VISUAL COGNITION, 2015
VOL. 23, NO. 6, 758–780
http://dx.doi.org/10.1080/13506285.2015.1085477

Routledge
Taylor & Francis Group

Dissociation of sensory facilitation and decision bias in statistical learning

Frédéric Simard and Christopher C. Pack

Montreal Neurological Institute, McGill University, Montreal, Canada

ABSTRACT
The detection of regularities in the sensory environment, known as statistical learning, is an important brain function that has been observed in many experimental contexts. In these experiments, statistical learning of patterned sensory stimulation leads to improvements in the speed and/or accuracy with which subsequent stimuli are recognized. That is, statistical learning facilitates the transformation of sensory stimuli into motor responses, but the mechanism by which this occurs is unclear. Statistical learning could improve the efficiency of sensory processing, or it could bias responses toward particular outcomes. The distinction is important, as these different hypotheses imply different functions and different neural substrates for statistical learning. Here we address this problem by studying statistical learning as a decision-making process, which allows us to leverage the extensive computational literature on this topic. Specifically we describe a method for applying the Diffusion Decision Model (DDM) to isolate different sensory and cognitive processes associated with decision-making. The results indicate that statistical learning improves performance on a visual learning task in two distinct ways: by altering the efficiency of sensory processing and by introducing biases in the decision-making process. By fitting the parameters of the DDM to data from individual subjects, we find that the prominence of these two factors differed substantially across the population, and that these differences were predictive of individual performance on the psychophysical task. Overall, these results indicate that different cognitive processes can be recruited by statistical learning, and that the DDM is a powerful framework for detecting these influences.

ARTICLE HISTORY Received 6 February 2015; Accepted 18 August 2015

KEYWORDS Statistical learning; Diffusion Decision Model; sensory processing; prospective Memory

Introduction

Outside of a laboratory setting, individual sensory stimuli do not occur at random but rather are predictable from the occurrence of other sensory stimuli. The brain has the ability to identify and assimilate such patterns by

CONTACT Christopher C. Pack christopher.pack@mcgill.ca

learning the underlying statistical structure. Such plasticity is called statistical learning.

Experimentally, statistical learning is typically assessed through tasks in which subjects make decisions about stimuli that are to varying degrees predicted by the presence of other stimuli. By learning these associations, subjects have been shown to manifest improvements in accuracy and reaction times on a number of tasks, including natural and artificial language acquisition (Pelucchi, Hay, & Saffran, 2009; Romberg & Saffran, 2010; Saffran, Aslin, & Newport, 1996; Saffran, Newport, Aslin, Tunick, & Barrueco, 1997), spatio-temporal visual familiarization (Bulf, Johnson, & Valenza, 2011; Fiser & Aslin, 2001, 2002, 2005; Howard, Howard, Dennis, & Kelly, 2008; Kim, Seitz, Feenstra, & Shams, 2009; Kirkham, Slemmer, & Johnson, 2002; Turk-Browne, Jungé, & Scholl, 2005; Turk-Browne, Isola, Scholl, & Treat, 2008), auditory familiarization (Creel, Newport, & Aslin, 2004; Gebhart, Newport, & Aslin, 2009; Newport & Aslin, 2004) serial reaction time task (Hunt & Aslin, 2001), visual search (Chun & Jiang, 1998), and categorization (Turk-Browne, Scholl, Johnson, & Chun, 2010).

Despite the widespread observation of statistical learning, the underlying mechanisms are a subject of debate. One hypothesis suggests that statistical learning primarily affects sensory processing, either by guiding the allocation of attention or by priming specific sensory templates (Altamura, Carver, Elvevåg, Weinberger, & Coppola, 2014; Chun, 2000; Chun and Jiang, 1998; Gavornik & Bear, 2014; Meyer & Olson 2011, Meyer, Ramachandran, & Olson, 2014; Olson, Chun, & Allison, 2001; Roser, Fiser, Aslin, & Gazzaniga, 2011; Summerfield, Trittschuh, Monti, Mesulam, Egner, 2008, Summerfield & Egner, 2009; Turk-Browne, Scholl, Chun, & Johnson, 2009, Turk-Browne et al., 2010; Zhao, Al-Aidroos, & Turk-Browne, 2013). An alternative explanation is that statistical learning creates a decision bias, in the form of a commitment to a particular outcome. This idea is supported by experiments in which participants are cued (Forstmann, Brown, Dutilh, Neumann, & Wagenmakers, 2010; Liston & Stone, 2008) or given specific instructions (Summerfield et al., 2008; Sterzer, Kleinschmidt, & Rees, 2009), or when the experimenter manipulates reward or stimulus probability (Chalk, Seitz, & Seriès, 2010; Voss, Rothermund, & Voss, 2004). The two hypotheses are not mutually exclusive, and it is possible that different tasks engage different mechanisms of statistical learning.

The interaction of sensory processing and decision biases features prominently in computational models of decision-making, which have long been fundamental to our understanding of the behaviour of subjects in many psychophysical paradigms (Gold & Shadlen, 2007). These models might therefore provide insight into the mechanistic basis of statistical learning. We have therefore studied the effects of statistical learning on behaviour, using an established model of decision-making: the Diffusion Decision Model (DDM;

Ratcliff, 1978). The DDM permits inferences about the contributions of sensory and decision-making mechanisms from the accuracy and reaction time data that is typically collected during statistical learning experiments.

In this paper, we begin by reviewing the DDM and describing how it can be applied to the study of statistical learning. We then describe the results of a psychophysical task that we used to probe statistical learning in human subjects. We then use the DDM to infer the processes responsible for statistical learning, and show that this approach can account for the performance of individual subjects and for variation across different subjects. These results suggest that statistical learning uses mechanisms associated with both sensory processing and decision-making; that these mechanisms are recruited differently in different subjects; and that such differential influences can account for a wide range of behaviour. Overall these results suggest that the DDM is a powerful mechanism for interpreting data on statistical learning.

Methods

To our knowledge the DDM has not been used previously in the context of statistical learning, so we include here a brief introduction to the approach. We then describe our experimental paradigm and the methods for fitting the resulting data to the parameters of the DDM.

Theoretical basis of decision-making

Decision-making is frequently modelled as a serial process, comprised of several stages (Gold & Shadlen, 2007). The first stage (leftmost box, Figure 1) is primarily sensory: its goal is to infer the state of the environment from noisy stimuli. In this context, the quality or reliability of sensory data are represented by continuous variables. A second stage (central cloud, Figure 1)

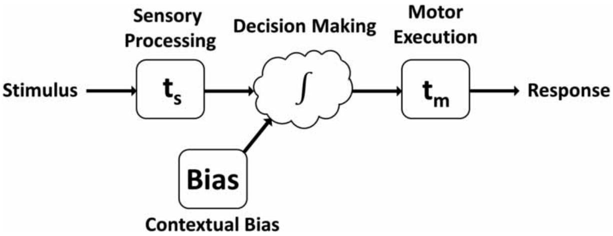

Figure 1. The stages of the decision-making process. A stimulus requires some time, t_s, to be processed. The accumulator integrates sensory information until a threshold, or decision boundary, is reached. Contextual information can bias the accumulator by adding an offset toward one response and away from the other. When the decision boundary is reached, the system commits to a specific motor action. The motor command requires some time, t_m, to be executed.

involves a decision variable that integrates the sensory information, which accumulates throughout the deliberation period (Platt & Glimcher, 1999). Contextual information about stimuli or rewards can bias the decision-making process at this stage. A decision rule stops the process once the accumulated value of the decision variable reaches a certain point. Finally, the outcome of the decision process is manifested through a motor action (rightmost box, Figure 1). This model is generally studied in the context of a standard two-alternative forced choice (2AFC) paradigm, in which the subject chooses one of two responses on each trial.

