

Active learning as a link between environmental statistics and the development of internal representations

József Arató

Central European University
Department of Cognitive Science

*In partial fulfilment of the requirements for the degree of
Doctor of Philosophy in Cognitive Science*

Supervisor: József Fiser

Secondary Supervisor: Gergely Csibra
External Advisor: Constantin Rothkopf

Submitted: Budapest, August 2018

Declaration of Authorship

I hereby declare that this submission is my own work and to the best of my knowledge it contains no materials previously published or written by another person, or which have been accepted for the award of any other degree or diploma at Central European University or any other educational institution, except where due acknowledgement is made in the form of bibliographical reference.

József Arató

Abstract

Although it is known that facing a dynamically changing sensory stream, people's perceptual decisions could be influenced not only by individual past stimuli, but also by extracted summary statistics of the stimuli, the effects of these long-term influences are underexplored. In the present thesis, I explored the impact of past stimulus statistics on two distinct types of visual decisions. In the first line of research, in Chapters 2-3, I focused on visual explorative decisions via eye-movements and investigated whether hidden statistical structures of complex scenes could influence visual exploration. I found that spatial regularities of visual stimuli influenced explorative eye-movement patterns, that this effect emerged over time, and it could predict the success in learning the underlying structure of the input. These findings suggest a strong relationship between visual exploration and learning, during which the two processes continuously influence each other. I also showed how this relationship depended on the explicit vs. implicit nature of the task. In the second line of research, in Chapters 4-5, I explored long-term statistical influences in perceptual decision making. To this end, I tested the influence of past probabilities of appearance on discrimination judgments about ambiguous stimuli. I found that statistics of past stimulus strongly influenced perceptual decisions independently of the well-documented short-term sequential effects. This past influence depended on the change-dynamics between long-term and recent stimulus probabilities, sometimes resulting in locally irrational biases. Taken together, the results in these two research domains are consistent with a framework, in which past stimulus statistics are perpetually and automatically built into complex internal representations, which in turn, depending on the task and type of regularity, can dramatically influence visual decisions.

Acknowledgements

I would like to mention some of the people who helped me over the years.

To start with, I want to thank current and former members of the Vision Lab. I am grateful to you on many levels, working with you guys has been a pleasure.

From the Lab, I owe special thanks to:

Oana Stanciu, for all the help and advice. Tünde Szabó, for the long discussions. Gábor Lengyel, Márton Nagy, Sára Jellinek, Ádám Koblinger and Benő Márkus in general, for many small and not so small things. Csilla Pakozdy, for the help with grammar.

I would also like to thank:

Ádám Koblinger (again), for the modeling work that forms an important part of the General Discussion.

Máté Lengyel, for advice on modeling.

Constantin Rothkopf, for his involvement in the project which forms Chapter 2 of the current thesis.

My secondary supervisor, Gergely Csibra, for the always straight to the point advice.

My supervisor, József Fiser, for all the help, motivation and support over the years.

Everyone at CEU Department of Cognitive Science.

My sponsors, that is: CEU in general, especially the CEU Global Teaching Fellowship and the travel grants that allowed me to visit the University of Hamburg, summer schools and conferences. The Sciex – Swiss Scientific Exchange Program which helped me to spend valuable time at the University of Fribourg.

I would also like to thank my family, friends, flatmates in Budapest and all over the world! I will not write a list, but you know who you are!

Finally, I would like to thank Lili Varga, for her unwavering support.

Table of Contents

Abstract	3
Acknowledgements	4
List of Figures and Tables	8
Chapter 1 Past Influences on Perceptual Decisions and Eye-movements	10
Introduction.....	10
Perceptual Decision-Making.....	11
Summary on past influences in perceptual decision making	26
Visual Exploration.....	27
Summary on visual exploration	36
Outlook on the goals of the thesis	38
Chapter 2 Active Statistical Learning.....	39
Summary	39
Introduction.....	40
General Methods.....	44
Stimuli and Structure:.....	44
Computational Modeling.....	49
Experiment 1: Explicit Active Statistical Learning.....	54
Introduction.....	54
Results	55
Discussion	63
Experiment 2 and 3.: Implicit Active Statistical Learning	64
Introduction.....	64
Methods	66
Results	67
Discussion	77
General Discussion	79
Conclusions.....	83
Chapter 3. The link between Statistical Learning, Eye-movements and Working Memory.....	85
Summary	85
Introduction.....	86
Experiments 4-5: Active Statistical Learning and Working Memory.....	87
Introduction.....	87
Methods	88

Results	92
Discussion	99
Experiment 6.: Eye-movements are attracted by temporal regularities	102
Methods	106
Data Analysis	108
Results	109
Discussion	112
General Discussion	114
Chapter 4 Effects of Past Probabilities on Perceptual Decision Making: Exps 1-4.....	116
Summary	116
Introduction.....	117
Methods	119
Results of Experiments 1-3.....	125
Balanced Test Sanity Check- Exp 4	129
Discussion	130
Chapter 5 Change-dependent weighting of past probabilities: Exps. 5-7.....	135
Summary	135
Introduction.....	136
Eliminating long-term effects by gradual changes: Experiment 5	137
Methods	138
Results	139
Inducing a long-term effect by a sudden change without shifting appearance probabilities	
Experiment 6	142
Methods	142
Analysis.....	143
Results	144
Scaling of change-point-related effects: Experiment 7	145
Methods	146
Results	147
Combined Prediction of RT and Choice in Exps. 1-7	148
Discussion	149
Chapter 6 General Discussion	152
Active learning as sequential decision making.....	153
Active learning as ecological decision making	154
Active statistical learning in the light of active learning theory.....	156
Modeling the learning of environmental regularities	159

Role of uncertainty in past probability effects.....	161
Sampling of episodic experience.....	163
Links to reasoning.....	164
Conclusions.....	167
Appendix A Additional Analysis for Chapters 2-3.....	169
Appendix B Additional Analysis for Chapters 4-5.....	173
Appendix Text B.1. Validation of Noise Model.....	174
Appendix Text B.2. Drift Diffusion Model.....	175
References.....	180

List of Figures and Tables

Figure 2.1. Stimuli and Procedure	44
Figure 2.2. Transition P-s for Null Model, Model Simulation	50
Table 2.1.: Numbering of cells.....	53
Figure 2.3. Test performance and transition predictability in Exp 1.....	55
Figure 2.4.: Pair structure influence on eye-movement measures in Exp 1.	59
Figure 2.5 Model based assessment of pair-structure influence on visual exploration.	60
Table 2.2: Combined prediction of Learning.....	61
Figure 2.6. Predicting the content of learning	62
Figure 2.7: Familiarity Test Performance distribution in Exp 2 & 3.	67
Figure 2.8: Entropy of Transition Distributions	70
Figure 2.9: Pair Exploration Rate	71
Figure 2.10: Pair Return Rate over time in experiments 2 & 3.	72
Figure 2.11. Model based assessment of Pair influence	74
Figure 3.1. Structure of Exp 4 & 5	89
Figure 3.2 Familiarity Test Performance Distribution Exp 4 & 5	92
Figure 3.3. Working Memory Task Performance in Exp 4 & 5	93
Figure 3.4. Pair Structure influence in Exps 4 & 5.	95
Table 3.1. Combined prediction of learning.....	96
Figure 3.5. Global Structure influence in Exp 4 & 5.....	97
Figure 3.6: Pair vs. Global structure	98
Figure 3.7. Stimuli and Structure of Exp 6.....	105
Figure 3.8. Task Performance in Exp 6	109
Figure 3.9: Statistical Influence on Eye-movements.	111
Figure 4.1. General Methods and Model Selection.....	125
Figure 4.2: Experiments 1-3	126
Figure 4.3. Immediate past and long-term probabilities	127
Figure 4.4. Contrast of Decision Bias and Response Times.....	128
Figure 4.5. Experiment 4	129
Figure 5.1. Experiment 5.	139
Figure 5.2 Reaction times and decisions are differently affected by change dynamics.	141
Figure 5.3 Experiment 6.	143
Figure 5.4: Emergence and Temporal Stability of Long-term effects.	144
Figure 5.5. Experiment 7	146
Table 5.1: Predicting Decisions Bias and RT in Exps. 1-7	148
Appendix Figure. A.1. Model fit likelihood with different null models.....	169
Appendix Figure. A.3. Fitting M1 with larger alpha range.	170
Appendix Figure A.2. Pair influence with alternative null models.	170
Appendix Figure. A.5.	171
Appendix Figure. A.4. Model Selection Results for Exp 1-3	171
Appendix Table A.1.....	172
Appendix Figure B.1	173
Appendix Table B.1.....	174
Appendix Figure B.2	175
Appendix Table B.2.....	177
Appendix Table B.3.....	177
Appendix Table B.4.....	178
Appendix Table B.5.....	178

Chapter 1

Past Influences on Perceptual Decisions and Eye-movements

Introduction

Intuitively, vision seems easy: we only need to open our eyes, look around, and we effortlessly see the world as it is. However, it has been recognized long ago, that the biological underpinnings of this smooth visual experience are very complex and there is a lot of processing occurring under the hood. We don't see the world exactly as it is, our visual system is susceptible to a large number of visual illusions, which in turn can reveal some of the internal mechanisms of visual processing (Gregory, 1970). The main reason for this deviation from faithful transmission of incoming sensory information is that our smooth visual experience is based on a process that integrates sensory information with a complex underlying internal model of the world that we spend years developing. The incorporation of this extensive knowledge into our momentary experience is what makes the usually effortless recognition of the huge variety of objects, faces, animals etc. possible.

While the influence of past visual experience on current sensation is acknowledged by most cognitive scientists, there is less agreement on the details of this influence. The dominant approach is to treat vision as perceptual inference: the visual system is using the ambiguous perceptual information to guess the most likely true state of the world (Yuille & Kersten, 2006). In this context, it is well-documented that past visual experience can influence this inference at different time-scales: from seconds to years, it can affect elaborated decisions as well as fine-tune the details of low-level perceptual mechanisms (Sagi, 2011; Thompson & Burr, 2009). Yet, both the nature of these effects on a longer time scale and the seamless integration of effects between different timescales are underexplored topics.

A second important aspect of human vision is that it is an active process (Hayhoe & Ballard, 2005). The amount of visual information that is continuously arriving to our eyes is so huge that we cannot fully process it. Therefore, a presumed set of attentional mechanisms is actively used to filter out the irrelevant aspects of the incoming stream. The foveated nature of the human eye makes the resolution at distinct parts of the visual field very different. As a consequence, where we focus our gaze has a strong influence on the amount of information we get from various parts of our surroundings. Our seemingly effortless eye-movements are in fact, continuously sampling our environment, selecting parts of our surroundings to receive more information from (Friston, Adams, Perrinet, & Breakspear, 2012). Once again, the link between this active process and the emerging internal representations is a topic with many open questions.

In this thesis, I will focus on the above two issues. The central argument of the thesis is that we build our visual experience continuously into our internal representations, which in turn influence subsequent momentary perceptual decisions. This influence has many manifestations, it shows up with or without a task, and it affects many aspects of visual processing, from visual discrimination to eye-movements. In the first part of this introductory chapter, I will review the most important findings of how past events influence vision, starting from short- and moving on to longer time-scales in perceptual decision making. In the second half, I will move on to reviewing how experience influences visual search and eye-movements. I will argue that a full treatment of visual exploration requires the integration of bottom-up and top-down information in a framework, which considers natural vision as an active sequential decision-making process.

Perceptual Decision-Making

Adaptation

Anyone familiar with the situation of entering a dark room from outdoors has experienced that past illuminations levels influence currently perceived brightness. Similarly, if after looking at a waterfall for a couple of minutes we look away, we see the world going upwards (Barlow & Hill, 1963). After looking at a masculine face, a neutral face seems to be more feminine (C. Zhao, Seriès, Hancock, &

Bednar, 2011). After staring at a left tilted grating, an intermediate grating will be seen as tilted more to the right (Wolfe, 1984). What is common to all these phenomena is, that recent visual experience can bias our perceptual system to see the opposite in the current input (right after left, female after male, up after down, dark after bright). These effects have all been loosely grouped under the label of adaptation, demonstrating the subjective nature of human visual experience in various contexts. Adaptation has also been described in other species (Vinken, Vogels, & Beeck, 2017), and it has been shown to have an influence at many different time-scales from milliseconds to days. While the dominant explanation for adaptation used to be neural fatigue, the various manifestations of the effect suggest that the term covers hugely different neural mechanisms (Thompson & Burr, 2009). One reason why similar effects can be described with such a diverse set of stimuli could be that encoding small differences in a large range is a fundamental task of a flexible visual system in all relevant dimensions. Adaptation is typically classified as a negative recency effect, since it describes a phenomenon of biasing the coding of the input to the direction opposite from the viewer's recent experience.

