat lower than the ratings in experiments 3 and 4, maybe due to an anchoring effect of the greyscale images, which also may eliminate the interaction effect between verbal hints and repetitions.

Experiment 6

So far, our results demonstrated a differentiation between the transient effects of initial hints and the gradually increasing effects of memory when revisiting the images over multiple presentations. In Experiment 6 we tested the possibility that familiarity alone, without increase in perceptual experience, may account for the gradual increase in vividness ratings. We therefore followed the same design as in experiments 4 and 5, but introduced a new condition with no-hinted noisy two-tone images, in which object identification was very difficult, in addition to the non-hinted two-tone images with identifiable objects and hinted, identifiable two-tone images we used in previous experiments.

Methods

We followed the same design as in experiment 4 and 5, with 30 images. Ten of these images were not accompanied with any hints, and ten were accompanied with hints as before. The other ten images were taken from the study of Samaha et al. (Samaha et al., 2018). The two-tone images we used from their study were the 10 images which got the lowest free identification ratings (below 10%). We used these images and not random noise as these were more similar to the images we used before in terms of spectral frequencies and therefore form a strongest criterion to verify the familiarity effect on ratings than random noise. We recruited 25 participants in experiment 5 (14 male participants, age: 36.72 ± 8.69) to match the sample size previous experiments.

Results

We first confirmed that individuals' ratings of each image were better described by linear model rather than a step model. As in previous experiments, we observed that the linear model gave a better fit to the vividness ratings than the step function (paired t-test: $t(24) = 6.4$, $p < 0.00001$). We then evaluated the vividness ratings using a mixed effect linear model, with repetitions and hint levels (no-hint, detailed hint, noise) and their interactions as fixed effects and images and participants as random effects (Figure 3D). We observed a

significant repetition effect (1.77 ± 0.24 , $t(3744) = 7.33$, $p < 10^{-12}$) as vividness ratings increased over repetitions. Detailed hints increased the vividness ratings (5.61 ± 1.12 , $t(3744) = 5.01$, $p < 10^{-6}$), while noisy images had a significantly lower vividness ratings (-10.38 ± 1.12 , $t(3744) = -9.22$, $p < 10^{-19}$) compared with no-hint identifiable images. In addition, a significant interaction between detailed-hints and repetitions (-0.72 ± 0.33 , $t(3744) = -2.15$, $p = 0.032$) indicated an attenuation of the increase in vividness ratings for hinted images compared with non-hinted images. A significant interaction between noisy images and repetitions (-1.43 ± 0.34 , $t(3744) = -4.23$, $p < 0.0001$) showed that the increase in vividness ratings for noisy images was even more attenuated.

Discussion

Experiment 6 demonstrated that the increase in vividness ratings is not purely based on familiarity, i.e. the fact that the participants observed the same image multiple times. In addition, images that were very hard to identify, as the ones used in the noisy condition, received consistently low vividness ratings, well below the non-hinted identifiable images which were used in all our experiments so far. This indicates that vividness ratings represent the participants' perceptual experience, which include identification of the object in the image and can increase as the experience becomes more detailed and vivid.

Experiment 7

The results from our experiments so far suggested that vividness ratings corresponded with perceptual experience and the identification of the object in the two-tone image. However, experiments 3-6 used vividness rating as a proxy for perceptual experience of the correct object. While experiments 1-2 indicated that vividness is linked with object detection, as indicated by the Give-Up rates, and experiment 6 demonstrated that unidentifiable images were given low vividness ratings, it is possible that participants might perceive an object that was not actually present in the original image. This certainly happened occasionally in an informal experiment in which participants were asked to draw round the object that they had

perceived in the image. For example, one participant found an octopus in the image of a kangaroo (Figure 4). Our last experimental investigation therefore sought to provide another validation of the vividness ratings as an indication that viewers had correctly detected the object in the image.

Methods

Figure 4 – Examples of Object Detection

In an informal experiment we presented participants with the two-tone images (on the left) and asked them to draw round the object they detect. Most participants correctly depicted the object -kangaroo in this example drawn by one of our participants (middle), while a few detected very different objects, like in the example from another participant who drew an octopus (right).