According to this model, the 2AFC decision process revolves around the dynamics of the accumulator (centre cloud in Figure 1). Several examples of ways in which a change in the properties of the accumulator can influence decisions are shown in Figure 2. Here the accumulated value of the decision variable through time is shown as the solid black lines in each panel. The horizontal grey lines correspond to the decision boundaries. Time unfolds from left to right in each plot. In this context, changes in decision-making are hypothesized to result from one or more of the following factors:

(1) Change in the drift rate

The drift rate is related to the quantity or quality of sensory information. It represents the average rate at which the decision variable travels, or drifts,

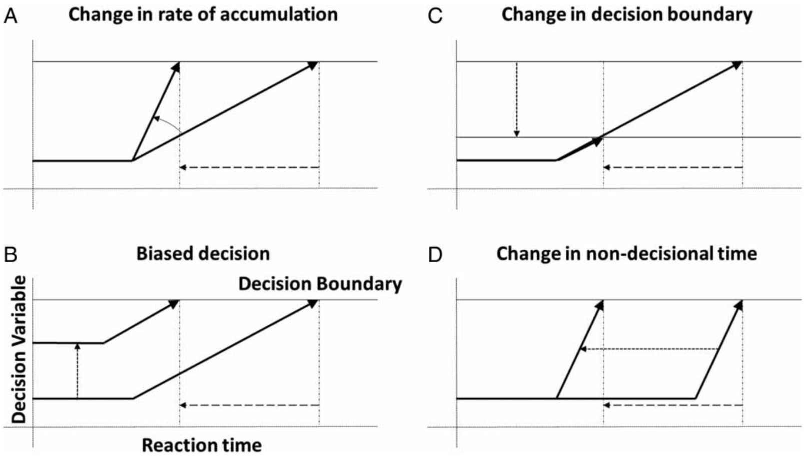

Figure 2. Effect of the manipulation of individual parameters on decision-making in the DDM. (A) The rate at which information is integrated is affected by the quantity and quality of the information. (B) The initial value of the decision variable can be biased toward one answer by contextual information. (C) The location of the decision boundary reflects the conservatism of the subject; a closer boundary means that decisions will be based on less information. (D) The non-decisional time measure the time required for sensory processing and motor execution and can change due to the complexity of the motor sequence to be executed.

toward the decision boundary (Figure 2A; Ratcliff, 1978; Voss et al., 2004). Thus any experimental manipulation that affects sensory stimuli could affect the drift rate and hence accelerate or delay the time required for the decision variable to reach the decision boundary. When the rate of accumulation is increased, decisions are generally taken faster and their accuracy is also improved (Ratcliff, 1978; Voss et al., 2004).

(2) Bias of the decision variable

A second aspect of the decision-making process is a bias in the initial value of the decision variable (Figure 2B). Under normal conditions, prior to sensory stimulation, each possible outcome of the decision process is equally probable, and no decision is favoured. However, biases can arise when one choice is associated with a higher reward (Voss et al., 2004) or when one choice is more probable (Chalk et al., 2010). This effect of a bias in the decision variable has the particular characteristic of decreasing the reaction time for the favoured outcome, while increasing the reaction time for the disfavoured outcome, resulting in an asymmetric reaction time distribution (Ratcliff & Rouder, 1998).

(3) Change in the decision boundary

The rule that defines the moment at which the value of the decision variable is committed into an action is typically implemented by the decision boundary. The decision boundary is a threshold that, once crossed, sets the end of deliberation period. A change in the decision boundary (Figure 2C) will result in a change in the reaction time and in the accuracy. If the decision boundary is closer to the starting point, decisions will be taken faster, but less time will be allowed for the accumulation of evidence. This, in turn, decreases the resilience to noise in the system and increases the error rate (Ratcliff & Rouder, 1998).

(4) Effects on non-decisional time

In addition to the dynamics of the accumulator, there are two other important aspects of decision-making. The first is the time required for the information to reach the accumulator, and second is the time required for the execution of the motor command ($t_s + t_m$ in Figure 1). The time for the information to reach the accumulator and the delay incurred by the execution of the motor command are together referred to as the non-decisional time. A change in the non-decisional time affects reaction times (Figure 2D), but has no effect on accuracy (Voss et al., 2004).

The Diffusion Decision Model

The DDM is a specific instance of the framework shown in Figures 1 and 2. The key assumption is that the observer represents a decision variable that is driven by sensory evidence toward two boundaries, each one corresponding to one of the two possible outcomes. At the beginning of a trial, the value of the decision variable is set to be somewhere in between the two boundaries, *a* and *0*. This is the starting point, *z*. Once the stimulus is made available, information is integrated and moves the value of the decision variable until it reaches one of the two boundaries. Once a boundary has been crossed, the decision is made. The average speed at which the decision value varies is the drift rate, *v*, and it depends on the quality of the information. The variable t_0 is the non-decisional time. A schematic of the model dynamics is shown in Figure 3, and the relevant parameters are listed in Table 1.

Implementation of the DDM usually involves estimation of the means and inter-trial variances of the drift rate, decision bias and non-decisional time parameters for each subject, along with the mean of the decision boundary, for a total of seven parameters. The estimation of the mean values aims at extracting the average behaviour, while the inter-trial variance aims at taking into account the variance in behaviour across trials. Additionally, the DDM has

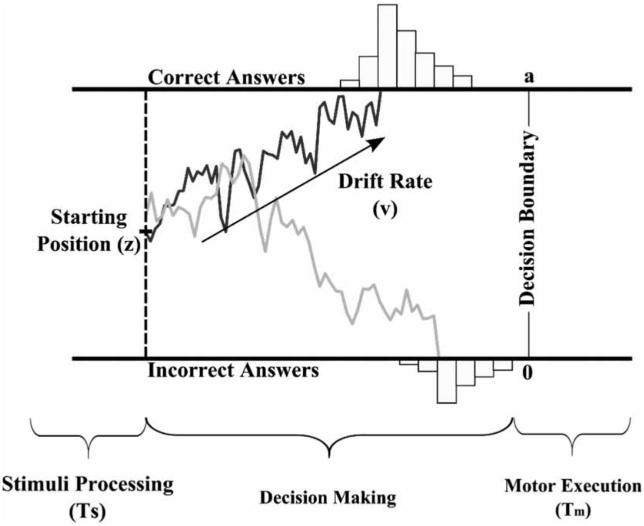

Figure 3. The Diffusion Decision Model proposes that the decision-making process behaves as a random walk process. At the beginning of each trial, the value of the decision variable is sampled from the uniform distribution, centred on the starting position, $U(z - s_z, z + s_z)$. The decision variable then drifts, following a Brownian motion, mean rate of v, sampled from a normal distribution N(v, s_v). This continues until one of the two boundaries {a, 0} is reached. Each boundary is associated with one of the choices of the two-alternate-forced-choice paradigm {top, correct; bottom, incorrect}, and the reaction time for that trial is determined by summing the decision time and the non-decision time, which sums up sensory (t_s) and execution latencies (t_m). It is the inherent noise in the system that explains errors in decision-making, but a bias in the starting position, change in the drift rate or decision boundaries all affect the system's resilience to noise.