Priming

Another basic type of past influences on perception has been collectively described as priming (Tulving & Schacter, 1990). Unlike adaptation that describes a negative influence by recent past, priming refers to a positive effect, a facilitation of identifying stimuli that occurred recently. Like adaptation, priming is also a very general term used to describe very different phenomena (eg. social vs. conceptual priming). The findings relevant for the present thesis have been collectively described as perceptual priming (Treisman, 1992). The classic definition of perceptual priming is an automatic enhancement in processing features or locations that were recently relevant. Manifested by, for example, enhanced feature based visual search (Maljkovic & Nakayama, 1994, 1996). For the sake of completeness, it is worth noting that while priming usually refers to a facilitating effect, negative priming has also been reported. This refers to the phenomenon, in which ignoring objects can make subsequent processing of the same objects slower (Tipper, 1985) .

Serial Dependence

While priming has been used to describe faster processing of repeated stimuli or features for a long time, in recent years, there has been a surge of interest in a related finding, showing that not only is the processing speed increased, but the perceived identity of the stimulus itself became biased toward the immediate past. While adaptation is a prevalent mechanism enabling us to see changes in the sensory input efficiently, a complementary goal is to detect persistent items reliably. This can be achieved by relying on the fact that the world is more or less stable. However, it is not trivial to capitalize on this fact since our eyes are constantly moving (a few times every second), changing the input to any part of the visual cortex with each saccade. A phenomenon that could underlie the stable perception of the constantly varying visual input has been recently described as serial dependence and was demonstrated with different stimuli (gratings, faces) (Fischer & Whitney, 2014; Liberman, Fischer, & Whitney, 2014). The main finding of this work is that perceptual estimates of the orientation of a grating are systematically biased toward the orientation of the immediately preceding trial (and to a weaker extent to the trials occurring before).

Serial dependence has originally been described as a low-level perceptual phenomenon, however this claim has been challenged by a recent paper (Fritsche, Mostert, & de Lange, 2017). Fritsche and colleagues used a slightly modified version of the original paradigm (Fischer & Whitney, 2014) with a larger sample size, and found that previous stimuli had a negative influence (consistent with classic literature on the tilt after-effect) and a positive serial dependence that was the consequence of past decisions (Fritsche et al., 2017). The latter finding is line with another recent report (Akaishi, Umeda, Nagase, & Sakai, 2014), which found that perceived motion direction of random dots was positively influenced by recent decisions. Decisions on easily distinguishable stimuli had a stronger influence on subsequent trials which was interpreted as a consequence of stimulus-independent internal states. Yet another recent paper (Bronfman et al., 2015) used the sequential sampling (Drift-diffusion) framework, and investigated the mechanism by which choices could influence subsequent evidence accumulation both with low-level perceptual (luminance) and high-level (numerical

evidence) accumulation paradigms. The main conclusion of this work was that choices not only affected subsequent decision criteria, but also the process of evidence accumulation itself by modulating the gain of the process.

The paradigm of Fischer & Whitney decoupled motor responses from the actual choice by using a response bar, which started from a random orientation at each trial. In many experiments, however, motor responses and choices are in fact connected, which posits the question: What is their respective influence of these two on serial dependencies? When both the effect of previous choices and motor responses was measured in a random dot motion paradigm, it was found that the motor responses only had a small -and not significant- influence on subsequent decisions, which were in turn strongly biased by past choices (Braun, Urai, & Donner, 2018). This finding is consistent with another recent paper, showing that most of the variability in perceptual decision making can be accounted for by noise in the inference mechanisms, while motor/selection noise accounts for only a small portion of response variance (Drugowitsch, Wyart, Devauchelle, & Koechlin, 2016).

One specific stimulus domain where various recent past influences have been described is face perception. Initially only positive sequential effects had been found with face stimuli (Lieberman et al., 2014). However, a more recent paper (Taubert, Alais, & Burr, 2016) reported that whether the effect of recent past in face perception was positive or negative depended on the facial attribute in question. Exactly the same stimulus could elicit a positive or a negative aftereffect depending on whether the task was to judge a stable (gender) or changing (mood) feature, with stable attributes eliciting positive, and changing attributes invoking negative aftereffects (Taubert et al., 2016). A related study showed that the same face stimuli could elicit an adaptation or a priming effect depending on whether they were followed by an ambiguous (adaptation), or by an unambiguous stimulus (priming) (Walther, Schweinberger, Kaiser, & Kovács, 2013).

While the debate whether serial dependence is a consequence of past stimuli or decisions is ongoing, it is also unclear whether these past effects influence the momentary percept or the decision only.

Although in principle, past decisions might influence decisions only, it has been suggested that serial dependence directly affects perception (Cicchini, Mikellidou, & Burr, 2017) in line with recent fMRI results showing serial influences in V1 (St John-Saaltink, Kok, Lau, & de Lange, 2016).

Hot Hand vs. Gambler's fallacy

While positive/negative serial effects have only recently become the focus of interest in psychophysics, in the decision-making literature similar effects have been known for a long time (Tversky & Kahneman, 1971). The hot-hand illusion during decision-making refers to the expectation that a streak of events will continue (framed as series of successes originally), and gambler's fallacy refers to the alternative expectation that an event having one of two possible outcomes will be followed next time by the alternative outcome. These tendencies to expect events to repeat or to alternate more than what would be expected from a true Bernoulli process have been described as repetition/alternation biases, respectively. Whether one or the other bias is found depends on what people assume about the generating process of the sequence (B. D. Burns & Corpus, 2004). For example, having to judge how random a sequence of events is, people have completely different expectations depending on whether the sequence is generated by a human (e. g. basketball throws) or by a random process (coin flips). Using event sequences with the exact same statistical properties, a Gambler's fallacy effect was found in human behavior when the sequence was supposedly generated by a random process, and a Hot Hand effect appeared when it was supposed to be the outcome of human actions (Ayton & Fischer, 2004).

Pattern Effects

In forming expectations on how a sequence of events will be continued, people clearly do not assume that events are independent or Markov since, among other factors, the pattern of the last few trials strongly influences perceptual decisions. For example, people can be faster in responding to either repetitions or alternations of recent events, if the event fits into the pattern of the recent past (Cho et al., 2002). Similarly, both repetitions and alternation responses become slower, when the pattern of recent events is violated. While this behavior has been described as irrational or even

superstitious, in fact, it can be rational in the real world, where most events are not independent of the past (A. J. Yu & Cohen, 2008). A related effect was found in relation to the perceived illusory motion direction (Maloney, Dal Martello, Sahm, & Spillmann, 2005). Here, participants had to respond to the motion direction of sequences of unambiguous events, followed by an ambiguous event. Maloney and colleagues found that people judged the movement of the ambiguous event as a continuation of the recent past. While in this experiment, the effect of past stimuli and decision were inseparable, Maloney and colleagues conducted a follow-up experiment, where participants did not have to respond to the inducers, only the ambiguous events. They found that even without having had to respond to past stimuli, the stimuli itself elicited a pattern effect, albeit a weaker one - with the effect disappearing for repetition sequences and preserved for alternating sequences (Maloney et al., 2005). The finding that repetition sequence effects rely on past decisions, but an alternating sequence of stimuli is sufficient to elicit a negative effect, is consistent with the proposal of past stimuli having negative and past decisions positive recency effects (Fritsche et al., 2017). A further interesting finding from this paper is the temporal dependency of the positive recency effect, which becomes stronger with a longer stimulus-response interval (Fritsche et al., 2017), further confirming that positive recency arises from high-level decision processes and not from the stimulus itself (which is the same regardless of stimulus- response interval).

Predictive Adaptation

While the last few trials or seconds before the stimulus already have a large and somewhat controversial influence on perceptual decision making, effects on such a short time-scale are not the whole story of contextual interactions. Events from many seconds or several minutes before could also influence a current perceptual decision. Since the influence of the last few trials is already very complex, one can expect that it could be even more challenging to draw conclusions about influences on longer time-scales. An interesting attempt was made a few years ago to look at long term influences on adaptation, using the phenomena of tilt aftereffect and binocular rivalry (Chopin & Mamassian, 2012). The results of this study showed that the strength of the well-known tilt after-

effect depended on long-term statistics, as if the perceptual system expected the statistics of recent experience to resemble past ones. The hypothesis that adaptation can be explained based on differences between recent probabilities and long-term prior probabilistic expectations is intriguing for many researchers treating perception as a Bayesian inference. However, the finding itself has proven to be controversial as the short-term negative influence combined with long-term random fluctuations of stimulus probability were sufficient to explain the obtained pattern of results (Chopin & Mamassian, 2013; Maus, Chaney, Liberman, & Whitney, 2013). It remains to be seen whether the perceptual system uses a certain period of past experience as a prior moving window based on which the statistics of sensory input is evaluated.

Probability matching and base-rate effects

Unlike in perception, the influence of long-term probabilities in the decision-making literature has been investigated for a long time. It is known that while people are sensitive to probabilities of events, they usually fail to adopt a strategy that maximizes potential pay-offs. Instead of a maximization strategy (picking the most probable outcome all the time), a common finding is probability matching: if the probability of a certain event A is e. g. 70%, people will choose A 70% of the time. The extent to which probability matching or maximization manifests itself in the participants' response can be strongly influenced by the framing of the problem. Faced with the same probabilistic outcomes in a gambling setup people are closer to maximizing, while facing a decision-making problem, a probability matching dominates their choices (Goodnow, 1955).

In perceptual decisions making, making one stimulus more frequent (also called as elevating its base-rate) in a binary discrimination task is a simple way to test the incorporation of probabilistic information. This approach was used in a categorization experiment, which found that people indeed incorporated this information into their judgments, with a bias to choose the more frequent option more often under perceptual uncertainty (Bohil & Wismer, 2014). While the authors did not interpret their results that way, the findings seem to be consistent with a probability matching strategy. Using a more complex spatial localization task that required integration of auditory and

visual information, Wozny and colleagues found that the majority of people exhibit probability matching when trying to infer the cause of the observed sensory input (Wozny, Beierholm, & Shams, 2010).

While probability matching could be viewed as a fallacy, as it fails to maximize immediate rewards, in some cases it is rational if the potential learning benefits are taken into account. Faced with a binary choice problem, choosing Option 1 all the time does not provide any information about Option 2, so if the potential reward probabilities change over time or the decision maker is uncertain about the pay-off structure, choosing the less good option can be a rational strategy. Therefore, considering the necessity to learn, probability matching can be good strategy that balances the trade-off between exploitation and exploration: it can reap sufficient rewards but is still flexible enough to allow further learning (K. J. Burns & Demaree, 2009; Gaissmaier & Schooler, 2008).

A further indication that probability matching should not be taken as an evidence of human irrationality comes from contrasting it with explicit reasoning studies. In probability estimation tasks, since the classic work of Tversky & Kahneman, it is known that people are remarkably insensitive to base-rate information (Kahneman & Tversky, 1972). Interestingly, presenting more information in an explicit probability estimation problem can facilitate ignorance of base-rates (Bar-Hillel, 1980). Therefore, considering the complete failure to use probabilities in reasoning, the incorporation of stimulus probabilities in perceptual decision-making with a probability matching strategy can be considered a remarkable feat and not a fallacy.

Although it seems established that the brain stores and uses probabilistic information, even enthusiastic supporters of the “Bayesian Brain” concept do not claim that the brain is performing exact Bayesian inference as it is known to become intractable as the number of variables gets large. An influential potential solution posits that the brain uses sampling to represent probability distributions. Sampling is mentioned here since probability matching is exactly the expected outcome of sampling from a probability distribution (Sanborn & Chater, 2016). Therefore, further investigating

the mechanisms of probability matching could provide insights to the fundamental problems of representing and using probabilistic information in the brain.