In experiment 7 participants observed the 20 images from experiment 3 for 2.5 seconds each, either accompanied with a detailed verbal hint or no hint. Each image was observed once. After the image presentation, participants were asked to rate the vividness of the object in the image. After rating, participants were asked to colour in the object that they perceived in the image. In this experiment we collected data from 36 participants (20 male participants, age: 39.72 ± 13.29) as we estimated a lower effect size of $d_z = 0.5$, with 90% power, for differences in *accuracy* of drawings.

We collected the drawings and calculated an overlap index by comparing the drawing area with the ground truth – the actual area of the animal depicted in the image (Figure 5A). We divided the intersection of the participant's drawing and the true area of the object, by the

union of these areas to generate the overlap index. This index is close to 1 when the participant's drawing matches perfectly the object in the image, and it is close to 0 when the drawing does not overlap with the object at all. Note that when the drawing overlaps with the object in the image, but includes many irrelevant parts as well, the index will decrease from 1 towards 0, like in the example in figure 5A.

Results

We used a paired t-test to examine the effect of hints on overlap index and ratings (Figure 5B, C). As in the previous experiments we found a significant effect of hints on vividness ratings (Hint (mean $\pm$ std): 22 ± 7.8 , No Hint: 17 ± 7.2 , $t(35) = 4.7$, $p = 0.00003$, $d_z = 0.79$). In addition, we observed a significant increase in overlap index with hints (Hint: 0.56 ± 0.12 , No Hint: 0.48 ± 0.12 , $t(35) = 3.4$, $p = 0.0016$, $d_z = 0.57$). To verify that both effects correspond, we examined the correlation between each participant's hint effect on ratings and hint effect on overlap index and found a significant correlation (Figure 5D, $R = 0.47$, $p = 0.004$).

Figure 5 – Object identification and vividness ratings

(A) Participants' drawings (example from one participant in red) were compared with the actual object area (ground truth, example in green). An overlap index was constructed by dividing the intersection between the two areas (in yellow) by the union of the areas (sum of blue and yellow). The closer the index was to 1, the more accurate the participant's depiction of the object. (B) We found a significant effect of 'Hints' on overlap index such that drawings of hinted images tended to have higher overlap index (error bars represent SEM, ** $p < 0.005$). (C) We also observed a significant effect of hints on vividness ratings, just as in the previous experiments (error bars represent SEM, **** $p < 0.00005$). (D) We found a significant correlation between the effect of hints on ratings and the effect of hints on overlap index. Each dot represents one participant, and the dashed lines represent 2 SEMs around the correlation line. Error bars represent standard error of the means (SEM).

Discussion

The observation that hinted images were more accurately depicted than non-hinted images, and that this difference corresponded with the differences in vividness ratings confirms that vividness ratings are related to the ability to accurately identify the object in the image. The findings from this experiment suggest a strong link between vividness ratings and the ability to detect the object in the ambiguous images. Note that vividness ratings in this experiment were made immediately after the short presentation of the two-tone image, and that participants then had more time with the image while drawing the object. This could lead to reduced effect of hints in drawing accuracy. However, the fact that even in this case vividness ratings and drawing accuracy were strongly related suggests that vividness ratings were closely linked to perceptual experience and object identification.

Perceptual Experience Model

We observed some repeating patterns across the experiments, in which hint level and repetitions had positive effects on perceptual experience, and that these effects were attenuated by memory load and inconsistencies in image presentation. We hypothesised that these effects stem from a common object detection process, in which the interaction

between priors, formed by attentional verbal (or non-verbal) hints, and the information carried by the observed evidence from trial to trial via memory, interact to affect perceptual experience. To test our hypothesis, we formalised an object identification model, inspired by evidence accumulation models such as drift diffusion model (Ratcliff & McKoon, 2008). For the sake of simplicity, we used a simple world including only nine objects which belong to one of three categories (Figure 6A). However, it is possible to scale up this model using continuous instead of discrete representations to include a more elaborate object space, using the same evidence accumulation mechanism.