Table 1. Parameters of the Diffusion Decision Model.

Parameter	Description
A	Location of the top boundary to the decision variable
V	Mean drift rate applied on the decision variable
Z	Starting point, normalized to the value of a. A value of 0.5 indicates a starting position halfway between the two boundaries and hence a lack of bias
t_0	Non-decisional response time
s_v	Inter-trial drift rate variance
s_z	Inter-trial starting point variance
s_t	Inter-trial non-decisional time variance

an eighth parameter, the intra-trial drift rate variance, which acts as a scaling parameter; for simplicity we set it to 0.1 (Voss, Nagler, & Lerche, 2013).

A key motivation for this study is that the parameters in the DDM correspond to potential mechanisms by which statistical learning might influence behaviour. Previous studies have shown that statistical learning improves behaviour, but in the context of the DDM, this could be associated with different parameters. Influences on information processing efficiency would manifest as an increase of the drift rate, v, while decision biases might displace the starting point, z, in favour of the most likely outcome. As detailed below, these parameters can be estimated for individual subjects, based on data of the kind typically recorded in experiments on statistical learning. We therefore collected new data in a 2AFC experiment aimed at inducing statistical learning.

Experimental paradigm

Participants

Fifty-four subjects were recruited from McGill University classified ads (33 women and 21 men) to participate in this study. Subjects ranged in age from 18 to 63 (mean 24.15 years), and all subjects reported normal or corrected-to-normal vision. The experimental protocol was approved by the IRB of the MNI, and all subjects gave informed consent prior to starting the experiments. Subjects were given 12 CAD for their participation.

Visual stimuli and procedure

Visual stimuli consisted of images of birds and cars, taken from the internet. All images were chosen to be larger than needed, to be of square ratio, and to represent only one car or one bird in focus. Images were then downscaled to a size of 256 × 256 pixels and transformed from colour to greyscale. The spatial frequencies, contrast and luminance were matched across all images using the SHINE toolbox (Willenbockel et al., 2010). A total of 16 images were preprocessed this way, eight of them representing a car and the remaining eight representing a bird. For a given session, however, only six images of each class were needed. A subset of 12 images was then randomly selected, from the complete set, for each participant.

In pilot analysis we found that the DDM was better constrained when the subjects made a sizeable number of errors during the experiments. We therefore blurred the images by manipulating the phase coherence, using the weighted mean phase method (Dakin, Hess, Ledgeway, & Achtman, 2002). For each trial a phase coherence value was sampled from a normal distribution and then used to blur the image. During the first 120 trials of a session, the phase coherence was progressively reduced, following a linear trend, from 65 ± 2.5% to 32.5 ± 2.5%. This allowed the participants to learn the images and to adjust to the level of difficulty of the perceptual task. The average phase coherence then remained at 32.5% for the rest of the session. Example images are shown in Figure 4B.

Images were presented on a CRT monitor (Hewlett Packard, A7217A) with a refresh rate of 70 Hz and a resolution of 1280 × 1024 pixels. Subjects viewed each image binocularly at a distance of 60 cm, while head position was stabilized with a chin rest. The stimuli were displayed using the Psychtoolbox (Brainard, 1997; Kleiner et al., 2007; Pelli, 1997), and the entire experiment was programmed in Matlab (MathWorks, USA).

Each session was comprised of a series of trials, each of which consisted of the presentation of a single image followed by the subject's response. Each image was taken from the selected set composed of six unique car images and six unique bird images. The images were centred on the screen and spanned eight degrees of visual angle. At the beginning of the trial, the image was shown for 500 milliseconds, and the subject had an additional 800 milliseconds, during which the image was hidden, to identify the corresponding category (bird or car; Figure 4A). Subjects were asked to answer as accurately as possible, but within the allowed timeframe. The response was given by pressing a button on a corded numeric keypad (Targus AKP10EU, Canada). Reaction times were recorded using the performance timer available in Matlab (MathWorks, USA).

The association between the keys and the classes (bird and car) was counterbalanced across participants. If, by the end of the prescribed delay, no answer had been provided, a message was displayed on the screen informing the participant that no answer was recorded. In addition, a two second penalty was added following each non-response. Participants were informed, prior to the experiment, that if they saw the non-response message, they should slightly increase their speed to prevent it from happening again. On average, the number of trials that ended with no answer for a participant represented less than 1% of the total number of trials. When a trial was completed and an answer had been provided within the instructed time limit, an inter-trial delay of 100 ms was inserted before the onset of the next trial.

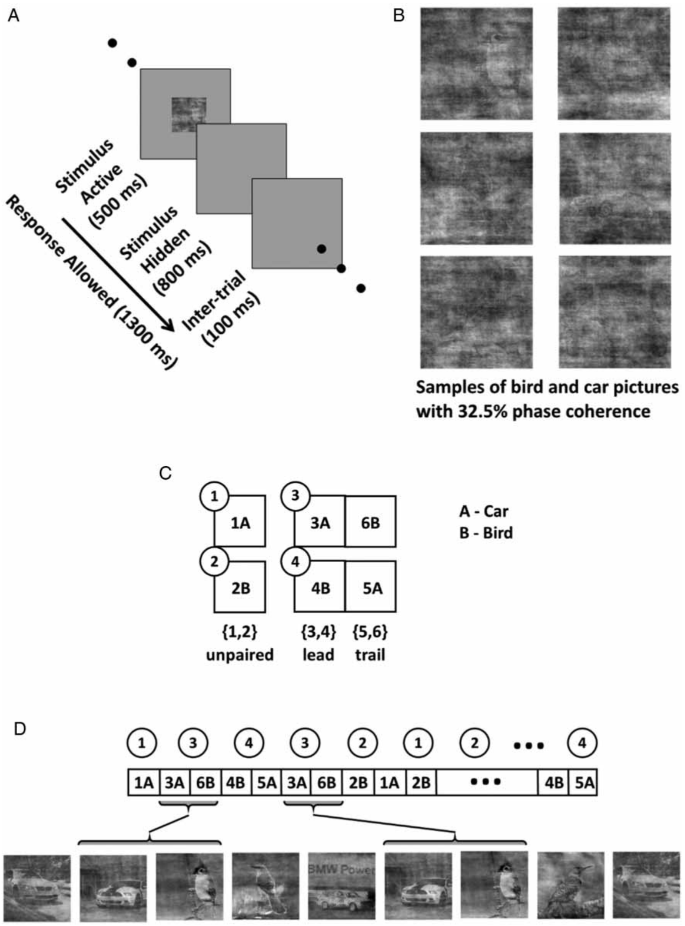

Figure 4. (A) Single trial timeline. The image is shown in the centre of the screen for 500 ms. The image is then hidden, but the subject has an additional 800 ms to indicate the category (bird or car) to which the image belongs. At the end of a trial, a 100 ms delay is added before the onset of the next trial. (B) Samples of images, with 32.5% phase coherence. It should be noted, however, that on a computer screen and in a dark environment, images are easier to distinguish than it might appear to the reader. Images on the left represent birds, and images on the right represent cars. (C) Statistical structure of the trial sequence. Unpaired images (1 and 2) are not associated with any other image; they neither predict nor are predicted. Images 3 and 4 are lead images. They predict what the next trial will be. These are followed by trail images (5 and 6), which are fully predicted by the corresponding leads. The capital letters in the boxes indicate whether the picture is a car or a bird, while the numbers index the image out of the full set of 12 images. Shown here is the statistical structure for only half of the images. The full set is merely a duplication of this summarized example. (D) A sample of the structured stream of trials is shown. It is generated by placing groups of elements 1 to 4 in a random sequence, while making sure that none is repeated twice. The numbers are then replaced by the images that correspond to them. Thus 1 becomes image 1A, 3 becomes images 3A and 6B and so on.