Perceptual and Statistical Learning

While the influence of long-term probabilities on visual perception became a focus of interest only recently (Chopin & Mamassian, 2012), other long-term influences on perception have been investigated for a long time under the label of learning. The field of perceptual learning is investigating specific improvements in a low-level perceptual task over several days of practice. Specific means that for example, practicing orientation discrimination at a given orientation at some retinal location, can only increase performance around that particular angle and location (Sagi, 2011). Interestingly, slight changes in the training regime can change that, as with a double training paradigm-with interleaved trials with a different task- the learning proved to be generalizable, casting some doubt on the prevalent low-level interpretations of perceptual learning (Xiao et al., 2008).

A different literature with potential long-term influences on perception arising from spatial or temporal regularities in the stimulus stream has been described as visual statistical learning (Fiser & Aslin, 2001). While statistical learning has been described as an automatic and implicit process, it does require attention to the stimuli to manifest (Toro, Sinnett, & Soto-Faraco, 2005; Turk-Browne, Jungé, & Scholl, 2005). Initially, it has not been clear whether the implicit knowledge acquired via familiarization and usually measured at subsequent familiarity test could influence perception. However, a recent study showed in a temporal statistical learning paradigm that people are faster in processing objects that have been predictable (Barakat, Seitz, & Shams, 2013). Interestingly, not only the processing speed of the participants but even their perceptual sensitivity (d') was lower for elements which had been predictable due to the previous statistical learning training. Notably, the effect persisted even when the element was not predictable during test, only in the previous learning block, suggesting a general increase in sensitivity to predictable items as a consequence of statistical learning (Barakat et al., 2013). The fMRI literature also suggests that even task irrelevant-implicit

statistical learning can influence the neural responses to predictable items (Turk-Browne, Scholl, Johnson, & Chun, 2010). However, unlike in Barakat et al.'s (2013) behavioral study, the effects of predictability have only been investigated for the items that were actually more predictable in the fMRI study. Returning to the behavioral results of Barakat et al. (2013), it is remarkable that simply due to a stimulus being more predictable by the statistical structure of the input during learning, it becomes easier to detect later, even in test situations when it is not predicted by the preceding stimuli anymore. This posits a long-lasting effect on perceptual sensitivity by statistical learning. While the reasons for this phenomenon remain to be fully explained, the results make the theoretical line that separates statistical from perceptual learning fuzzy (Fiser, 2009; Gold & Stocker, 2017).

A sensitive measure of statistical learning uses reaction times; accelerated responses to predictable events can show learning even when the statistical information is not recallable on a subsequent familiarity test (Kim, Seitz, Feenstra, & Shams, 2009). An alternative paradigm uses semi-predictable event sequences, with interleaved predictable and random events (Howard Jr & Howard, 1997; Nemeth, Janacek, & Fiser, 2013). In this paradigm, a speeding up in responses to predictable parts of the sequence is the most common finding, without explicit awareness of any regularity in the sequence. In a related study, the interaction between long-term stimulus predictability and recent patterns of input was measured (Wilder, Jones, Ahmed, Curran, & Mozer, 2013). The short-term influence of repetitions/alternations was measured similarly to the method used in Yu & Cohen (2008). However, unlike in previous studies, the long-term frequency of repetitions was also manipulated. Wilder et al.'s main finding was that effects on decision making due to long- and short-term patterns interacted in an additive manner, and thereby the relative influence of short-term patterns was further enhanced by long-term influences, suggesting that people can simultaneously track environmental regularities on multiple time-scales (Wilder et al., 2013).

Long-term influences on neural responses

There is a vast neuroscience literature on short-term influences of past stimuli, and it has been found that short-term effects could depend on long-term learning and expectations. For example, fMRI signals show that the prevalent effect of repetition suppression depends on the long-term probability of stimulus repetitions (C. Summerfield, Trittschuh, Monti, Mesulam, & Egner, 2008). In this experiment faces were used as stimuli, and repetition suppression was stronger when a face image was expected and weaker when its appearance was surprising given the long-term experience. While this result has since been replicated with faces (Kovács, Iffland, Vidnyánszky, & Greenlee, 2012), repetition-probability-dependent repetition suppression effects in fMRI signals were not found for common objects (Kovacs, Kaiser, Kaliukhovich, Vidnyanszky, & Vogels, 2013) nor with monkey single-cell recordings (Kaliukhovich & Vogels, 2010). A possible resolution of this discrepancy was suggested by a study using a paradigm based on a regular alphabet and a novel font: the dependence of repetition suppression on long-term repetition probabilities emerged only for stimuli with which people had a vast amount of experience (Grotheer & Kovacs, 2014).

Although the effect of long-term stimulus probability on single-cell responses at different areas is debated (Bell, Summerfield, Morin, Malecek, & Ungerleider, 2017; Vinken & Vogels, 2017), the emerging picture seems to suggest that while V1 responses (in an oddball paradigm with rats) are only influenced by short-term adaptation, higher-level visual areas also show specific increase in neural response for rare stimuli (Vinken et al., 2017). Findings from the monkey IT cortex suggest that neurons in higher visual areas represent stimulus probability with reduced overall response but with enhanced information content for decoding highly probable stimuli (Bell, Summerfield, Morin, Malecek, & Ungerleider, 2016). Using an MVPA method, a human fMRI study showed that expected stimuli elicit a smaller overall BOLD response but a sharper stimulus representation already in the primary visual cortex (Kok, Jehee, & de Lange, 2012). In sum, there exist ample evidence suggesting that past expectations influence the processing of visual stimuli in the brain, but a stark discrepancy emerged between results based on animal recordings and human fMRI data. While the human

literature reports influence of long-term expectations both at high level (C. Summerfield et al., 2008) and in early visual areas (Kok et al., 2012), the animal literature suggest that those influences appear only in higher areas (Vinken et al., 2017), or they can be completely absent (Kaliukhovich & Vogels, 2010).

Value Based Decision Making

While the literature linking perceptual decision making and long-term probability effects is scarce (but see Bohil & Wismer, 2014; Chopin & Mamassian, 2012), there is a vast literature on learning effects in value-based decisions making (Behrens, Woolrich, Walton, & Rushworth, 2007; Schrater & Acuna, 2010; Steyvers, Lee, & Wagenmakers, 2009). Although there are obvious differences between value based- and perceptual decisions, investigating the two phenomena in the same framework could be fruitful, as there are remarkable similarities in the underlying mechanisms. In fact, an analogy between sampling sensory information and sampling from memory has been recently proposed (Shadlen & Shohamy, 2016). The idea of treating value-based and probability-related effects in the same framework is promoted by evidence that prior probability and economic value would bias perceptual decisions in a similar manner (Mulder, Wagenmakers, Ratcliff, Boekel, & Forstmann, 2012). Using a modeling approach based on drift-diffusion, several studies reported that both economic value and prior probability changed the starting point (bias) of the evidence accumulation process without affecting the speed of evidence accumulation itself (Mulder et al., 2012; C. Summerfield & Koechlin, 2010). This finding is in contrast with the short-term influences described above, which can modify the speed of evidence accumulation (Bronfman et al., 2015).

Results from long-term effects in value-based decision-making show that people are not only capable of adapting to changes of the pay-off probabilities, but they can also learn the probability of changes, in other words, the volatility of the environment (Behrens et al., 2007). This can be achieved by changing the relative weighting of past vs. present information through adjusting the learning-rate: stable environments require a lower, while volatile environments a higher learning rate. Using a hierarchical modeling approach to a numerical prediction task, people were shown to increase their

learning rate when the prediction error was large and decrease it when the prediction error was small (Nassar, Wilson, Heasly, & Gold, 2010). A possible explanation of this finding is that a large prediction error enables detecting a change-point in the environmental probabilities, and after the quick adjustment occurs due to enhanced learning, the learning rate can decay quickly to a low value to reflect that the conditions became stable again.

Volatility and perceptual decision making

In the past few years, there were several attempts to link the influence of environmental volatility to the problem of perceptual decision making (Glaze, Kable, & Gold, 2015; Norton, Fleming, Daw, & Landy, 2017; C. Summerfield, Behrens, & Koechlin, 2011). In this complex problem, the decision maker faces uncertainty at two distinct levels: first, there is uncertainty in the stimulus identity itself at any given moment; second, there is uncertainty in the stimulus probabilities or category boundaries, which can change over time. Studies using a random dot motion paradigm with within trial direction changes and a normative approach reported that when the conditions are stable, people can integrate evidence optimally (Glaze et al., 2015). However, human behavior can be approximated better by a leaky accumulator in a volatile environment suggesting that people can take into account probability of changes in the environment as they accumulate evidence (Glaze et al., 2015). While traditional psychophysical approaches use Signal Detection Theory and psychometric curves to assess decision criteria from participants' responses, human adults can also be explicitly queried about their decision criteria. This latter approach might obtain somewhat different thresholds from what would be inferred directly from the responses derived from SDT, Furthermore, different measures suggest different learning rates: decision criteria based on explicit queries were updated faster than criteria calculated from the actual responses in a categorization task (Norton et al., 2017). However, explicitly query allows the trial-by-trial assessment of decision criteria, which would not be possible with the more conventional method. In a recent study, such an explicit querying approach was compared to traditional psychophysics measures in an orientation discrimination task embedded in both static and volatile environments (Norton et al., 2017). As

expected, participants were faster to readjust when the conditions were dynamic. However, the results also showed that even when conditions were stable, people still adjusted their decision criteria continuously - and sub-optimally. Such a suboptimality can be a consequence of a rational decision making process, where the decision-maker is uncertain about the structure of the task (Schrater & Acuna, 2010). A very recent paper tried to disentangle individual differences in decision making in a volatile environment (Glaze, Filipowicz, Kable, Balasubramanian, & Gold, 2018). In this paradigm -as in a previous study of (Glaze et al., 2015) - people had to detect change points in the binary underlying source that places stimuli at one of two positions corrupted by positional noise, and found that the variability of human decision making could be well explained by a bias-variance trade-off: some people “overfitted” the noise in the data by adapting to random fluctuations in the underlying process, while others were insensitive to changes that could be informative about the underlying generative process.

A related but somewhat less investigated field is the explicit estimation of environmental probabilities. In a recent experiment, participants had to estimate the proportion of differently colored items in a box, based on individually presented samples. The main finding was that while people made small adjustments to their probability estimates continuously (even when conditions were stable, as in: Norton et al., 2017), sometimes, they completely reset their estimate and discarded all information from the past. This behavior cannot be explained by the conventional high-learning rate models (Gallistel, Krishan, Liu, Miller, & Latham, 2014).

Perception vs Decisions

The studies above described different ways of how long-term past might influence perceptual decision-making. Many of these proposals used a shift in decision criteria or bias (as the starting point in drift diffusion) to incorporate these long-term expectations. This approach stems from Signal Detection Theory, separating perceptual decision making into a perceptual and decision phase (Stanislaw & Todorov, 1999). While this proved to be a fruitful approach, there are also suggestions that sensitivity and bias are in fact intimately related (Wei & Stocker, 2017). According to this

proposal, separating expectation effects based on whether they affect perceptual or post-perceptual decision processes might be impossible in many cases. Furthermore, there is ample evidence from neuroscience that top-down expectations influence early visual processing (Kok, Bains, van Mourik, Norris, & de Lange, 2016; for a review see: van Kerkoerle, Self, & Roelfsema, 2017). In contrast, an alternative proposal based on a behavioral approach has suggested recently that expectations solely affect post-perceptual cognitive processes, and other top-down effects on perception can be explained away by attention or some other factors (Firestone & Scholl, 2014). While it is true that the separation between expectations and attention based top-down effect can be non-trivial, it has been shown that they affect neural representation of visual stimuli differently (J. Jiang, Summerfield, & Egnér, 2013; Kok et al., 2012).