Our model assumed that a viewer holds a belief vector b , with values for possible objects in the image. The belief vector is formed by accumulating evidence from the image over time in a stochastic manner, with an update parameter γ . The belief vector was normalized after each update step.

$$[3] \quad b_t = b_{t-1} + \gamma Ev_t$$

Evidence were generated in a stochastic manner according to the information distribution of the image. Each image was represented by a vector of probabilities to elicit evidence in favour of different objects. On each time step the evidence for each object was generated from a Bernoulli distribution governed by the value of the image vector for this object.

$$[4] \quad Ev[i] \sim Ber(Image[i])$$

We assumed that an ambiguous image would generate evidence in favour of multiple objects. For example, an ambiguous image of a parrot might generate evidence describing a parrot, but might also provide category level evidence, such as beak or talons, that could be attributed to all birds. In addition, noisy parts of the image might provide evidence pointing to other unrelated objects, as seen in the informal examples of the kangaroo being perceived as an octopus (Figure 4). The more ambiguous the image, the lower the rate of object specific information it would elicit. We formalized this notion by generating each image

object probabilities vector as a normalised weighted sum of three vectors (Figure 6A). The first vector was set to 1 at one object, and 0 in for the rest of the objects, and represented the probability of observing evidence for a specific object in the image. The second vector was set to 1 in all objects with the same category as the specific object in the first object, which represented the probability of seeing a category level information. The last vector was set to 1 in all objects, and represented the basic level of getting evidence for every object. The specificity of the information in the image was governed by the ratio between the weights for these three vectors: $[\beta_{object}, \beta_{category}, \beta_{general}]$. Images with more pronounced β_{object} will emit evidence supporting a specific item most of the time. Images with pronounced $\beta_{category}$ will lead to lower accumulated evidence for a specific object, as all category objects will be promoted in a similar rate. Finally, the values in the belief vector in each time point were taken as a level of object identification, associated in our experiments with vividness ratings.

To demonstrate the ability of this model to reproduce the vividness temporal dynamics we found in our experiments, we simulated this model with images containing evidence probabilities for 9 objects, which were clustered in three categories (Figure 6A). Note that a similar mechanism could be extended to much more complicated semantic structure, using continuous instead of discrete probability distributions. Each process lasted 25 time-steps in which evidence was randomly generated according to the image evidence distribution and accumulated. We examined the accumulation of evidence supporting the object represented in the image (the vector multiplied with β_{object}) over time.

Hint level effects were introduced by setting the prior values of the belief vector. ‘No hint’ condition was modelled by a flat prior across objects (all objects = 0.11), ‘general hint’ condition was modelled by assigning similar values to all category items (category objects = 0.16, non-category objects = 0.08), and ‘detailed hint’ condition was modelled as high prior to a specific object (hinted object = 0.25, other objects = 0.09). The ‘wrong hint’ condition

was modelled as a 'detailed hint', but with high prior on the wrong object, resulting with lower prior for the correct object. Finally, greyscale prime was modelled as a very high prior (hinted object = 0.9).

Figure 6 – Model and simulations

To gain a better understanding of possible mechanism driving object identification and perceptual experience, we constructed a computational model. The model assumes a stochastic evidence accumulation process, in which evidence favouring objects are accumulated over multiple exposures to the image. (A) This process is governed by the probability of an image to elicit evidence in favour of items. This distribution is generated by the weighted sum of three vectors - object specific evidence, category level evidence, and noisy evidence. (B) We simulated the model with two sets of weights, a high noise set (left) and a low noise set (right). Hint-level was set by changing the prior beliefs for

each of the nine objects. Using the high noise weights set (left) we recovered the attenuated ratings observed in Experiment 4, and the decrease in vividness for greyscale images. Using the high-noise weights set (left) we recovered the behavioural pattern of vividness reports from Experiment 3, in which vividness increased over time in all hint levels except the greyscale image. (C) We sampled the simulated data to examine the pattern of early improvements, captured in Experiments 1 and 2. We found that the low noise weights set (right) led to the observed pattern of vividness ratings in Experiment 1, and the high noise weights set (right) led to the attenuated ratings in Experiment 2.