The sequence of images during the first 120 trials was completely random, with the limitation that an image could not be shown twice in a row. During this initial period, the experimenter stayed with the subject to provide any clarification about the task. During this same initial period, the task difficulty was progressively increased through the manipulation of the average image phase coherence. Unbeknown to the participant, the next 1200 trials were ordered according to a statistical rule.

The statistical structure of the experiment involved grouping eight of the 12 images into pairs. The two images of a pair were always presented according to the same sequence. To balance the order in which categories were presented, two of the pairs had a car preceding a bird and the other two had a bird preceding a car (Figure 4C). Because the first element of the pair always came first in the sequence, we called it the *lead,* while the second element was termed the trail. The four images remaining were called unpaired images.

As a result of this statistical structure, there were three types of trials in the experiment: four of the 12 possible trials were lead trials; four were trail trials; and four were unpaired. Accordingly, each type of trial had a different statistical property: the lead trials were predictive of the next trial; the trail trials were predicted by the previous trial, and the unpaired trials did not predict nor were predicted by any other trial. Thus the trail trials served as a probe of statistical learning, while the unpaired trials served as a baseline. The sequence of trials was generated randomly; an example sequence is shown in Figure 4D.

Each trial could conclude with one of the following outcomes: correct, if the answer provided corresponded to the category of the image that had been presented; incorrect, if the answer did not match the stimulus; and failed, if no answer was provided. Failed trials were dropped from the dataset.

After the experiment, the subjects were asked two questions, in the following order: (1) Did you notice any patterns during the experiment? If yes, elaborate. (2) What do you think this experiment was about? The first question was aimed at assessing that statistical learning had taken place implicitly. The second question was to identify subjects who might be familiar with statistical learning or who might have adjusted their strategy based on knowledge of some aspect of the experimental psychology literature. No subject reported awareness of non-random structure in the image sequences, nor did they report having employed an overt, predictive strategy.

Model-based analysis of experimental data

Data pre-processing

Before implementing the DDM-based analysis of our data, we performed a few pre-processing steps. We removed outliers from the dataset, using standard methods described elsewhere (Ratcliff & Tuerlinckx, 2002; Vandekerckhove & Tuerlinckx, 2007; Voss et al., 2013). Accordingly, we removed any trials that had a response time below a reasonable time (150 milliseconds), along with trials in which reaction times were three standard deviations above the mean. On average about 1.5% of the trials were removed this way.

Because we were interested in investigating the steady state of statistical learning, we dropped the first 400 trials (~8 minutes of exposure). This decision followed an analysis of the dynamics, which revealed that, on average, statistical learning effects were stationary after roughly 400 trials (Suppl. figure 1). The remaining trials were merged according to their type: lead, trail and unpaired. Data from car and bird trials were not different in any obvious way, and so they were combined. This left us with three sets of trials per participant, corresponding to the three types of images in our experiment.

Model analyses

The goal of our analysis was to infer the parameters that best accounted for the behavioural data from each individual subject. Specifically, we asked which parameters of the DDM best accounted for the empirical cumulative distributions of reaction times recorded during the experiments.

As mentioned above, the DDM is defined by four parameters that capture the average behaviour of the participant over all trials and three parameters that account for variation in the behaviour across trials (Voss et al., 2013). To determine the values of these parameters that best fitted a given dataset, we employed *fast-dm* (Voss et al., 2004), a programme designed for that purpose. We configured the software to rely on the chi-square criterion to measure the quality of the fit. We inferred the value of each of the parameters independently for each of the participants and each image type. Therefore, we fitted a total of 162 independent models (54 participants x 3 image types), as is commonly done in similar studies (Voss et al., 2004). We then examined the average parameter values across participants and across image types.

Results

In this section we describe the behavioural effects of statistical learning on accuracy and reaction time. We then perform a model-based analysis, using the DDM to isolate parameters that are most strongly correlated with variations in these behavioural metrics. This allows us to relate our data to specific hypotheses about statistical learning: For example, effects of learning on sensory processing could be manifested as a change in the DDM drift rate v, while effect on decision biases could be more associated with changes in the starting point z. Effects on other parameters, such as the decision threshold, would also be captured by our analysis.

Overview of psychophysical data

Figure 5A compares mean accuracy on trail and unpaired images (See Methods). Here each point corresponds to one participant, and the tendency for accuracy to be greater for the trail trials indicates the occurrence of statistical learning. This will be examined further in the next section, but the data also reveal an unexpected result: overall performance was bimodally distributed across subjects. This was particularly evident when the data were projected along the diagonal of the scatterplot (Figure 5A bimodality coefficient:

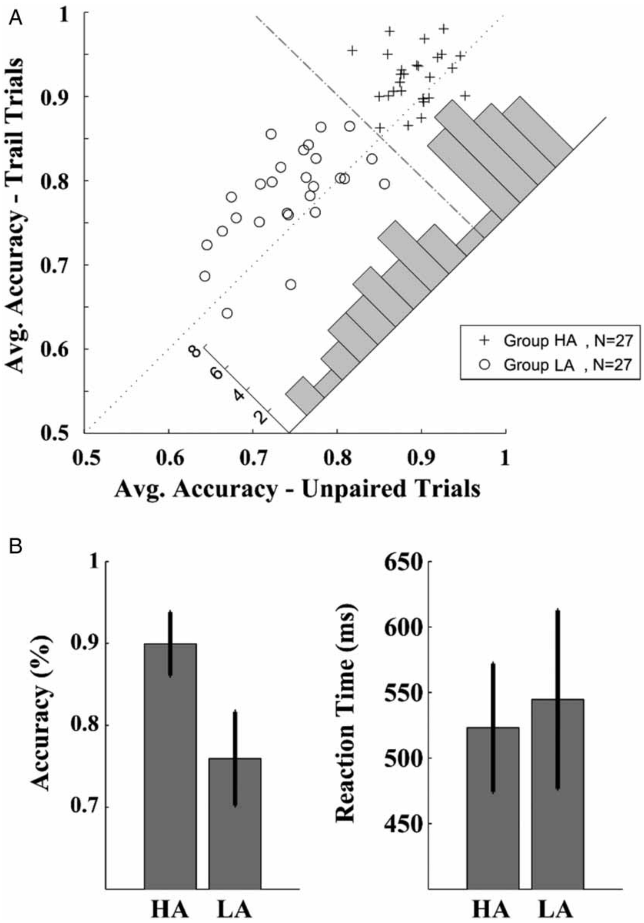

Figure 5. (A) Scatter plot showing accuracy for trail trials vs. accuracy for unpaired trials. Each point represents a single subject's average response (+s and Os indicate membership in the high- and low-accuracy groups). The inset shows a histogram of the same data projected onto an axis oriented at 45° and the boundary used as criterion to split the population into two groups. (B) Summary of accuracy (left) and reaction time (right) data for the high-accuracy (HA) and low-accuracy (LA) groups.