Optimality of Perceptual Decisions

A discussion in the literature relevant to this thesis deals with the question of whether people are Bayes optimal decision makers in the sense that they can combine prior expectations with uncertain sensory input in a statistically optimal fashion (Kersten, Mamassian, & Yuille, 2004). Unlike results in “cognitive” or economic decision making tasks, where people seem to be subject to a large number of biases and fallacies (Tversky & Kahneman, 1974), a first glance at the sensory-motor literature suggests that people are, indeed, optimal in perceptual decision making (Ernst & Banks, 2002; Körding & Wolpert, 2004). A more direct comparison of economic and motor decision-making suggests that the same pay-off odds can evoke different probability distortions depending on whether they are framed in a classic economic- or in a motor decision problem (Wu, Delgado, & Maloney, 2009). However, the evidence that people are optimal in integrating priors and likelihoods during sensory-motor decisions is mostly limited to stimuli with simple Gaussian probability distributions, while people have severe difficulties to learn and use more complicated prior distributions for guiding their perceptual decisions (Acerbi, Vijayakumar, & Wolpert, 2014). Thus, it is not clear how optimally humans can cope with more complicated prior distributions if at all. An

alternative reason why people were reported to behave sub-optimally in a perceptual experiment could be the artificial nature of the applied experimental paradigms, in which stimuli deviated greatly from what would be expected from the statistics of the natural environment.

Summary on past influences in perceptual decision making

To summarize findings from the various paradigms described above, previous trials influence perceptual decisions in different ways: past stimuli often elicit a negative, decisions a positive influence, with stimulus type, presentation time, and inter-trial interval as potentially crucial factors. This influence of the recent past is suboptimal if trials are independent, however it could be crucial in adjusting to changes when the environment is temporally correlated but volatile. Looking at the influence of the past at longer time-scales, most models assume a gradual discarding of the past that is implemented by the learning rate of e. g. reinforcement learning models, but sudden discarding of all past information has also been proposed (Gallistel et al., 2014). On the other hand, in a different framework, the long-term experience is not discarded, but instead used as a prior, based on which recent information is evaluated (Chopin & Mamassian, 2012). The exact relation between these models, and the extent to which differences in the experimental paradigms (eg.: stimulus timing, response method, feedback) could explain differences in how people handle past information in perceptual decisions making is unclear at present time.

Despite the fact that most of cognitive psychology and psychophysics uses the assumption of independent experimental trials, the picture emerging from the literature is that sequential perceptual decision-making trials are anything but independent. There seems to be a complex looping interaction, where incoming stimuli and our momentary decisions about them - via internal representations - will influence the future perception of similar stimuli. These effects are relatively well explored in shorter time-scales – but not without controversies. In contrast, there are only a few established findings from longer time-scales mostly from the last few years, while the exact relationship between short- and long-term effects in perceptual decision making is largely unresolved.

Visual Exploration

Moving away from perceptual decision making, the field of visual attention and visual search has a number of relevant findings about the influence of past information on visual processing from various domains and experimental paradigms. Below, I will start by reviewing findings from simple visual search and move on to more complex real-life scenarios, continue with models of visual search, and finish with integrating the results on the active nature of vision into my survey.

Visual Search

Before the use of eye-tracking technology became prevalent, reaction times from searching for a target in an array of distractor stimuli could be used as indicators of visual search. Increasing the number of distractors usually makes the task harder, manifesting in longer search times. The extent to which search times change as a function of the number of distractors is called the search slope and is a dominant measure for studying visual search (Wolfe, 1998). Based on these search slopes measures, the provocative claim has been put forward that visual search has no memory at all (Horowitz & Wolfe, 1998). To support this surprising conclusion, the authors had analyzed visual search in a standard unchanging and dynamic search arrays. They found that search slopes were identical in dynamic and stable arrays, despite the past displays having no predictive power in dynamic arrays. This has been interpreted as support for the memory-less nature of search. However, a more careful look at the proportion of correct answers instead of search times challenges this claim (M. Peterson, Kramer, Wang, Irwin, & McCarley, 2001). Furthermore, moving away from search slopes and using eye-tracking data by analyzing the distribution of revisited location during exploration of similar search arrays suggests that visual search does have memory (M. Peterson et al., 2001).

The effect of past probabilities can be easily investigated in visual search by making the target more frequent in an area of the screen. There is evidence from visual search times that people manifest probability matching when faced with biased target location probabilities (van der Heijden, 1989), suggesting that visual search is sensitive to environmental probabilities, but does not utilize this information in the naively optimal manner, which is line with several findings from perceptual decision making results described in the first part of this chapter. However, the influence on long-term probabilities on visual search does not always manifest itself. For example, although people are faster to react and make more saccades to the more frequent locations, when short-term probabilities are controlled for, these effects can be fully explained by repetition priming (Walthew & Gilchrist, 2006). Importantly, only slight modifications in the paradigm of Walthew & Gilchrist completely changed these results, and demonstrated that long-term target probabilities do attract visual search, even when short-term influences are controlled for (J. L. Jones & Kaschak, 2012). Yet another study found that long-term probabilities affected visual search direction only when the target location probabilistically depended on the direction of saccades. Without this gaze-contingent manipulation, a high probability of a region containing the target was insufficient to bias visual search (Paeye, Schütz, & Gegenfurtner, 2016).

Contextual Cueing

Beyond simple target probability effects, a paradigm that allows the investigation of more complex statistical influences on visual search is called contextual cueing. In a typical contextual cueing paradigm, participants are responding to a target (eg: left/right oriented T) in an array of - seemingly randomly arranged- distractors. If the arrangement of some of the distractor arrays is repeating over time, people become faster in responding to targets within the repeating displays. This effect emerges even when participants are unable to tell apart repeating patterns from random ones, suggesting that visual search is a sensitive measure of implicit statistical knowledge about complex stimuli (Chun & Jiang, 1998). While the original finding was interpreted as faster deployment of visual

attention to targets predicted by the context, this explanation of the underlying mechanism has been challenged, as other factors such as response selection were indicated to play an important role in contextual cueing (Kunar, Flusberg, Horowitz, & Wolfe, 2007).

More recently, a study using a similar paradigm of manipulating the probability of target occurrence at different areas of the screen found that participants responded faster to targets at the areas with high target appearance probability (Y. V Jiang, Won, & Swallow, 2014). This effect of spatial probability cueing on search times without explicit awareness also influenced a second measure of visual search, the location of the first saccade. This bias was relatively persistent as participants searched the initially rich quadrant above chance, even when the actual distribution of targets later during the experiment became balanced. The effect emerged both implicitly and explicitly and it proved to be persistent by surviving an explicit instruction about balanced test probabilities. Contextual cuing does not only work by facilitating search at a given spatial location but also by a spatial co-variation at different positions (Chun & Jiang, 1999). In this paradigm a large number of shapes and distractors were used, and the orientation of distractors was predictive of the target location. This predictive information facilitated visual search times, as compared to a random mapping, despite the location of search targets and distractors varying from trial to trial (Chun & Jiang, 1999). This result further suggests that visual search is sensitive to complex statistical relationships.

Visual Search, Memory, Reading

The next level of generalization requires switching from well-controlled but artificial search arrays, to studying visual search by using images of real-world scenes. A classic finding using such tasks is that the effect of bottom-up saliency is weaker, and top-down factors can dominate as people have strong expectations about what kind of objects to expect at different locations in real scenes (Loftus & Mackworth, 1978). In real world settings, different memory components influence search patterns

differently. First, observers have semantic knowledge about typical locations of real world objects, which has an enormous influence on where they search for them (Võ & Wolfe, 2012). Second, observers might have episodic memories about the locations where they have seen a given target during an experiment. Interestingly, the influence of episodic memories on visual search in real world scenes can be weak: when observers searched for different objects in the same scene repeatedly, there was only a small advantage in search times (Võ & Wolfe, 2013). Not surprisingly, during repeated search for the same objects in the same scenes, the benefit of episodic experience was very strong (Võ & Wolfe, 2013), as saccades could go directly to the target if it was already encountered at a given location (J. J. Summerfield, Lepsien, Gitelman, Mesulam, & Nobre, 2006).

While most research would agree that fixations and looking times are sensitive measures of memory-related processes during search in real-world scenes, there seems to be a debate about the role of implicit and explicit memories in influencing eye-movements. A common finding is that eye-movements are influenced by knowledge that is not available to conscious awareness/explicit report (a review: Hannula, 2010). However, other studies found that memory effects on eye-movements do not reflect implicit processes: looking times at manipulated parts of recently presented scenes (Smith & Squire, 2008) or even at previously studied images during old/new discrimination (Urgolites, Smith, & Squire, 2018) display learning effects only if those memories are also amenable for explicit reports. Similar conclusions were drawn from experiment with children (Koski, Olson, & Newcombe, 2013).

The temporal order of earlier presentation of stimuli also influences the order of fixations when the same stimuli are presented simultaneously at different spatial locations (Ryan & Villate, 2009). This suggests that eye-movements can be a sensitive measure of serial episodic experience. A specific domain that allows well-controlled investigations of eye-movements is reading. Reading studies have shown that eye-movements are not only sensitive to word frequency (Rayner & Raney, 1996), but also to transitional probabilities between words (McDonald & Shillcock, 2003), suggesting that eye-movements are a good potential measure of acquired statistical properties of the environment.

Visual attention, learning and prediction

The above-mentioned memory effects on visual exploration can explain only a small fraction of the variance in eye-movements looking at either artificial or natural scenes, since there are many other factors that can influence the location of fixations. First, there are inherent biases when people do laboratory tasks on a computer, for example to look more at the middle of a scene (Tatler, 2007) or proceed with visual search from left to right (Spalek & Hammad, 2005). There are also domain specific biases, such as the typical fixation patterns observed when looking at faces (Walker-Smith, Gale, & Findlay, 2013). There are many top-down factors, which are hugely task dependent (Hayhoe & Ballard, 2005) , while without a task there is a strong influence of bottom-up/saliency driven effects on fixation locations (Parkhurst, Law, & Niebur, 2002). Investigating the interaction of bottom-up and top-down factors in saccade target selection it has been shown that with increasing saccade latency people are more influenced by the value of targets and less by visual saliency, confirming the intuition of stronger impact of bottom-up factors on faster actions (Schutz, Trommershauser, & Gegenfurtner, 2012). As all these different mechanisms can influence momentary eye-movements to a variable extent, it is challenging to quantify all of them for an integrated model of visual exploration of even simple artificial scenes, not to mention real world scenarios. Nevertheless, there have been several attempts towards modeling human visual attention during the exploration of artificial and natural scenes (for a review Borji & Itti, 2013).

Using artificial scenes, and modeling only the number (and not the location) of fixations until a target is detected in visual noise, visual search was found to be close to optimal in selecting fixations that minimize uncertainty about the possible target locations (Najemnik & Geisler, 2005). Interestingly, for a near optimal performance in this task, it is sufficient to use a strategy based on the “inhibition of return”, and there is no need to integrate any other information across saccades. Although the work by Najemnik & Geisler was a breakthrough in the sense that it managed to link visual search to a Bayesian Ideal Observer, the limited scenario of searching for a single Gabor target in $1/f$ noise and the fact that it could assess only the number of fixations without their locations leaves many

questions open. A more complex approach modeling eye-movements in a visual discrimination task also found that sequential fixation locations were close to what followed from an optimal strategy (Renninger & Coughlan, 2007). However, a local uncertainty reduction model was sufficient to approximate this behavior: people looked at the regions of the scene they were the most uncertain about. The proposal that people use extrafoveal visual input to guide search optimally has been further challenged by an experiment which directly manipulated the availability of extrafoveal information (Morvan & Maloney, 2012). By manipulating the eccentricity of the targets, these researchers found that humans were far from being optimal; in fact, they were not sensitive to the experimental manipulation at all. The opposite conclusion has been reached by a study that used a similar approach but with events of different lengths instead of different eccentricities (Hoppe & Rothkopf, 2016). This paper found that people could adapt to task requirements and adjust their visual sampling behavior in order to maximize target detection performance given their limited perceptual sensitivity. A potential reason for the contradicting results could be that the failure to optimize was found with respect to fixation location, while the successful adaptation to task requirements required adjusting the length of the fixations (Hoppe & Rothkopf, 2016; Morvan & Maloney, 2012). Interestingly, a very recent study found that even the very basic behavior of timing visual blinks was adaptive to environmental regularities (Hoppe, Helfmann, & Rothkopf, 2018).