Our simulations recovered the patterns of behavioural results in experiments 1-6 (Figure 6 B, C). We used two sets of evidence weights, one with high category level noise ($[\beta_{object}, \beta_{category}, \beta_{general}] = [0.05, 1, 0.05]$), and another with low levels of noise ($[\beta_{object}, \beta_{category}, \beta_{general}] = [1, 0.4, 0.05]$). These represented different types of images, one which is more difficult to detect as it emits a low rate of item-specific evidence, and the other is more readily detectable as it emits high rate of object-specific evidence. We simulated the model with all hint types in both settings (Figure 6 B, C). We found that the model was able to recover the gradual increase in ratings over presentations in all hinted and non-hinted conditions. However, making the image noisier imposed a constrain on the level of vividness ratings achieved. This is somewhat similar to the effect of making the images less readily recognised using horizontal flip or increasing the number of images to remember, which imposed a memory constrain in experiments 2 and 4-6.

In addition, our model captured the decrease in vividness over observation when an image was presented with a very high prior (Figure 6B, red lines), in our case the greyscale image in experiment 5, and even the detailed hint in experiment 4. As updates to perceptual experience are based on evidence accumulation, vividness ratings become more similar to the information in the image, represented by the evidence weights, as these govern the rate and quality of evidence sampled by the observer. In such a case a prior which is more accurate than what the image permits will be attenuated over time

Taken together, our model suggests that perceptual experience may follow a stochastic and dynamic process of evidence accumulation, which is affected by attentional hints that form a prior for object perception, and by accumulation of the information acquired from each interaction with the image. The information captured from the image can be modulated due to the ambiguity of the image itself, or by the constraints due to image presentation, such as inconsistency in presentation. Manipulating the priors and image evidence probability distribution, while keeping the model intact, was sufficient to recover the behavioural patterns from the four experiments.

General Discussion

We investigated the contributions of expectations based on hints and repetitions on perceptual experience. By using ambiguous two-tone images over multiple repetitions, we were able to evaluate the evolution of perceptual experience over time, and distinguish between the initial effect of hints and the gradual effect of previous experience on object detection and perceptual experience. In six experiments we observed increased vividness rating with hints concerning the objects' identity, and increased vividness ratings with repeated image presentation. In experiment 1 and 2 object recognition was tracked directly, showing that hints and repetitions promoted object detection. In addition, the time spent with the image and vividness ratings were affected by hints and repetitions, suggesting that perceptual experience can increase even after object recognition, as it become more vivid. This observation was replicated in the follow up experiments which focused on vividness ratings over multiple, fixed-length, presentations. In all these experiments initial vividness ratings were higher when the image was accompanied with a verbal hint or a visual prime, and increased gradually from one presentation to the next (also indicated by our linear and step-function model fitting).

A number of observations indicated that expectations based on hints, either verbal or greyscale images, had only a transient effect, and wore off after their initial presentation,

when memory-based recollections were attenuated. In experiment 2, where images were horizontally flipped across presentations, we observed an increase in Give-Up responses to the verbal detailed hint images, indicating that object detection deteriorated when recognition of the initial image was harder. A similar attenuation effect was observed when the number of images to track was increased from 20 to 30 between experiments 3 and 4. This increase in memory load led to a slower increase in vividness ratings, most marked for hinted images. Finally, the initial boost in vividness ratings observed for images which were primed with a greyscale version of the image wore off after the first presentation, indicating that not all of the information obtained from the greyscale prime could be recovered when the image was encountered in a second presentation. An evidence accumulation model of perception was able to capture these different dynamic profiles, highlighting the role of hints in setting priors on perception, the role of memory in accumulating evidence across presentations, and the limits on top-down effects imposed by the noise level of the image.

Gradual increase in vividness ratings

In all experiments we found that our measures of perceptual experience kept on increasing even after the object in the image was identified. This was evident in experiments 1 and 2, where both object identification and vividness were measured, and in experiments 3-6 using the model fitting procedure to compare step-like vividness profile to gradually increasing profile. The gradual increase was more pronounced when the number of images in the set was small (in Experiment 3), and when the images were originally not accompanied by hints. This indicates the role of repetition-based expectations on visual perception, as the visual evidence is accumulated, stored, and recalled, over presentations, until all information is extracted and vividness plateaus. We refuted the possibility that this effect was solely dependent on familiarity with an image, as noisy images did not lead to increased perceptual experience. This process was captured by our evidence accumulation model.