$b > 0.555$; SAS Institute, 1989). The bimodal distribution of performance was not related to age (*t*-test, $p = 0.61$, sup. table 1) or to gender (*z*-test, $p = 0.40$, sup. table 2), and it was largely independent of statistical learning.

To facilitate further analysis, we split the population into two groups by tracing a boundary at the location of the minimum bin count situated in between the two modes, as illustrated in Figure 5A. One group (labelled with + s in Figure 1, $N = 27$) had very high accuracy for both types of images, while the second group (labelled with Os in Figure 1, $N = 27$) had lower accuracy. Mean accuracy for the first group was 90%, while that for the low-accuracy group was 75% (two-sample *t*-test, $t(52) = 14.19$, $p < .0001$).

The variation in accuracy could reflect different response strategies used by the two groups. In particular, the high-accuracy group could have used a more conservative approach in which they sacrificed reaction time for accuracy. Contrary to this idea, the data show that the high-accuracy group actually exhibited lower reaction times for correct answers (mean = 523 vs. 545 ms, two-sample *t*-test, $t(52) = 1.\ 32$, $p < .2$; Figure 5B). These results indicate that the disparity between the two groups is not due to a different speed–accuracy trade-off.

Psychophysical measures of statistical learning

The effects of statistical learning were evident in both groups. For the high-accuracy group, the responses to the trail images were, on average, 3.2% more accurate than those to the unpaired images (paired *t*-test, $t(26) = 3.99$, $p < .001$; Table 2 and Figure 6A, left). Similarly, for the low-accuracy group, performance for the trail trials was 3.9% more accurate than for the unpaired images (paired *t*-test, $t(26) = 4.10$, $p < .001$; Table 2 and Figure 6A, right).

We also observed effects of statistical learning on reaction times. For the high-accuracy group, correct answers for the trail images were delivered 10 ms faster than those for the unpaired images (paired *t*-test, $t(26) = 3.31$, $p < .01$). Interestingly, subjects in this group typically gave incorrect answers to

Table 2. Summary of psychophysics results depicted in Figure 6.

	Accuracy	RT (correct)	RT (incorrect)
Group High Accuracy ($N = 27$)			
Unpaired	0.891 ± 0.03	0.526 ± 0.05	0.529 ± 0.05
Lead	0.884 ± 0.04	0.527 ± 0.05	0.533 ± 0.07
Trail	*0.923 ± 0.03	*0.516 ± 0.05	*0.551 ± 0.07
Group Low Accuracy ($N = 27$)			
Unpaired	0.744 ± 0.06	0.544 ± 0.06	0.551 ± 0.10
Lead	0.751 ± 0.05	0.545 ± 0.07	0.545 ± 0.09
Trail	*0.783 ± 0.06	0.543 ± 0.07	0.552 ± 0.09

*$t(26) > 2.05$, $p < .05$, two tailed paired *t*-test, $N = 27$ in both groups.
±indicates standard deviation.

trail trials more slowly than for unpaired trials (delay of 22 ms; paired *t*-test, $t(26) = 2.65$ $p < .02$; Figure 6B, left). This will be discussed further in the DDM analysis below. For the low-accuracy group, the reaction times for correct and incorrect trials did not differ significantly between trail and unpaired image types ($t(26) = 0.17$, $p > .86$ over all image pairs; Figure 6B, right). Responses to lead images did not exhibit any statistically significant trend (high-accuracy group: accuracy, $t(26) = 0.86$, $p > .4$, reaction time, $t(26) = 0.67$, $p > .5$; accuracy low-accuracy group: $t(26) = 0.66$, $p > .5$, reaction time, $t(26) = 0.45$, $p > .65$), when compared to the unpaired images for either group.

Diffusion Decision Model

We fit the parameters of the DDM to each participant's observed distributions of reaction times for correct and incorrect responses using the *fast-dm* software package (Voss & Voss, 2007). Overall this model provided an excellent

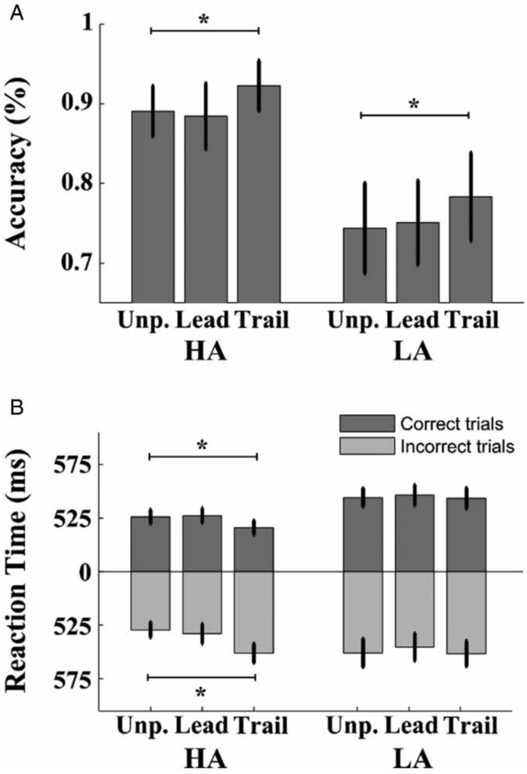

Figure 6. (A) Average accuracy values across the three types of images, with the data on the left corresponding to the HA group and the data on the right to the LA group. (B) Average reaction times for correct and incorrect trials for all image types, with the data on the left corresponding to the HA group and the data on the right to the LA group.

Table 3. Inferred parameters for the Diffusion Decision Model.

Type of trials	a	v	s_v	z	s_z	t_0	s_t
Group: High Accuracy ($N = 27$)							
Unpaired	°0.782	°3.710	1.621	0.515	0.321	0.433	0.177
Lead	0.797	3.413	1.272	0.527	0.302	0.423	0.173
Trail	*0.851	*4.111	1.835	*0.549	0.256	0.430	0.160
Group: Low Accuracy ($N = 27$)							
Unpaired	°0.708	°2.110	1.417	0.514	0.267	0.452	0.228
Lead	0.748	2.060	1.316	0.494	0.294	0.435	0.200
Trail	*0.764	*2.459	1.717	0.513	0.249	0.440	0.208

*$p < .05$, one tailed paired t-test.
°$p < .05$, two-sample t-test.

fit to the data for all of our participants ($\chi^2(200) < 55.6$, $p > .995$). Average parameter values within each accuracy group are listed in Table 3.

We first used the DDM to explore possible parametric differences between the two accuracy groups. To this end we used data from the unpaired images to separate inter-group differences from effects on statistical learning. The results (Table 3) show that the drift rate, v, and the decision boundary, a, were significantly different between groups (parameter a, paired t-test, $t(52) = 2.20$, $p < .04$; parameter v, paired t-test, $t(52) = 5.82$, $p < .0001$). The difference in the drift rate implies a difference in the ability to extract information from the noise in the images. As depicted in Figure 2C, this would lead to a faster accumulation of information, which is consistent with the faster reaction times exhibited by these subjects (Figure 5B). The higher decision boundary indicated by the DDM analysis suggests that high-accuracy subjects also waited longer to accumulate information before committing to a response. This accounts for their higher accuracy (Figure 5A). This suggests that the reaction times for this group could have been even faster, but that they sacrificed some speed for additional gains in accuracy (Figure 7A). Thus, the co-variation of drift rate and decision boundary works toward taking a more informed decision.