Despite the fact that modeling human visual attention in simple search tasks is already quite challenging, there have been attempts to investigate visual attention to more complex influences arising from stimulus statistics. An interesting proposal built on visual search reaction times suggests that the presence of structured visual information could attract human visual attention (J. Zhao, Al-Aidroos, & Turk-Browne, 2013). In a combined statistical learning and visual search paradigm, these authors found that the mere presence of statistical regularities can attract spatial attention more than areas containing only random stimuli (J. Zhao et al., 2013). Notably, it has not been clarified whether this effect had an influence on eye-movements as well. It is interesting to contrast this finding with earlier proposals stating that unpredictability and surprise attract visual attention

(Duncan & Humphreys, 1989). One problem with this latter proposal is that the most unpredictable stimulus possible is pure white noise, which is clearly not the best candidate for attracting human visual attention. To account for the problem that purely predictability-based descriptions are not sufficient to describe the allocation of human attention, a Bayesian measure of surprise was developed based on Kullback-Leibler divergence, and was calculated on different features of real video clips (Itti & Baldi, 2009). The same videoclips were played to human observers as their eye-movements were recorded, it was found that this Bayesian measure of surprise is a successful predictor of human visual attention, better than traditional saliency based metrics (Itti & Baldi, 2009).

Although they investigate similar questions with very different paradigms, there seems to be a fundamental contrast between the findings of Zhao et al (2013) and Itti & Baldi (2009). While Zhao's work shows that regularities attract human attention, Itti & Baldi's findings demonstrate that the more surprising an event is (i.e. the less regular it is), the more it will attract eye-movements. An intriguing idea to resolve this contradiction is that human observers prefer looking at input that is complex enough not to be trivial, but not too complex or completely unpredictable. This idea was tested with human infants both with visual event sequences and auditory input, and it was found that infants pay attention the longest to event sequences that have an intermediate level of complexity as measured by information entropy (Kidd, Piantadosi, & Aslin, 2014, 2012). In a similar vein, another infant study found that anticipatory looking only occurs if the visual event is probabilistic, and not when it is fully deterministic, since presumably deterministic events do not carry enough information to be interesting (Téglás & Bonatti, 2016). Importantly, the proposal of Kidd et al. (2012) shifted the focus from pure stimulus complexity (i.e. information content) to information content relative to the knowledge of the observer. Unfortunately, their study did not measure learning, and therefore, could not test whether the shift in internal knowledge of the observer has any effect on their behavior thereby confirming this preference for stimuli of intermediate complexity.

Despite the above described advances using eye-movements as indicators of human information gathering from the environment, an important hurdle in the exploration of active vision is that eye-movements do not fully reflect what information is processed by the observer, as a lot of information can be gathered from the visual periphery. In fact, information from the visual periphery is one of the main underlying factors in the identification of saccadic targets (Yamamoto & Philbeck, 2013). In order to study how people sample information from the environment based solely on their internal representations, a possible solution is to use a gaze-contingent “window”. In such a paradigm, participants have access to visual information only from the local vicinity of areas on the screen where they fixate, while the rest of the image is not displayed. With this method, eye-movement patterns will be a true reflection of human information gathering as it is impossible to get information from the visual periphery at any time. This approach was used in a paradigm where participants had to perform a binary categorization of textures as stripy or patchy (Yang, Lengyel, & Wolpert, 2016). Efficient distinguishing of these textures required exploratory eye movement patterns that were differentially influenced by the incoming visual information. The eye movements of the participants in the experiment largely followed the optimal strategy, and while their performance fell short of that of an ideal observer, the exploration patterns were remarkably similar to predictions from the ideal model, confirming that human are capable of efficient active sensing (Yang, Lengyel, et al., 2016).

Active Learning

The review above shows that while there is literature suggesting that people attend to informative parts of the sensory environment, these studies either did not measure learning or did not link learning to the attentional biases (Kidd et al., 2014; Kidd, Piantadosi, et al., 2012; J. Zhao et al., 2013). Therefore, these studies cannot address the question of whether and how these effects might emerge with experience, if at all. On the other hand, studies on active sensing showed that people can query the environment in a manner that maximizes information gain to solve a particular task on a given trial (Yang, Lengyel, et al., 2016). A missing point in the literature is the link between these

two phenomena: do people use their active sensing to maximize information gain even without an explicit task, and does this maximization emerge through experience? It has been suggested that there is no need for an explicit task or rewards to initiate search for maximal information, since acquiring learnable information is itself a reward for humans (Gottlieb, Oudeyer, Lopes, & Baranes, 2013). There is also some fMRI evidence supporting this claim, showing that visual information that disambiguates perceptual uncertainty has a rewarding effect on the brain (Jepma, Verdonchot, van Steenbergen, Rombouts, & Nieuwenhuis, 2012).

These findings on “information as reward” suggest that people could not only be capable of active sensing, which focuses on reducing uncertainty in a given moment (trial), but that they might also be active learners. As opposed to active sensing, an active learner would not only search efficiently to reduce uncertainty at a given trial but would also be capable of searching for information efficiently to reduce uncertainty over a longer period (across trials). This would require constantly integrating information from the sensory environment with the developing internal representations and changing search behavior accordingly. There is already some evidence based on a yes/no question game that people are capable of active searching: children asked questions that efficiently reduced the search space, in a manner that is sensitive to the environmental probabilities (Nelson, Divjak, Gudmundsdottir, Martignon, & Meder, 2014). However, it is important to distinguish between the concept of active learning as used in education, and the related but different computational notion used in the present thesis, which is coming from the machine learning literature (Winterbottom, 2012). The educational concept simply posits that being actively engaged with the material facilitates learning more than passively receiving it. The computational approach proposes a reason for this learning advantage: when people are actively engaged with a material, they can select from the stimuli what is the most informative for them at the current stage of learning and that is why they can learn better.

Curiosity

A human tendency to seek learnable information is related to curiosity, which might be a hard-to-define concept, but is at the core of a rapidly emerging new area of research (Kidd & Hayden, 2015). The similarities between search for visual information and curiosity is apparent if we compare the above described findings about infant visual attention (Kidd et al., 2014; Kidd, Piantadosi, et al., 2012) and findings about curiosity to answers to trivia questions. In the latter, an inverted U-shape relationship was found between confidence and curiosity: people's desire to find out the answer to a question is the highest when they have an intermediate level of confidence about the answer (Kang et al., 2009). This pattern is remarkably similar to the findings of Kidd et al. (2012, 2014), where an inverted U-shape relationship was found between stimulus complexity and the length of maintaining attention. Reinforcing this similarity between attention and curiosity, Kang et al.'s work (2009) linked curiosity to learning in a way that confirmed predictions of active learning theory: a few weeks later, participants remembered the information proportionally to the extent they were curious about it beforehand (Kang et al., 2009).

Summary on visual exploration

Taken together, the above described findings and theories suggest that people are inherently curious and active learners: if there is any useful structured information in the sensory environment, people pay attention to this usable part of the sensory input, and this attentional bias could facilitate their learning about that information selectively. In the case of vision, it is likely that this selection is done—at least partially—by eye-movements, which will influence what information the observer gets, and how the observer updates his/her internal representation of the environment. The updated representation, in turn, influences further information sampling, thus the two processes of learning and information sampling interact in a loop-like manner guided by the different parts of the sensory input that are more or less informative in the light of the momentary internal representations and goals. While this proposal is intriguing, there is little experimental evidence to support its main claims about the loop between visual attention and learning. It is not even established whether there exists

such a constant rewarding effect of finding learnable information in the sensory input. Furthermore, the extent to which people search for usable regularities in the input and the efficiency of this search is unknown. Finally, similarly to perceptual decision making, the sensitivity of visual search to sequential environmental statistics both on long- and short time-scales is largely unexplored.

Outlook on the goals of the thesis

In this introductory chapter, I reviewed two broad research areas with their current results and open questions. In my thesis, I will focus on two main topics related to these two areas: eye-movement-related active learning during visual statistical learning (Study 1) and long terms serial effects in visual perceptual decision-making (Study 2).

In the first domain (Chapter 2-3, Study 1, 6 Experiments), I will seek answers to the following questions: what is the relationship between learned statistics of the visual environment and visual search? Can people use implicitly learned regularities to guide eye-movements during visual exploration? Can eye-movements be used as an indicator of learning hidden regularities? Do interactions between learning regularities and visual search happen automatically, or do they rely on an explicit task? What kind of statistical representations influence eye-movements?

In the second domain (Chapters 4-5, Study 2, 7 Experiments) I will explore the following questions: To what extent and time scale can past probabilities influence perceptual decision making? Is there an interaction between long-term probability influences and short-term serial effects? How are such interactions influenced by changes in stimulus probabilities over-time? Can change dynamics influence how people update their internal models on the statistics of the sensory input?

Chapter 2

Active Statistical Learning

Summary

To investigate the learning process when people receive sensory input with underlying regularities, we combined statistical learning with eye-tracking in a set of three experiments. We used a novel gaze-contingent spatial statistical learning paradigm that enabled tracking the influence of stimulus statistics on visual exploration patterns. In the first experiment, using an explicit learning paradigm, we found that several different temporally emerging measures of visual exploration can predict learning performance, thereby validating our novel paradigm. To test whether our findings generalize to implicit learning, we ran two additional experiments that were almost identical to Experiment 1, differing only in instructions (Experiment 2) or in instructions and length (Experiment 3). Using Implicit instructions, we found that robust statistical knowledge can emerge without any easily detectable effect on eye-movements. However, based on a more sophisticated analysis of the eye movement statistics, we could still follow observers' learning, identify the best implicit learners, and track a late emergence of direct eye movement patterns found during explicit learning. These results suggest that there is a smooth link between implicit and explicit statistical learning, and that eye movements in our novel method are appropriate to trace this learning process and the transfer from implicit to explicit knowledge.

Introduction

There is a vast literature on statistical learning documenting that human adults are capable of acquiring the regularities of the sensory input in the auditory, haptic (Conway & Christiansen, 2005) and visual modalities implicitly (Fiser & Aslin, 2001, 2002). Apart from adults, infants (Saffran, Aslin, & Newport, 1996) and many different animals species (for a review: Santolin & Saffran, 2018) are also sensitive to the regularities of the environment. This broad spectrum of findings suggests that implicit statistical learning could be a crucial mechanism, which enables efficient processing of the probabilistic properties of the sensory input. However, despite the widely held assumption that statistical learning is a fundamental, automatic and modality independent mechanism, we know surprisingly little about how this mechanism works. Specifically, there are two intertwining main questions that prevent statistical learning from being coherently integrated in the larger scheme of human learning: first, its unclear relation to explicit learning, and second, lack of knowledge about the characteristics of its gradual emergence as the function of the sequentially accumulating sensory information.

Regarding the first question, the literature of explicit and implicit learning is enormous (Ellis, 2009; Reber, Walkenfeld, & Hernstadt, 1991; Willingham & Goedert-Eschmann, 1999), and even within the field of statistical learning there is some confusion on what is meant by implicit and explicit learning. Statistical learning is implicit, because the learner does not know about the existence and nature of regularities in advance and can only discover them over-time in an unsupervised manner. This is very different from explicit learning, where participants are instructed on what they are supposed to remember. The representation that emerges via an implicit learning process might still become “conscious” and can still be similar to what is learned through an explicit task. However, the extent, to which the representations emerging via the two different learning processes are similar, and how they relate to the classic explicit/implicit memory division (Graf & Schacter, 1985) is unknown. A way

to assess the relationship of the learned representations could be to use different behavioral measures of learning, and assess if they are affected similarly by implicit/explicit learning.

Relating to the second question, even if what is acquired via the two learning methods is alike, the *process* of learning is probably very different. Implicit learning could be more gradual, with a continuous incorporation of sensory information, while explicit learning could proceed in a more step-wise manner. However, the actual similarities and differences between explicit and implicit learning dynamics are not known and could only be assessed by developing sensitive measures of learning processes.