We were able to uncover a gradual increase in vividness since we used an experimental design with repeated image presentations after the initial hint and used a continuous measure of perceptual experience. Previous studies tended to use only one presentation of an image after visual prime (Chang et al., 2016; Dolan et al., 1997; Hegde & Kersten, 2010) and could not detect gradual changes. In addition, these studies only tested recognition with a binary yes/no question, which could not detect any levels of vividness of recognition.

It is important to note the limitations of using vividness ratings as proxies for perceptual experience. Unlike other studies which examined whether prior knowledge enhance or bias performance in perceptual tasks (Neri, 2014; Pearson et al., 2011; Samaha et al., 2018; Summerfield et al., 2006; Teufel et al., 2018), vividness ratings rely on self-report of a perceptual experience. Such report may be inaccurate and vary across individuals (Fleming, Weil, Nagy, Dolan, & Rees, 2010). However, in the research of visual imagery, self-reports of perceptual experience were shown to provide consistent and meaningful description of subjective experience, to correlate with neural responses in visual cortex and the retina (N. Dijkstra et al., 2017; Nadine Dijkstra, Bosch, & van Gerven, 2017; Hertz & Amedi, 2014; Zokaei et al., 2019), and to predict performance in subsequent perceptual tasks (Pearson et al., 2011). In addition, vividness ratings did not directly measure the correct detection of objects in the images. As indicated in Figure 4, some participants did not identify the objects correctly. However, our results regarding Give-Up rates and the incentive given to participants to correctly detect objects (see Experiment 1 -methods) suggest that participants mostly identified the objects correctly. In addition, experiment 7 suggests that images that were assigned with higher ratings were also correctly delineated by the participants. Finally, in our results vividness ratings lead to consistent and replicable effect across 6 experiments and were shown to be related to the ability to depict the objects on the two-tone images. The usefulness of vividness ratings in detecting the dynamic increase in perceptual experience over repetitions was not compromised by the limitations associated with self-report.

The transient effect of hints

In our results show that semantic hints consistently lead to faster object identification and increased vividness, as an attentional template could be used as a basis for identification in the first presentation. However, this effect did not always transfer completely to subsequent image presentations. This was most apparent in experiment 5, where greyscale priming led to initial very high vividness ratings, which later decreased and converged with the ratings of the hinted images, which slowly increased over additional presentations. This pattern suggests that such hint-based predictions have short term effect. What determines if this initial boost will affect later presentations is the likelihood of recruiting the previous experience in future presentations from the image itself. If the previous experience was very detailed, but the image was very noisy, it will be hard to recover the previous experience, and vividness ratings will be attenuated. This was the case when image presentation was not consistent, and in the greyscale primes. However, this does not mean that all the boost from the hints was lost, as vividness ratings never fell to the level of non-hinted images. As our model suggest, this may be a result of limitation of vividness imposed by the image characteristics, how noisy it is, which will determine how much information can be carried from one presentation to the next.

Taken together, our results suggest that perceptual experience over time is shaped by interaction between expectations, formed by hints, and the memory of previous interactions with the image. The notion of predictive coding suggests a recurrent flow of information, as predictions are delivered from top to bottom to shape early visual areas, and bottom up information is used to update the models generating these predictions (Friston, 2010; Huang & Rao, 2011). Here we were able to characterise how two different sources of predictions, hints and repetition, affect perception. The role of attention in setting priors for perception has been highlighted in behavioural (Carlisle et al., 2011; Nako et al., 2015) and neuroimaging studies (Summerfield & Egner, 2009; Summerfield et al., 2006). The notion of attentional template, which sets a prior for detecting visual objects for example, was