We next used the DDM to discern the effects of statistical learning within each group, by comparing the parameters recovered for the trail images to those for the unpaired images. Since predicted stimuli were associated with a higher accuracy, we tested the hypothesis that the drift rate had increased following statistical learning. The results show that participants in the high-accuracy group did display a significantly higher drift rate, v (Table 3, one-tailed paired t-test, $t(26) = 1.96$, $p < .03$). We also tested the possibility that the decision boundary had increased, as this would account for the increased accuracy following learning. This effect was also observed in the data (Table 3, one-tailed paired t-test, $t(26) = 1.77$, $p < .05$) In other words statistical learning led to faster accumulation of information, and a slightly higher decision threshold, as described above to explain intergroup differences. This leads to higher accuracy for trail images (Figure 6A).

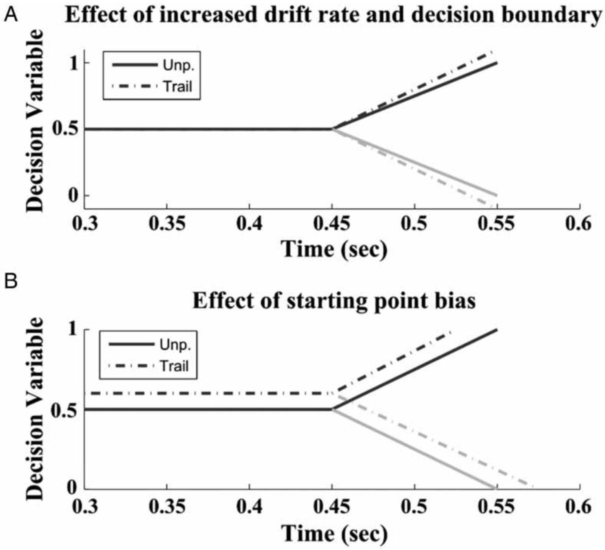

Figure 7. Graphical depiction of the effects of statistical learning on DDM parameters. Black traces show the average path of a correct decision, while the light grey traces show the path of an incorrect one. (A) The effect of an increase drift rate combined with a higher decision boundary. The main effect of this is the accumulation of more information, which improves accuracy. (B) The effect of a starting point bias toward the correct response. This results in an asymmetric reaction time for correct and incorrect answers. Correct answers are delivered much faster, but incorrect answers require more time to reach commitment.

An additional parameter of some importance in these subjects was the starting point, *z*, of the decision variable. If this variable were significantly biased toward the correct answer for the trail images relative to the unpaired images, it would account for the improved performance following statistical learning. The results of the analysis confirm this hypothesis (Table 3, one-tailed paired *t*-test, $t(26) = 2.58$, $p < .01$). This result explains a peculiar feature of the reaction times for trail images, namely that statistical learning resulted in correct answers being faster and incorrect answers being slower (Figure 6B). A divergence away from 0.5 for the starting point parameter is indicative of a decision bias and has the effect of creating an asymmetry between the reaction times for correct and incorrect answers (Figure 7B).

For the low-accuracy group, we also found a significantly greater drift rate (Table 3, one-tailed paired *t*-test, $t(26) = 1.99$, $p < .05$) and decision boundary (Table 3, one-tailed paired *t*-test, $t(26) = 2.85$, $p < .005$) for the trail images relative to the unpaired images. In contrast to the high-accuracy group, however, no change in the starting position was observed (paired *t*-test, $t(26) = 0.06$, $p > .95$). Thus the effect of statistical learning on the low-accuracy group was manifested in a higher drift rate and a higher decision boundary, but no involvement of decision bias.

Finally, we tested the possibility that a change in noise levels could account for the observed effects of statistical learning. These would be manifested as changes in the inter-trial variance parameters s_x, s_y, and s_t. However, we found that none of these parameters were significantly different across conditions or groups of subjects (Table 3; paired *t*-test, high accuracy $t(26) < 1.52$, $p > .14$; low accuracy $t(26) < 1.19$, $p > .24$). We also did not find a significant difference in non-decisional time (Table 3; paired *t*-test, high accuracy $t(26) < 1.04$, $p > .3$; low accuracy $t(26) < 1.67$, $p > .11$).

In summary, the use of the DDM has provided a mechanistic examination of two common hypotheses in the literature on statistical learning. The first, that statistical learning affects sensory processing, is supported for all subjects in our paradigm, as shown by influences on drift rates and decision boundaries. The second, a learned decision bias, is clearly evident in subjects that tend to perform the task well overall. This might be an indication that the statistical learning process can recruit two independent, but complementary, mechanisms to improve decision-making. We will detail the cognitive processes and brain regions that are likely to support these effects below.

Discussion

We have examined the ability of a popular decision-making model, the DDM, to account for psychophysical data on statistical learning. Our results suggest that, in our experimental paradigm, statistical learning improves the efficiency of sensory processing. For some subjects it also results in learned biases toward predictable outcomes. This suggests that statistical learning can engage different cognitive mechanisms and, more generally, that the DDM can be used to test quantitative hypotheses about the mechanistic basis of statistical learning.

Mechanisms supporting the effect on the drift rate

Our model-based analysis suggests that several aspects of the decision-making process are affected by statistical learning. The first factor corresponds to an increase in the drift rate, which could be a sign that the efficiency of sensory processing is improved by statistical learning. The second factor, the starting point, corresponds to the tendency of subjects to be biased toward one decision or the other. A third factor, the decision boundary, is generally associated with a change in the trade-off between speed and accuracy. We begin here by investigating the mechanisms that could be responsible for the increase in the drift rate.

The drift rate parameter of the DDM has frequently been associated with stimulus quality or reliability (Ratcliff, 1978; Voss et al., 2004). Consequently, the increase in drift rate parameter could be a sign of an improvement of sensory processing. The idea that statistical learning could be used to improve sensory processing via changes in low-level sensory processes has been suggested previously (Fiser, Berkes, Orbán, & Lengyel, 2010; Orban, Fiser, Aslin, & Lengyel, 2008; Roser et al., 2011; Turk-Browne, Simon, & Sederberg, 2012; Turk-Browne, Isola, Scholl, & Treat, 2008).

Supporting this view, many studies have reported an effect of statistical learning on neural activity in sensory brain areas. Meyer and Olson (2011; Meyer et al., 2014) presented monkeys with repeated pairings of the same visual images, such that the second stimulus in the pair was predicted by the first. Following long-term exposure to these stimulus pairings, inferotemporal neuron responses to the second stimulus were lower when it was preceded by the first stimulus than when it was presented in isolation. Similarly, recent imaging studies have identified, in early visual areas, BOLD signal correlates of statistical learning (Turk-Browne et al., 2010, 2009). Thus participants in our study might have used statistical information to filter the noisy images we presented to them and to extract meaningful information with more efficiency.

An alternative interpretation, which does not contradict the previous one, is that statistical learning could affect the distribution of attentional resources (Turk-Browne et al., 2012). Attention has already been suggested to increase the drift rate in a different context (Summerfield & Egner, 2009). Regularities in the environment are known to spontaneously capture attention (Zhao et al., 2013), and contextual cuing has been showed to guide visual attention (Chun, 2000). Thus visual attention may have played a role in our paradigm, but we have not investigated this issue systematically.

Statistical learning is also known to operate at the level of categorical knowledge (Brady & Oliva, 2008) and to affect the representation of stimuli in the Medial Temporal Lobe (Schapiro, Kustner, & Turk-Browne, 2012). Thus our use of stimuli that could be labelled at the superordinate level might have lead subjects to recruit other brain areas that are less involved in visual stimulus processing. Future investigations could be aimed at differentiating and quantifying the relative contribution of each of these factors to the change in drift rate.