Regarding the emergence of implicit knowledge, the large majority of statistical learning paradigms uses a learning/training phase followed by a separate test phase. While this approach has been successful to show statistical learning in many different studies, it can reveal little about the learning process. However, it is challenging to measure a learning in a continuous manner, and indeed, until recently, it has only rarely been attempted. One notable exception used a self-paced presentation method, to investigate how people learn temporal regularities in the order of appearance of shapes (Karuza, Farmer, Alex, Smith, & Jaeger, 2014). This study showed that the decrement in reaction times needed to identify predictable elements in a sequence followed the learning process. Similar methods could help to investigate learning dynamics by tracking how statistical representations emerge in individual participants' behavior during learning and how these measures could predict the final learning outcome (Siegelman, Bogaerts, Christiansen, & Frost, 2017).

The above described approach (Karuza et al., 2014) is suitable if the regularities in the sensory environment are temporal, which similarly to implicit learning studies (Nissen & Bullemer, 1987), makes reaction times a suitable measure of learning. However, in more realistic scenarios, people are faced with environmental regularities that are more complex than simple temporal order, since most visual stimuli also have spatial regularities. People are well-known to be sensitive to such regularities,

even for novel abstract shapes (Fiser & Aslin, 2001). While the original finding has been reported relatively long time ago, to date, little is known about the mechanisms of spatial statistical learning.

An intriguing possibility is to use eye-movements during learning, which as a continuous measure could reveal some of the mechanisms. Yet, to our knowledge, no study has analyzed eye-movements during spatial statistical learning. This is surprising, as eye-movements are a widely used measure of attention and memory processes. It has been shown that eye-movements can be sensitive to memories that are not yet available for verbal report (Hannula, 2010; Hollingworth, Williams, & Henderson, 2001), thus it is feasible that looking patterns during spatial statistical learning could also indicate implicit learning processes. Other papers have found that eye-movements reflect learning effects only when the relevant memory trace is already explicitly reportable (Smith & Squire, 2008). This suggests that eye-movements could indicate the emergence of explicit knowledge during spatial statistical learning, but implicit representations might not have an effect on the eye-movements.

A potential difficulty with investigating statistical learning through eye-movements arises from the fact that people gain a lot of information from the visual periphery, making the link between eye-movements and processed visual information non-trivial. In an experimental set-up, it is possible to make the link between processed visual input and eye-movements tighter, by using a gaze-contingent presentation method, where little or no information is presented on the visual periphery, with stimuli appearing continuously wherever people focus their gaze. Such a manipulation makes eye-movements a measure of information gathering, and has been successfully used to study how people search for useful visual information during classifying noisy patterns (Yang, Lengyel, et al., 2016).

Following the exposition above, the goal of the current chapter is to investigate two questions. First, to establish whether a gaze contingent visual exploration paradigm is sensitive enough to measure the explicit and implicit learning of abstract spatial regularities in a unified manner. Second, to clarify whether human implicit and explicit learning of statistical regularities proceeds in a similar manner.

We investigated these issues in a set of three spatial statistical learning experiments using eye-tracking. In the first experiment, we introduced a novel gaze-contingent active spatial statistical learning method within an explicit learning paradigm, and quantified eye movement signatures during explicit statistical learning. In the second experiment, we investigated eye movement patterns on the same statistical structures while making the task implicit, therefore testing whether the paradigm is suitable to gain insights into the mechanisms of both explicit and implicit statistical learning. In the third experiment, which differed from the second one only in length, we examined whether extended exposure in an implicit setup would make eye movement behavior converge to that observed during explicit learning.

General Methods

Stimuli and Structure

The experiment was created in PsychoPy, on a Windows 7 PC with a 27" screen, with a resolution of 1600*900 and refresh rate of 60Hz. A set of twelve abstract shapes was randomly divided into 6 pairs (**Fig 2.1 A**). The shapes within a pair had a fixed spatial orientation throughout the experiment: when one of the shapes in the pair were present in a scene, the other was always present, too, and the spatial relation between the two shapes was identical throughout the entire experiment. Two pairs were arranged horizontally, two vertically, one pair had a diagonal up and one diagonal down orientation. From the six pairs, scenes were created (Such as on **Fig 2.1B**), each containing 3 pairs on the 3 by 3 presentation grid. All possible scenes were created, with the constraint that each scene consisted of 1 horizontal, 1 vertical and 1 diagonal pair. This constraint results in 144 possible unique




Figure 2.1. Stimuli and Procedure **A-B**). the paradigm of Fiser & Aslin (2001). **A**) 12 shapes are randomly arranged into 6 pairs: 2-vertical, 2-horizontal and 2-diagonal. **B**) One possible arrangement of 3 pairs on the presentation grid (144 possible arrangements) **C**) An example of what participants see in our paradigm. In this example the observer looked from the bottom-middle to the bottom-left cell. If gaze was in the mid-region of the cell, that contained the image, the shape appeared, and remained visible at full contrast until gaze was in the cell. The shape in the previously visited cell gradually faded out over the course of 1.5 sec. Participants had 6 seconds to explore each scene.

scenes, identically to a previous study on visual spatial statistical learning (Fiser & Aslin, 2001). The probability of each cell containing a shape overall was $2/3$, and each shape was present in half of the scenes. The presentation grid had a black frame and had a size of 810*810 pixels (~28.4 deg of visual angel), meaning each cell spanned approx. 9.62 visual angels vertically and horizontally. Images were presented within the central region of each cell spanning an area of 5.7*5.7 visual angles (as on **Fig 2.1C**).

Procedure

The experiments were conducted in a dimly lit and sound attenuated room. A Tobii EyeX 60Hz eye-tracker was calibrated using a seven-point calibration, from a viewing distance of 60cm. After calibration, participants completed ten 6-second-long practice trials. On each practice trial, 6 images were randomly selected from of a set of 12 images of dogs. The images were arranged at random locations inside the 3*3 grid and were revealed in a gaze contingent manner: the content of the cell was visible only when the location of the observer's gaze was at the central region of the cell, otherwise the given cell was shown empty. Specifically, the content of a cell was revealed only if two subsequent eye position samples (taken approx. 15 ms apart) were within the central gaze contingent region of a cell (5.7*5.7 visual angels). The goal of the practice trials was to familiarize participants with the method of using their gaze to reveal images in the grid. After the practice trials, calibration of the eye-tracker was double checked, and recalibrated if necessary, before the start of the learning phase. The trials in the learning phase were also 6-second-long, following the same the gaze contingent rule as during practice. The experiment had 144 unique trials that were presented in a randomized order once in Exp 1,2 and twice in Exp 3.

Each trial started with an empty grid and a fixation point where the observers had to fixate to initiate the trial. The position of the fixation cross was pseudorandom-uniformly distributed, appearing at the middle of each cell of the 3 x 3 grid equal number of times across the experiment (16-times in Exps 1 ,2, 32-times in Exp 3). Unlike previous spatial statistical learning studies, the full scenes were never visible at once. Instead, individual shapes were revealed in a gaze-contingent manner, when

the participants' gaze was at the mid-region of cell. After participants looked at a cell containing the shape, the shape was visible at full contrast until gaze was in the same cell, but gradually faded away becoming invisible in ~1.5 sec, if the participant looked away to another cell. The shapes did not start to fade away until the gaze was within the outer region of the same cell, but already not in the gaze-contingent central region. If a participant visited three cells over a short period of time, after arriving to the third cell, the first shape abruptly disappeared, and the second started to fade out. Therefore, maximum two shapes of the grid were displayed at any given time and only one at full contrast. If the observer's gaze was in the mid-region of a cell, which did not contain a shape in the given trial, a gray square (same size as the shape images) was revealed to show that the cell is empty to reduce the observer's uncertainty whether s/he managed to fixate the cell. These gray cells remained visible until the trial was over, thereby ensuring that the end of each trial was easily noticeable. Participants were free to visit or revisit with their gaze any of the cells during the trial. When the trial was over after 6 seconds, all images disappeared, and after a 500ms inter-trial-interval the next fixation-cross appeared at one of the cells to initiate the next trial.

At the end of the learning phase, after a short break, a two-interval forced choice (2-IFC) test session followed. Before the test, participants were instructed to select the more familiar pair based on what they have seen during the learning phase, and to concentrate on the combinations and not solely on the individual shapes. For the test, 6 foil pairs (which did not appear in the same arrangement during learning) were created from the original shapes and were tested against each of the real pairs, resulting in 36 test trials, which were presented in a random order. The order of the real pair versus foil intervals on each trial was pseudo-randomly controlled: half of the trials started with a true pair, the other half with the foil. On each trial, participants had to select which pair was more familiar using the left/right arrow key for the 1st/2nd pairs respectively.

Data Analysis & Measures

All data was analyzed in Python, statistics were calculated using the SciPy and the *scikit-learn* libraries. As the experimental set-up was gaze controlled using the central areas of the cells of the

presentation grid, eye-movement data was analyzed in a discretized manner, based on whether eye-movement samples were within the gaze contingent region of one of the cells or not, since where exactly gaze fell within this region of interest had no functional consequence. Looking into the non-contingent outer regions of the cells had no influence on stimulus presentation. On average participants made more than 7 (7.2 ± 1^1) transitions per trial, adding up to more than 1000 transition events over the course of the experiments. From these transition events, we selected the ones that are potentially learning related, in a way detailed below. Since the number of transitions could also change over time, we were interested in proportions and not the absolute number of the events.

We separated the eye-movement transition data into two different measures, because they could indicate different behaviors: explorative looks and returns. Explorative looks were defined as transitions to cells for the first time on a trial. Returns were defined as transitions to cells that had already been visited before on the current trial. The difference between these events is important, since in case of returns, the participant could be more certain what s/he would see at a given location, since s/he had already seen it in the last couple of seconds. In case of explorative looks, no such information was available, the content was predictable only for shape pairs, and only if the participant had already learned about the spatial relationships between shapes.

Within explorative looks, we wanted to separate transitions from shapes that could be indicative of statistical learning. Starting from a cell containing Shape₁ there are three transition possibilities:

X₁. : Transition to another shape that is the pair of Shape₁

X₂. : Transition to another shape (that is not the pair of Shape₁)

X₃. : Transition to an empty cell

We defined our measure as Pair Exploration Ratio: $X_1/(X_1 + X_2 + X_3)$.

¹ From here on, in the present thesis this format represents Mean \pm SD

From returns, we were only interested in events, when the gaze returns to a shape that had already been visited within the ongoing trial. We separated three such possible events:

Y_1 : return to a shape directly from the other shape of the pair

Y_2 : return to a shape from to another shape (that is not the pair)

Y_3 : return to a shape from an empty cell

We defined our measure as Pair Return Ratio: $Y_1/(Y_1 + Y_2 + Y_3)$.

We calculated both of these measures trial by trial for each participant. For analyzing and visualizing temporal changes, we split the data into consecutive equal-length bins of 36 trials each. The above defined measures do not have a trivial chance level, as the probability of transitioning to a pair depends on the number of neighbors of the currently fixated cell, and also on the typical behavior of the participant. In order to obtain a chance-level we kept the exploration data as it was, and we randomly shuffled the order of the presented stimuli 100 times for each participant. We calculated our measures on each shuffled combination of exploration data and stimuli and averaged over shuffles and participants to obtain an overall chance. The advantage of this method is that effects of stimulus independent temporal patterns (eg.: exploring more cell over-time) in the exploration data are preserved in the chance measure. Since this is not an undisputable measure of chance, as there are similarities between the scenes which we shuffle, we do not base any of our main conclusions on this measure, but we include it as a baseline on the figures below.

Since there are several different eye-movement measures that are partially correlated and predictive of familiarity test performance, we used cross-validated Lasso regression (Tibshirani, 1996) to select the relevant predictors and account for over-fitting. Lasso regularizes regression weights by a parameter λ times the absolute value of the predictor. We selected the value of λ by cross-validation, resulting in decreased regression weights. Lasso is useful for feature selection from correlated predictors as unlike Least squares or Ridge regression, predictors that cannot predict the hold-out sample will often have zero weights assigned (Tibshirani, 1996) .

For the correlational measures, apart from the Pearson correlation, we calculated the exact p values (randomization test), by randomly permuting the data (shuffling the X and Y pairing) 5000 times, calculating the r value for each permuted sample, and looking at where the obtained r value falls within the permuted distribution (two-tailed). The advantage of this method is the weaker sensitivity to outliers and non-normality in the data (Bishara & Hittner, 2012). In general, with this exact test we obtained p -values very similar to the results of Pearson correlation, always supporting the same conclusions.