introduced to capture these effects (Carlisle et al., 2011). Hint effects are short lived, as giving different hints before observing a bi-stable image, which includes two images imposed on one another, can elicit different response (Summerfield & Egnér, 2009). In the memory domain the notion of memory schema was introduced to capture the facilitation of detection after prior exposure to stimuli (Alba & Hasher, 1983; Marlieke T.R. van Kesteren et al., 2012). In this account, a stimulus is encoded in a memory schema which is retrieved and updated in each interaction with the stimulus, leading to more accurate description over time. Such dynamics can explain how perception of a bi-stable image can become biased over time, as each perceptual experience serves to enhance the memory schema and affect future encounters, and future deployment of attentional templates (Pearson & Brascamp, 2008). While the temporal profile of these two accounts is implicit in the experimental designs and measures used to examine these two processes, many studies do not address these differences explicitly. Here we were able to demonstrate how recollection and hints provide distinct contributions to perception, show how attentional hints may facilitate the formation of a detailed memory schema, and demonstrate the limits memory-based predictions impose on the extent of attentional-template effects.

Our framework, supported by behavioural results across multiple experiments and the evidence accumulation model, can provide a unified account of the interaction between hints and memory-based expectations. It provides a mechanism that can determine when image properties can lead to long-lasting attentional expectation effects on perception, and when their effect will be short lived. Such a framework can be used to explain perceptual effects in other tasks and designs, for example in visual search or bi-stable perception. Finally, it provides more depth to the predictive coding model, in which different flavours of predictions may rely on different cognitive and neural mechanisms, and lead to different perceptual and behavioural consequence.

Acknowledgment

This study was made possible through the support of a grant from the John Templeton Foundation. The opinions expressed in this publication are those of the author(s) and do not necessarily reflect the views of the John Templeton Foundation.

References:

Alba, J. W., & Hasher, L. (1983). Is memory schematic? *Psychological Bulletin*, 93(2), 203–231.

<https://doi.org/10.1037/0033-2909.93.2.203>

Baayen, R. H., Davidson, D. J., & Bates, D. M. (2008). Mixed-effects modeling with crossed random effects for subjects and items. *Journal of Memory and Language*, 59(4), 390–412.

<https://doi.org/10.1016/j.jml.2007.12.005>

Carlisle, N. B., Arita, J. T., Pardo, D., & Woodman, G. F. (2011). Attentional Templates in Visual Working Memory. *Journal of Neuroscience*, 31(25), 9315–9322.

<https://doi.org/10.1523/JNEUROSCI.1097-11.2011>

Chang, R., Baria, A. T., Flounders, M. W., & He, B. J. (2016). Unconsciously elicited perceptual prior.

Neuroscience of Consciousness, 2016(1), niw008. <https://doi.org/10.1093/nc/niw008>

Chen, W. J., & Krajbich, I. (2017). Computational modeling of epiphany learning. *Proceedings of the National Academy of Sciences*, 114(18), 4637–4642. <https://doi.org/10.1073/pnas.1618161114>

De Lange, F. P., Heilbron, M., & Kok, P. (2018). How Do Expectations Shape Perception? Perceptual Consequences of Expectation. *Trends in Cognitive Sciences*, xx, 1–16.

<https://doi.org/10.1016/j.tics.2018.06.002>

Dijkstra, N., Zeidman, P., Ondobaka, S., Van Gerven, M. A. J., & Friston, K. (2017). Distinct Top-down and Bottom-up Brain Connectivity during Visual Perception and Imagery. *Scientific Reports*, 7(1), 1–9. <https://doi.org/10.1038/s41598-017-05888-8>

Dijkstra, Nadine, Bosch, S. E., & van Gerven, M. A. J. (2017). Vividness of visual imagery depends on