Mechanisms supporting the effect on the decision bias

Effects on starting position have been documented in tasks in which the participants are explicitly biased toward favouring one answer over the other. This can be done by giving them specific instructions (Summerfield et al., 2008; Sterzer et al., 2009), cuing them (Forstmann et al., 2010; Liston & Stone, 2009) or varying stimulus or reward probability (Voss et al., 2004; Chalk et al., 2010). These experiments generally call for the involvement of declarative memory: an explicit and conscious process. In contrast, statistical learning is an implicit process and there is a debate on whether it can make use of the mechanisms associated with explicit processes (Reber, 2013). Recent imaging studies, however, have revealed an overlap between brain regions that are activated during statistical learning tasks and those that are activated during explicit memory processing; relevant brain areas include the medial temporal lobe, hippocampus, striatum and caudate nucleus (Schapiro et al., 2012, Schapiro, Gregory, Landau, McCloskey, & Turk-Browne, 2014; Turk-Browne et al., 2009). The activation of these brain regions is triggered in anticipation of the predicted item, rather than in response to it, hence the prospective nature of this memory recall. One neural manifestation of this memory function was found in the lateral intraparietal area, where the neural activity that preceded stimulus onset was correlated with a bias of the starting position in a decision-making task (Coe, Tomihara, Matsuzawa, & Hikosaka, 2002). The use of an implicit prospective memory recall is, thus, likely to support the decision bias we observed in our experiment.

Mechanisms supporting the effect on the decision boundary

The effect on the decision boundary is likely to be a corollary of the effect on the drift rate, as it has been observed previously that speed and accuracy trade off in various tasks (Ratcliff, 1978). Had the drift rate increased, but the decision boundary remained the same, answers would have been provided faster, but with only a marginal improvement of accuracy. Given that our instructions specified that participants should emphasize accuracy, it is likely that subjects sacrificed the potential gain in reaction time for increased accuracy, a strategy mediated by an increase of the decision boundary parameter (Ratcliff, 1978).

Dissociation of information processing and implicit prospective memory

Statistical learning could affect information processing and memory through a common mechanism or through independent mechanisms. Our DDM analysis for the group with higher accuracy revealed an effect on both the drift rate and decision bias. For the group with lower accuracy, however, there were no effects of statistical learning on reaction times, nor was there an effect on the starting point parameter. The latter is typically associated with memory, as it involves an experience-dependent change in strategy, in advance of the presentation of a stimulus.

Together, these results show that the effects on sensory processing and memory can be dissociated. This in turn supports the view that statistical learning is implemented by a combination of mechanisms, rather than being a central process with distributed effects. Our observation that sensory processing is affected in the absence of memory related processes is an indication that, in our paradigm, the improvement of sensory processing does not involve the establishment of memory traces. This interpretation would be in accordance with the hypothesis that lower stages of sensory processing are responsible for organizing and amplifying statistically relevant information before it can be processed by higher stages (Barakat, Seitz, & Shams, 2013; Conway & Christiansen, 2006; Fiser et al., 2010; Krause & Pack, 2014; Roser et al., 2011; Tsui, Hunter, Born, & Pack, 2010; Turk-Browne et al., 2008).

Acknowledgements

This work was supported by the Natural Sciences and Engineering Research Council Grant [grant number 341534-12] to C.C.P. and a scholarship from the Fonds de recherche nature et technologies du Québec to F.S.

Disclosure statement

No potential conflict of interest was reported by the authors.

Supplementary material

Supplementary Figures/Tables are available via the "Supplementary" tab on the article's online page (http://10.1080/13506285.2015.1085477.YEAR.2015).

References

Altamura, M., Carver, F. W., Elvevåg, B., Weinberger, D. R., & Coppola, R. (2014). Dynamic cortical involvement in implicit anticipation during statistical learning. *Neuroscience Letters*, *558*, 73–77.

Barakat, B. K., Seitz, A. R., & Shams, L. (2013). The effect of statistical learning on internal stimulus representations: Predictable items are enhanced even when not predicted. *Cognition*, *129*(2), 205–211.

Brady, T. F., & Oliva, A. (2008). Statistical learning using real-world scenes extracting categorical regularities without conscious intent. *Psychological Science*, *19*(7), 678–685.

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433–436.

Bulf, H., Johnson, S. P., & Valenza, E. (2011). Visual statistical learning in the newborn infant. *Cognition*, *121*(1), 127–132.

Chalk, M., Seitz, A. R., & Seriès, P. (2010). Rapidly learned stimulus expectations alter perception of motion. *Journal of Vision*, *10*(8), 2, 1–18.

Chun, M. M. (2000). Contextual cueing of visual attention. *Trends in Cognitive Sciences*, *4*(5), 170–178.

Chun, M. M., & Jiang, Y. (1998). Contextual cueing: Implicit learning and memory of visual context guides spatial attention. *Cognitive Psychology*, *36*(1), 28–71.

Coe, B., Tomihara, K., Matsuzawa, M., & Hikosaka, O. (2002). Visual and anticipatory bias in three cortical eye fields of the monkey during an adaptive decision-making task. *The Journal of Neuroscience*, *22*(12), 5081–5090.

Conway, C. M., & Christiansen, M. H. (2006). Statistical learning within and between modalities pitting abstract against stimulus-specific representations. *Psychological Science*, *17*(10), 905–912.

Creel, S. C., Newport, E. L., & Aslin, R. N. (2004). Distant melodies: statistical learning of nonadjacent dependencies in tone sequences. *Journal of Experimental Psychology: Learning, memory, and cognition*, *30*(5), 1119–1130.

Dakin, S. C., Hess, R. F., Ledgeway, T., & Achtman, R. L. (2002). What causes non-monotonic tuning of fMRI response to noisy images? *Current Biology*, *12*(14), R476–R477.

Fiser, J., & Aslin, R. N. (2001). Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychological Science*, *12*(6), 499–504.

Fiser, J., & Aslin, R. N. (2002). Statistical learning of higher-order temporal structure from visual shape sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*(3), 458–467.

Fiser, J., & Aslin, R. N. (2005). Encoding multielement scenes: Statistical learning of visual feature hierarchies. *Journal of Experimental Psychology: General*, *134*(4), 521–537.

Fiser, J., Berkes, P., Orbán, G., & Lengyel, M. (2010). Statistically optimal perception and learning: From behavior to neural representations. *Trends in Cognitive Sciences*, *14*(3), 119–130.

Forstmann, B. U., Brown, S., Dutilh, G., Neumann, J., & Wagenmakers, E. J. (2010). The neural substrate of prior information in perceptual decision making: A model-based analysis. *Frontiers in Human Neuroscience*, *4*(40), 1–12.

Gavornik, J. P., & Bear, M. F. (2014). Learned spatiotemporal sequence recognition and prediction in primary visual cortex. *Nature Neuroscience*, *17*(5), 732–737.

Gebhart, A. L., Newport, E. L., & Aslin, R. N. (2009). Statistical learning of adjacent and nonadjacent dependencies among nonlinguistic sounds. *Psychonomic Bulletin & Review*, *16*(3), 486–490.

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, *30*, 535–574.

Howard Jr, J. H., Howard, D. V., Dennis, N. A., & Kelly, A. J. (2008). Implicit learning of predictive relationships in three-element visual sequences by young and old adults. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(5), 1139–1157.