To obtain Bayes Factors (BF) for paired and between group t -test, we used the BayesFactor package (Rouder, Speckman, Sun, Morey, & Iverson, 2009) with a non-informative Jeffrey-Zellner-Siow Prior on possible effects sizes. To calculate Bayes Factors (BF) for correlations we used the JASP package (Wagenmakers et al., 2018) with a two-tailed test and again a uniform Jeffrey-Zellner-Siow prior. By convention, Bayes factors below $1/3$ provide evidence for the null, in the range $= [1/3, 3]$ insensitive, above 3 evidence for the alternative.

For visualizing and interpreting our results, we separated participants into three groups based on performance on the familiarity test in the following way: Low Learners $\leq 58.3\%$ (21 correct out of 36 test trials) < Medium Learners < 86.1% (31/36) \leq High Learners. Our main conclusions are based on relationships between visual exploration data and test performance on the entire data-set, therefore are not affected by this grouping, which we used for demonstration purposes.

Computational Modeling

To obtain a measure that can be fitted to all gaze transitions, thereby not relying on the selection of certain events, and could determine the extent to which the exploration data of each participant is influenced by the pair structure, we developed a one-parameter computational model (M_1). We compare this model to a random null model (M_0). Since there are three type of statistical regularities in the stimuli (horizontal, vertical, and diagonal), we also developed a 3-parameter extension of M_1 :

Model 2 (M_2), which is sensitive to direction specific influences. Below we describe the models in detail.

Model Description

Null Model

For the null model, we used the individual empirical transition probability matrix (See **Fig 2.2A**) over the course of the whole experiment for each participant separately to predict their gaze transitions.

Eg: for a transition from Cell 1 to Cell2:

$$p(\text{Cell2}|\text{Cell1})_{\text{null}} = p(\text{Cell2}|\text{Cell1})_{\text{empirical}}$$

To test how well the null model can predict the data, we calculated the natural logarithm-based likelihood of each transition, given the empirical transition probability matrix. We only included trials that had at least two transition events ($\approx 97\%$ of Trials). This model had no free parameters and, therefore, it represented how well the average behavior of each participant could predict his/her own single trial exploration data.

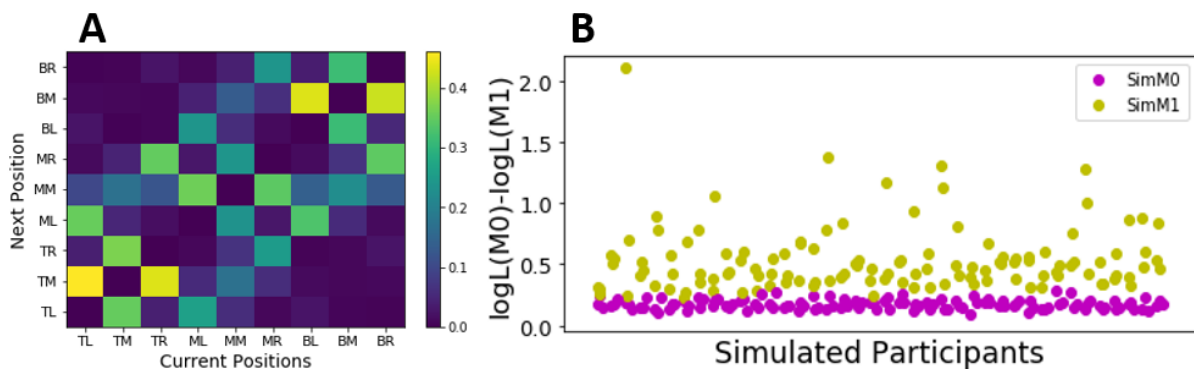


Figure 2.2. Transition P-s for Null Model, Model Simulation. **A)** Average Transition Probability Matrix of Experiment 1: average probability of transitions between the 9 cells of the grid, from cells on x-axis to cells on y axis. T: Top, M: Middle, B: Bottom, R: Right, L: Left, as an example: we can see that from the top-left cell (TL, first column) participants most often switch to the top-middle cell (TM). **B)** Example Simulation: One run of 120*2 simulated participants, each participant was simulated once with M_0 that is based on the individual Transition Probability Matrix and once with M_1 pair influence model, both models were fitted to both simulated data-sets trial by trial. The average per trial negative log likelihood difference between the fit of M_0 and M_1 is plotted for the two data set. As expected M_1 always fits better (all dots are above zero). When the data was simulated with M_1 , however, the benefit of fitting M_1 is reliably greater, showing that these models can be separated participant-wise. The variance in the advantage of M_1 in the simulated data-set can be explained by the fact that the empirically fitted alpha values were used for the simulation, varying across participants and trials.

Model 1: Overall Statistical Influence Model

The M1 learning model used the same empirical transition probability matrix as the null model, but was extended with a single parameter: α . This parameter represents pair-structure-bias, that is an increased probability of transitioning from each shape to its pair. For example, if Cell 1 contained Shape5, and Cell 2 contained Shape7, and the two shapes formed a pair, the likelihood of transition was increased by α .

Eg: If shapes in Cell 1 and Cell2 belonged to the same pair:

$$p(\text{Cell2}|\text{Cell1}) = p(\text{Cell2}|\text{Cell1})_{\text{null}} + \alpha$$

Transitions probability from Cell1 to other cells remained as before:

$$p(\text{Cell3}|\text{Cell1}) = p(\text{Cell3}|\text{Cell1})_{\text{null}}$$

Transition probabilities from each cell that contained a shape on a given trial were updated with α and then normalized. Transition probabilities from empty cells remained unchanged. We fitted the value of α trial by trial, by minimizing the negative log likelihood over the transitions observed on any given trial. To fit the model, we used the *fminbound* function of *scipy* with limiting the value of α in the 0-1 range². The goal of this model was to parametrize the extent to which participants look more within pairs, than what would be expected from average behavior.

Alternatives: to confirm that the choice of null model is not responsible for our results, we fitted M₁ to a directed random walk null model yielding essentially the same results for the pair influence parameter α ³.

² In the Appendix Fig.A.3 we consider an alternative by letting Alpha vary on a broader range, finding support for the same conclusions.

³ Result of fitting M₁ to a different M₀ can be found in Appendix Fig A.1 & A.2.

Validating approach with Simulation

To validate that we can separate the two models (M_0 & M_1), when the true underlying model is known, we conducted a simulation-based analysis. We generated data with both models using the empirical transition probability distribution of each participant from Experiments 1-3. For each trial, we simulated as many transitions as the participant had performed on that given trial. For simulating data with M_1 , for each simulated participant and trial, we used the empirically fitted α value of any given trial. We calculated the log-likelihood for both simulated data-sets under the M_0 (Null model does not have to be fitted). We fitted M_1 to both data-sets and calculated the log-likelihood of the fit. Since M_1 includes M_0 (if $\alpha=0$, $M_1=M_0$), M_1 always fits the data better, but the extent to which M_1 fits better is highly dependent on whether it is the true underlying model (See **Fig 2.2B**). To perform model selection on whether M_1 is better than M_0 we used the likelihood ratio test, with the number free parameters as the number of fitted non-zero alphas ($\alpha>0.01$) over all trials. The likelihood ratio test's output is the p -value of the obtained likelihood difference under the null-model (with an assumption that nested models likelihood differences follow a X^2 distribution, with degrees of freedom as the number of parameters M_1-M_0) (Huelsenbeck & Crandall, 1997). By using a threshold of $p < 0.001$, we identified whether each simulated participant is better fitted by M_0 or M_1 . After repeating this process ten times for all $120*2$ participants, the likelihood ratio test identified virtually all (119.9 ± 0.3 out of 120) null model simulated participants as better fitted by the null model, and more than 90% of M_1 (108.6 ± 2.1 out of 120) simulated participants as better fitted by M_1 . The reason for sometimes not recovering the generating M_1 model for a few simulated participants was the often-low alpha values used in the simulation (since we used the empirically fitted values). Repeating the process by simulating the data with alpha values randomly sampled from the fitted values of high learners achieved a perfect recovery ($120/120$) of M_1 as the better model for this simulated data-set. This success in separating M_0 and M_1 validates our approach and shows that we can draw participant-wise conclusions on whether their overall visual exploration data has been influenced by the pair structure of the stimuli.

Model 2: Specific Statistical Influence Model

In Model 1 statistical influence is represented by a single parameter α , representing the overall effect of pair structure. Since on each trial there are three pairs with different orientations, it is likely that some observers use only some of this information, or that a participant would learn about these aspects of the stimuli during different parts of the experiment. In order to analyze the influence of the different orientations separately, we extended our M_1 into a three-parameter model (M_2).

Cell1	Cell2	Cell3
Cell4	Cell5	Cell6
Cell7	Cell8	Cell9

Table 2.1.: Numbering of cells in the 3*3 Grid

In M_2 , transition probabilities between cells were updated in the way described below:

If on a given trial, shapes in Cell 1 and Cell2 (Cell numbers as in **Table 2.1.**) were part of the same *horizontal* pair:

$$p(\text{Cell2} | \text{Cell1}) = p(\text{Cell2} | \text{Cell1})_{\text{null}} + \alpha_1$$

If on a given trial shapes in Cell 1 and Cell4 were part of the same *vertical* pair:

$$p(\text{Cell4} | \text{Cell1}) = p(\text{Cell4} | \text{Cell1})_{\text{null}} + \alpha_2$$

If on a given trial shapes in Cell 1 and Cell5 were part of the same *diagonal* pair:

$$p(\text{Cell5} | \text{Cell1}) = p(\text{Cell5} | \text{Cell1})_{\text{null}} + \alpha_3$$

Transitions probabilities from Cell1 to other cells remained as before:

$$p(\text{Cell3} | \text{Cell1}) = p(\text{Cell3} | \text{Cell1})_{\text{null}}.$$

Transitions probabilities from all cells that contained shapes were updated with α_{1-3} on each trial and renormalized. Transition probabilities from empty cells remained unchanged. The values of α_{1-3} were fitted trial by trial by using the *minimize* function of *scipy* by minimizing the negative $\log(\text{Likelihood})$

over the gaze transitions of each trial. We only used trials with at least three transition events ($\approx 90\%$ of trials). To test the predictive power of this analysis we separated the 36 test trials based on the orientation of the presented true pair. This separation yields 12-12-12 horizontal/vertical/diagonal test trials, where the performance could be linked to the three parameters.

Experiment 1: Explicit Active Statistical Learning

Introduction

To gain insight about the process of spatial statistical learning, we combined visual statistical learning with a novel gaze contingent presentation method. Experiment 1 had three main goals: 1) To test if people can learn the pair structure during visual statistical learning if the underlying generative structure of the scenes is explicitly defined for them but they can never see the full scenes at once, only parts of it presented in a gaze-contingent manner. 2) To investigate whether the (known or learned) pair structure influences visual exploration patterns. 3) To test if individual learning outcome of the subsequent familiarity test could be reliably predicted from visual exploration patterns.

Methods

40 students (age: 25.5 \pm 4.6 years, 13 male, 35 Right Handed) participated in the study, after giving written informed consent. They were recruited via a Hungarian student organization and received ~ 1500 HUF or food vouchers as compensation. No participants were excluded. We chose a relatively large sample size, as we expected considerable individual variability in learning outcome.

Before the start of the learning phase, participants were told that their task was to pay attention and find several pairs of shapes in the scenes about which they will be questioned afterwards. They were also told that the pairs were defined so that their constituting shapes always appeared next to each other arranged either vertically, horizontally or diagonally. They were not told about other details on

how the scenes are constructed (such as the number of pairs or that each scene was constructed from one vertical, one horizontal and one diagonal pair). The learning phase consisted of the 144 unique gaze-contingently presented scenes followed by 36 familiarity test trials⁴.

Results

Familiarity Test Performance

Mean performance in the 36 2-IFC familiarity test trial was 70.56 +/-18 %, clearly above chance ($t_{39}=7.09$, $p<.0001$) (See **Fig 2.3A**). The distribution was not normal (Shapiro-Wills Test: .916, $p=.006$), suggesting that participants could be meaningfully separated into sub-groups based on their performance.