- the neural overlap with perception in visual areas. *Journal of Neuroscience*, 37(5), 1367–1373.
<https://doi.org/10.1523/JNEUROSCI.3022-16.2016>
- Dolan, R. J., Fink, G. R., Rolls, E., Booth, M., Holmes, A., Frackowiak, R. S. J., & Friston, K. J. (1997).
 How the brain learns to see objects and faces in an impoverished context. *Nature*, 389(6651),
 596–599. <https://doi.org/10.1038/39309>
- Eimer, M. (2014). The neural basis of attentional control in visual search. *Trends in Cognitive
 Sciences*, 18(10), 526–535. <https://doi.org/10.1016/j.tics.2014.05.005>
- Fleming, S. M., Weil, R. S., Nagy, Z., Dolan, R. J., & Rees, G. (2010). Relating introspective accuracy to
 individual differences in brain structure. *Science (New York, N.Y.)*, 329(5998), 1541–1543.
<https://doi.org/10.1126/science.1191883>
- Friston, K. (2010). The free-energy principle : a unified brain theory ? *Nature Reviews. Neuroscience*,
 11(2), 127–138. <https://doi.org/10.1038/nrn2787>
- Fulford, J., Milton, F., Salas, D., Smith, A., Simler, A., Winlove, C., & Zeman, A. (2018). The neural
 correlates of visual imagery vividness – An fMRI study and literature review. *Cortex*, 105, 26–
 40. <https://doi.org/10.1016/j.cortex.2017.09.014>
- Gilbert, C. D., & Li, W. (2013). Top-down influences on visual processing. *Nature Reviews
 Neuroscience*, 14(5), 350–363. <https://doi.org/10.1038/nrn3476>
- Godbole, N. R., Delaney, P. F., & Verkoijen, P. P. J. L. (2014). The spacing effect in immediate and
 delayed free recall. *Memory*, 22(5), 462–469. <https://doi.org/10.1080/09658211.2013.798416>
- Greve, A., Cooper, E., Tibon, R., & Henson, R. N. (2019). Knowledge is power: Prior knowledge aids
 memory for both congruent and incongruent events, but in different ways. *Journal of
 Experimental Psychology: General*, 148(2), 325–341. <https://doi.org/10.1037/xge0000498>
- Hegde, J., & Kersten, D. (2010). A Link between Visual Disambiguation and Visual Memory. *Journal of*

- Neuroscience*, 30(45), 15124–15133. <https://doi.org/10.1523/JNEUROSCI.4415-09.2010>
- Hertz, U., & Amedi, A. (2014). Flexibility and Stability in Sensory Processing Revealed Using Visual-to-Auditory Sensory Substitution. *Cerebral Cortex*. <https://doi.org/10.1093/cercor/bhu010>
- Hsieh, P.-J., Vul, E., & Kanwisher, N. (2010). Recognition Alters the Spatial Pattern of fMRI Activation in Early Retinotopic Cortex. *Journal of Neurophysiology*, 103(3), 1501–1507. <https://doi.org/10.1152/jn.00812.2009>
- Huang, Y., & Rao, R. P. N. N. (2011). Predictive coding. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2(5), 580–593. <https://doi.org/10.1002/wcs.142>
- Judd, C. M., Westfall, J., & Kenny, D. A. (2017). Experiments with More Than One Random Factor: Designs, Analytic Models, and Statistical Power. *Annual Review of Psychology*, 68(1), 601–625. <https://doi.org/10.1146/annurev-psych-122414-033702>
- Lieberman, A., Fischer, J., & Whitney, D. (2014). Serial dependence in the perception of faces. *Current Biology*, 24(21), 2569–2574. <https://doi.org/10.1016/j.cub.2014.09.025>
- Liu, Z., Grady, C., & Moscovitch, M. (2018). The effect of prior knowledge on post-encoding brain connectivity and its relation to subsequent memory. *NeuroImage*, 167(November 2017), 211–223. <https://doi.org/10.1016/j.neuroimage.2017.11.032>
- Moore, C., & Cavanagh, P. (1998). Recovery of 3D volume from 2-tone images of novel objects. *Cognition*, 67(1–2), 45–71. [https://doi.org/10.1016/S0010-0277\(98\)00014-6](https://doi.org/10.1016/S0010-0277(98)00014-6)
- Nako, R., Smith, T. J., & Eimer, M. (2015). Activation of New Attentional Templates for Real-world Objects in Visual Search. *Journal of Cognitive Neuroscience*, 27(5), 902–912. https://doi.org/10.1162/jocn_a_00747
- Neri, P. (2014). Semantic Control of Feature Extraction from Natural Scenes. *Journal of Neuroscience*, 34(6), 2374–2388. <https://doi.org/10.1523/JNEUROSCI.1755-13.2014>