Hunt, R. H., & Aslin, R. N. (2001). Statistical learning in a serial reaction time task: Access to separable statistical cues by individual learners. *Journal of Experimental Psychology: General*, *130*(4), 658–680.

Kim, R., Seitz, A., Feenstra, H., & Shams, L. (2009). Testing assumptions of statistical learning: Is it long-term and implicit? *Neuroscience Letters*, *461*(2), 145–149.

Kirkham, N. Z., Slemmer, J. A., & Johnson, S. P. (2002). Visual statistical learning in infancy: Evidence for a domain general learning mechanism. *Cognition*, *83*(2), B35–B42.

Kleiner, M., Brainard, D., Pelli, D., Ingling, A., Murray, R., & Broussard, C. (2007). What's new in Psychtoolbox-3. *Perception*, *36*(14), 1–1.

Krause, M. R., & Pack, C. C. (2014). Contextual modulation and stimulus selectivity in extrastriate cortex. *Vision Research*, *104*, 36–46.

Liston, D. B., & Stone, L. S. (2008). Effects of prior information and reward on oculomotor and perceptual choices. *The Journal of Neuroscience*, *28*(51), 13866–13875.

Meyer, T., & Olson, C. R. (2011). Statistical learning of visual transitions in monkey inferotemporal cortex. *Proceedings of the National Academy of Sciences*, *108*(48), 19401–19406.

Meyer, T., Ramachandran, S., & Olson, C. R. (2014). Statistical learning of serial visual transitions by neurons in monkey inferotemporal cortex. *Journal of Neuroscience*, *34*(28), 9332–9337.

Newport, E. L., & Aslin, R. N. (2004). Learning at a distance I. Statistical learning of non-adjacent dependencies. *Cognitive Psychology*, *48*(2), 127–162.

Olson, I. R., Chun, M. M., & Allison, T. (2001). Contextual guidance of attention. *Brain*, *124* (7), 1417–1425.

Orbán, G., Fiser, J., Aslin, R. N., & Lengyel, M. (2008). Bayesian learning of visual chunks by human observers. *Proceedings of the National Academy of Sciences*, *105*(7), 2745–2750.

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*(4), 437–442.

Pelucchi, B., Hay, J. F., & Saffran, J. R. (2009). Statistical learning in a natural language by 8-month-old infants. *Child Development*, *80*(3), 674–685.

Platt, M. L., & Glimcher, P. W. (1999). Neural correlates of decision variables in parietal cortex. *Nature*, *400*(6741), 233–238.

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, *85*(2), 59–108.

Ratcliff, R., & Rouder, J. N. (1998). Modeling response times for two-choice decisions. *Psychological Science*, *9*(5), 347–356.

Ratcliff, R., & Tuerlinckx, F. (2002). Estimating parameters of the diffusion model: Approaches to dealing with contaminant reaction times and parameter variability. *Psychonomic Bulletin & Review*, *9*(3), 438–481.

Reber, P. J. (2013). The neural basis of implicit learning and memory: A review of neuropsychological and neuroimaging research. *Neuropsychologia*, *51*(10), 2026–2042.

Romberg, A. R., & Saffran, J. R. (2010). Statistical learning and language acquisition. *Wiley Interdisciplinary Reviews: Cognitive Science*, *1*(6), 906–914.

Roser, M. E., Fiser, J., Aslin, R. N., & Gazzaniga, M. S. (2011). Right hemisphere dominance in visual statistical learning. *Journal of Cognitive Neuroscience*, *23*(5), 1088–1099.

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science*, *274*(5294), 1926–1928.

Saffran, J. R., Newport, E. L., Aslin, R. N., Tunick, R. A., & Barrueco, S. (1997). Incidental language learning: Listening (and learning) out of the corner of your ear. *Psychological Science*, *8*(2), 101–105.

SAS Institute, Inc. (1989). *SAS/STAT user's guide* (4th edition). Cary, NC.

Schapiro, A. C., Gregory, E., Landau, B., McCloskey, M., & Turk-Browne, N. B. (2014). The necessity of the medial temporal lobe for statistical learning. *Journal of Cognitive Neuroscience*, *26*(8), 1736–1747.

Schapiro, A. C., Kustner, L. V., & Turk-Browne, N. B. (2012). Shaping of object representations in the human medial temporal lobe based on temporal regularities. *Current Biology*, *22*(17), 1622–1627.

Sterzer, P., Kleinschmidt, A., & Rees, G. (2009). The neural bases of multistable perception. *Trends in Cognitive Sciences*, *13*(7), 310–318.

Summerfield, C., & Egner, T. (2009). Expectation (and attention) in visual cognition. *Trends in Cognitive Sciences*, *13*(9), 403–409.

Summerfield, C., Trittschuh, E. H., Monti, J. M., Mesulam, M. M., Egner, T. (2008 Sep). Neural repetition suppression reflects fulfilled perceptual expectations. *Nature Neuroscience*, *11*(9), 1004–1006.

Tsui, J. M., Hunter, J. N., Born, R. T., & Pack, C. C. (2010). The role of V1 surround suppression in MT motion integration. *Journal of Neurophysiology*, *103*(6), 3123–3138.

Turk-Browne, N. B., Isola, P. J., Scholl, B. J., & Treat, T. A. (2008). Multidimensional visual statistical learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(2), 399–407.

Turk-Browne, N. B., Jungé, J. A., & Scholl, B. J. (2005). The automaticity of visual statistical learning. *Journal of Experimental Psychology: General*, *134*(4), 552–564.

Turk-Browne, N. B., Scholl, B. J., Chun, M. M., & Johnson, M. K. (2009). Neural evidence of statistical learning: Efficient detection of visual regularities without awareness. *Journal of Cognitive Neuroscience*, *21*(10), 1934–1945.

Turk-Browne, N. B., Scholl, B. J., Johnson, M. K., & Chun, M. M. (2010). Implicit perceptual anticipation triggered by statistical learning. *Journal of Neuroscience*, *30*(33), 11177–11187.

Turk-Browne, N. B., Simon, M. G., & Sederberg, P. B. (2012). Scene representations in parahippocampal cortex depend on temporal context. *Journal of Neuroscience*, *32*(21), 7202–7207.

Vandekerckhove, J., & Tuerlinckx, F. (2007). Fitting the Ratcliff diffusion model to experimental data. *Psychonomic Bulletin & Review*, *14*(6), 1011–1026.

Voss, A., Nagler, M., & Lerche, V. (2013). Diffusion models in experimental psychology: A practical introduction. *Experimental Psychology*, *60*(6), 385–402.

Voss, A., Rothermund, K., & Voss, J. (2004). Interpreting the parameters of the diffusion model: An empirical validation. *Memory & Cognition*, *32*(7), 1206–1220.

Voss, A., & Voss, J. (2007). Fast-dm: A free program for efficient diffusion model analysis. *Behavior Research Methods*, *39*(4), 767–775.

Willenbockel, V., Sadr, J., Fiset, D., Horne, G. O., Gosselin, F., & Tanaka, J. W. (2010). Controlling low-level image properties: the SHINE toolbox. *Behavior Research Methods*, *42*(3), 671–684.

Zhao, J., Al-Aidroos, N., & Turk-Browne, N. B. (2013). Attention is spontaneously biased toward regularities. *Psychological Science*, *24*(5), 667–677.