- Ozubko, J. D., Moscovitch, M., & Winocur, G. (2017). The influence of recollection and familiarity in the formation and updating of associative representations. *Learning & Memory*, 24(7), 298–309. <https://doi.org/10.1101/lm.045005.117>
- Pearson, J., & Brascamp, J. (2008). Sensory memory for ambiguous vision. *Trends in Cognitive Sciences*, 12(9), 334–341. <https://doi.org/10.1016/j.tics.2008.05.006>
- Pearson, J., Rademaker, R. L., & Tong, F. (2011). Evaluating the mind’s eye: The metacognition of visual imagery. *Psychological Science*, 22(12), 1535–1542. <https://doi.org/10.1177/0956797611417134>
- Ratcliff, R., & McKoon, G. (2008). The Diffusion Decision Model: Theory and Data for Two-Choice Decision Tasks. *Neural Computation*, 20(4), 873–922. <https://doi.org/10.1162/neco.2008.12-06-420>
- Samaha, J., Boutonnet, B., Postle, B. R., & Lupyan, G. (2018). Effects of meaningfulness on perception: Alpha-band oscillations carry perceptual expectations and influence early visual responses. *Scientific Reports*, 8(1), 1–14. <https://doi.org/10.1038/s41598-018-25093-5>
- Summerfield, C., & De Lange, F. P. (2014). Expectation in perceptual decision making: Neural and computational mechanisms. *Nature Reviews Neuroscience*, 15(11), 745–756. <https://doi.org/10.1038/nrn3838>
- Summerfield, C., & Egner, T. (2009). Expectation (and attention) in visual cognition. *Trends in Cognitive Sciences*, 13(9), 403–409. <https://doi.org/10.1016/j.tics.2009.06.003>
- Summerfield, C., Egner, T., Greene, M., Koechlin, E., Mangels, J., & Hirsch, J. (2006). Predictive Codes for Forthcoming Perception in the Frontal Cortex. *Science*, 314(5803), 1311–1314. <https://doi.org/10.1126/science.1132028>
- Sun, S. Z., Shen, J., Shaw, M., Cant, J. S., & Ferber, S. (2015). Automatic capture of attention by conceptually generated working memory templates. *Attention, Perception, and Psychophysics*,

77(6), 1841–1847. <https://doi.org/10.3758/s13414-015-0918-1>

Teufel, C., Dakin, S. C., & Fletcher, P. C. (2018). Prior object-knowledge sharpens properties of early visual 1 feature-detectors. *Scientific Reports*, (September 2017), 1–12.

<https://doi.org/10.1038/s41598-018-28845-5>

Tse, D., Langston, R. F., Kakeyama, M., Bethus, I., Spooner, P. A., Wood, E. R., ... Morris, R. G. M. (2007). Schemas and Memory Consolidation. *Science*, 316(5821), 76–82.

<https://doi.org/10.1126/science.1135935>

Van Essen, D. C., & Maunsell, J. H. R. (1983). Hierarchical organization and functional streams in the visual cortex. *Trends in Neurosciences*, 6, 370–375. [https://doi.org/10.1016/0166-](https://doi.org/10.1016/0166-2236(83)90167-4)

[2236\(83\)90167-4](https://doi.org/10.1016/0166-2236(83)90167-4)

van Kesteren, M. T. R., Fernandez, G., Norris, D. G., & Hermans, E. J. (2010). Persistent schema-dependent hippocampal-neocortical connectivity during memory encoding and postencoding rest in humans. *Proceedings of the National Academy of Sciences*, 107(16), 7550–7555.

<https://doi.org/10.1073/pnas.0914892107>

van Kesteren, Marlieke T.R., Ruiter, D. J., Fernández, G., & Henson, R. N. (2012). How schema and novelty augment memory formation. *Trends in Neurosciences*, 35(4), 211–219.

<https://doi.org/10.1016/j.tins.2012.02.001>

Zokaei, N., Board, A. G., Manohar, S. G., & Nobre, A. C. (2019). Modulation of the pupillary response by the content of visual working memory. *Proceedings of the National Academy of Sciences*,

201909959. <https://doi.org/10.1073/pnas.1909959116>