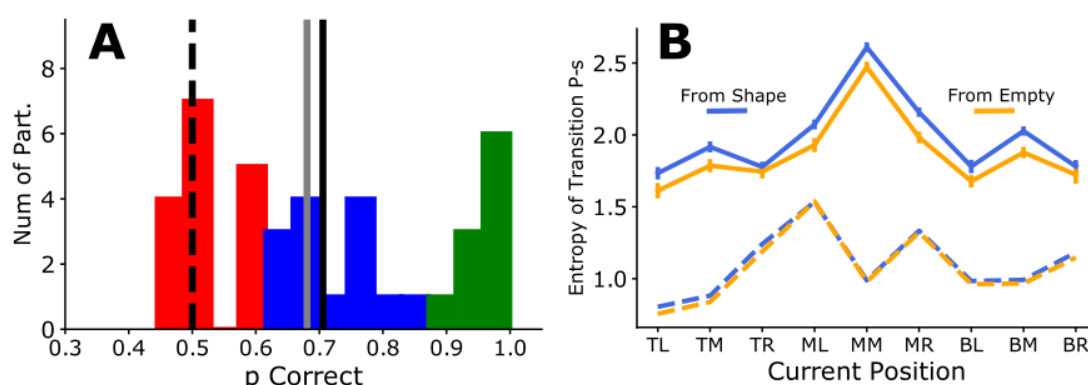


Figure 2.3. Test performance and transition predictability in Exp 1. **A)** Familiarity Test performance distribution in Exp 1. Mean/Median Performance: solid vertical black/gray lines, Chance: Dashed black line. Colors separate participants by familiarity test performance, with red, blue, green, as low medium and high learners respectively. **B)** Entropy of transition probability distributions shows that transitions from each cell were more unpredictable if the cell contained an object as opposed to when it was empty, excluding the possibility of a visual information independent scanning strategy. X-axis: The 9 cells of the presentation grid (eg: TL: top left, BR: bottom right) , the solid lines are the log₂ based entropy of transition probability distribution for each cell. The two lines separated based on whether the cell was empty or contained a shape. The dashed lines are predictions from a simulated scanner (description in main text). Transitions from the center (MM) are the most unpredictable, which can be explained by the highest number of adjacent cells. Errorbars=SEM

Visual Exploration Behavior Descriptives

The descriptive statistics provided in this section are the mean +/- SD across participants over the per trial average of each participant. These measures do not address our main research questions, but

⁴For the details of the stimuli and procedure, please see the General Methods above

since we introduce a novel gaze-contingent method under which participants were free to decide how they explore the scenes, these measures provide an overview of participants' behavior in this paradigm.

Number of visited cells and transitions

Participants visited an average 6.93 ± 1.5 of the 9 cells per trial (77% of cells in the grid). Since some of the cells were empty, the number of shapes seen was necessarily lower: 4.66 ± 0.99 of the 6 shapes, but the ratio was the same (77% of the shapes in the scene). To test if the number of the visited cells changed over time, we fitted a linear regression with trial number as a predictor and number of visited cells as a dependent variable. The average slope in this analysis was greater than zero ($t_{39}=2.8103$, $p=.0077$) showing that observers visited more cells over-time. The number of transitions did not directly follow from the number of visited cells as participants could return to a previously visited cell during a trial, which would, therefore, increase the number of transitions, but not the number of visited cells. The average number of transitions made between cells in a trial was 8.12 ± 1.8 , adding up to a total of 1169 ± 259 transition per participant over the course of the whole experiment. The vast majority ($95.1 \pm 2.2\%$) of transitions were made to adjacent cells. The most common transition to adjacent cells was horizontal (3.93 ± 1.0 transition per trial), followed by vertical (2.77 ± 0.82 per trial) and diagonal (1.01 ± 0.38 per trial).

Looking Times

Average looking time at a single visit to gaze-contingent mid-region of a cell was 531 ± 142 ms. Looking time were becoming shorter over time as shown by the slope which was significantly below zero ($t_{39}=3.6001$, $p=.0009$). Notice, that this looking time does not necessarily reflect a single fixation, only the continuous time spent in the central region of a given cell before moving to a next location that could include a single fixation or a main fixation accompanied by several corrective fixations or even multiple fixations at different parts of same shape. $68.56 \pm 9.1\%$ of the overall looking duration was spent inside the gaze-contingent mid-regions, despite it only covered $1/3^{\text{rd}}$ of the area

of the grid. The amount of time spent in the outer regions of the cells can be explained by the fact that the shapes, once uncovered, did not disappear while the gaze was in the outer region of the cell.

Predictability of Transitions

One possible strategy in a gaze-contingent setup is that participants scan the grid in some regular or random manner regardless of what they see inside the cells, and since shapes constituting pairs are visible more often next to each other in these scenes, participants passively learn this co-occurrence information. To exclude this possibility of visual information-blind scanning, we separated the transition probability distribution from each cell depending on whether the cell, from where the transition was initiated was empty or contained a shape (Example of Transition P distribution on **Fig 2.2A**). Next, we calculated the \log_2 -based entropy of these transition distributions to quantify how spread out each distribution was. As expected, entropy was the highest in the middle cell of the grid, as there are more possible adjacent cells to transition to (**Fig 2.3B**). More importantly, the entropy of transitions was higher if the cell contained a shape, excluding a visual-input-independent scanning strategy ($t_{39} = 6.4809, p < .0001$). As a baseline for this measure, we calculated the same measure for simulated scanner model. The scanner model always searches horizontally first, then vertically, with left to right, and top to bottom preference, furthermore it keeps track of the visited locations within a trial, and only returns to previously visited cells if all adjacent cells have already been visited on that trial. The scanner model was simulated for each participant on each trial, to visit the same number of cells as the participant did on that trials. The results of this comparison showed that participants behavior was far more unpredictable than simple scanning as can be seen by the difference between the simulated (dashed) and solid (empirical) lines on **Fig 2.3B**.

Statistical structure learning measures

Pair Exploration Rate

We have found that the proportion of explorative transitions (looking at a shape for the first time in a given trial^{5,6}) made within pairs (i.e. so that the previous look in the trial was at the corresponding other shape of a pair) was correlated with learning ($r_{38}=0.3907$, $p=.0127$, $p_{\text{exact}}=.0118^7$, $\text{BF}=3.947$). Furthermore, the slope of this measure was also correlated with learning performance ($r_{38}=.5731$, $p=.0001$, $p_{\text{exact}}=0.0$), showing that the more pair-exploration rate increased, the better participants learned the pairs. To confirm the temporally emerging pattern of this measure (**Fig 2.4A-B**), we analyzed it in temporal bins, finding that within pair exploration was not predictive of learning in the first half, but was highly predictive in the second half of the experiment (Quarter 1: $r_{38}=-0.105$, $p=.5179$, $p_{\text{exact}}=.5180$, $\text{BF}=0.24$; Q2: $r_{38}=.242$, $p=.1332$, $p_{\text{exact}}=.1358$, $\text{BF}=0.59$; Q3 $r_{38}=.456$, $p=.0031$, $p_{\text{exact}}=.0022$, $\text{BF}=13.36$; Q4: $r_{38}=.498$, $p=.0011$, $p_{\text{exact}}=.0006$, $\text{BF}=33.8$, Bonferroni corrected $p_{\text{crit}}=.0125$).

⁵ for definition see General Methods

⁶ For descriptive stats of these measures see Appendix Table A.1.

⁷ Hereby on r denotes the Pearson correlation coefficient, with the p value, p_{exact} denotes the exact p value obtained by permutation

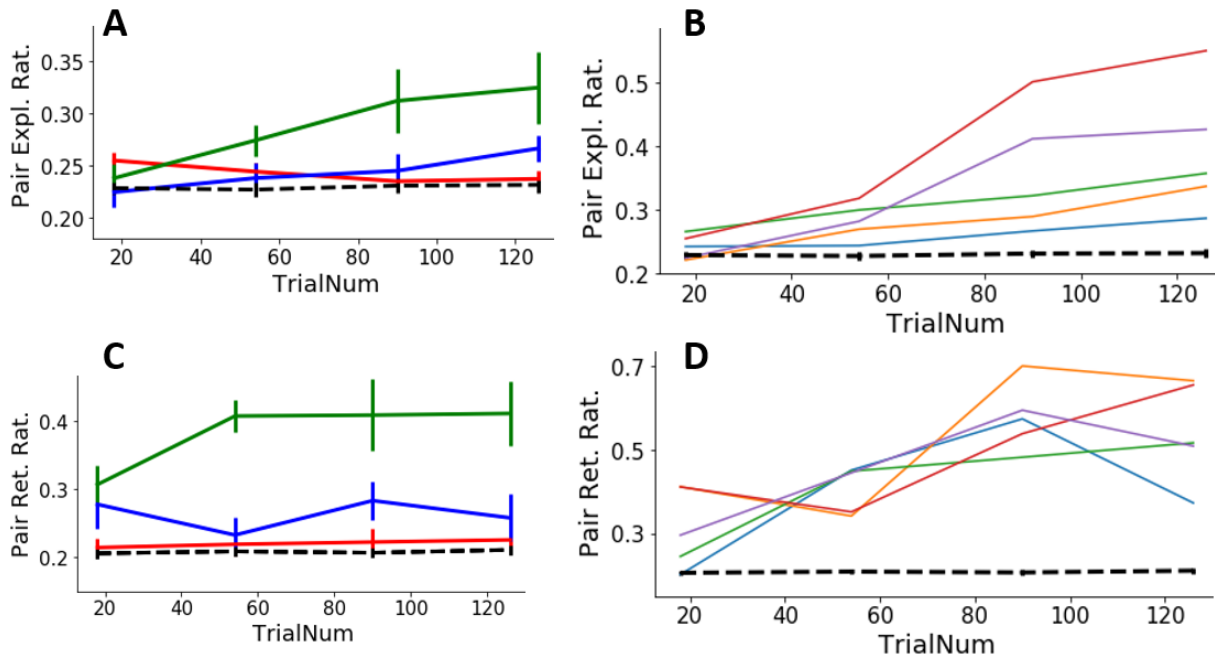


Figure 2.4.: Pair structure influence on eye-movement measures in Exp 1. Both measures show a temporally emerging influence of pair structure that is predictive of familiarity test performance. **A-B)** Pair exploration rate **A)** Participants grouped by test performance (red-blue-green as low -medium-high learners). Dashed line represents shuffled chance. Each bin contains 36 consecutive trials. **B)** Pair Exploration Rate of five Example individual high learners. **C-D)** Pair Return rate **C)** Participants grouped by learning performance. **D)** Example high learners (same participants as in B). While pair exploration rate tends to increase until the end for most high learners, return rate often drops at the end, suggesting that the pairs had already been confidently acquired. Error bars: SEM.

Confirmatory Looks: Pair Return Rate

To assess learning further, we calculated a different measure, the proportion of returns made within pairs⁸. Pair Return Rate was highly predictive of learning performance at the familiarity test ($r_{38}=.7014$, $p<.0001$, $p_{\text{exact}}=0$, $\text{BF}=41036$) (**Fig 2.4 C-D**). Unlike in the case of the previous measure, the slope of within-pair returns was not predictive of learning outcome ($r_{38}=.2852$, $p=.0745$, $p_{\text{exact}}=.0748$) because for some participants, pair return rate increased fast in the middle part of the experiment and dropped-off toward the end (see **Fig 2.4B** vs. **2.4D**). The reason for this pattern could be that after learning the pairs, it is not necessary to do confirmatory returns. Confirming this, looking at the correlation separately in time-bins, we have found that the within-pair returns were predictive of learning outcome early in the experiment (strongest in the second quarter), and lost some predicate power by the end (Correlation with learning in temporal quarters Q1: $r_{38}=.49$, $p=.0013$, $p_{\text{exact}}=0.001$,

⁸ for definition see General Methods

BF=28.31; Q2: $r_{38}=.6155$ $p<.0001$, $p_{\text{exact}}=0$, BF=1082.99; Q3: $r_{38}=.5827$, $p=.0001$, $p_{\text{exact}}=0$, BF=357.76; Q4: $r_{38}=.5601$, $p=.0002$, $p_{\text{exact}}=0$, BF=178.77; Corrected $p_{\text{crit}}=.0125$).

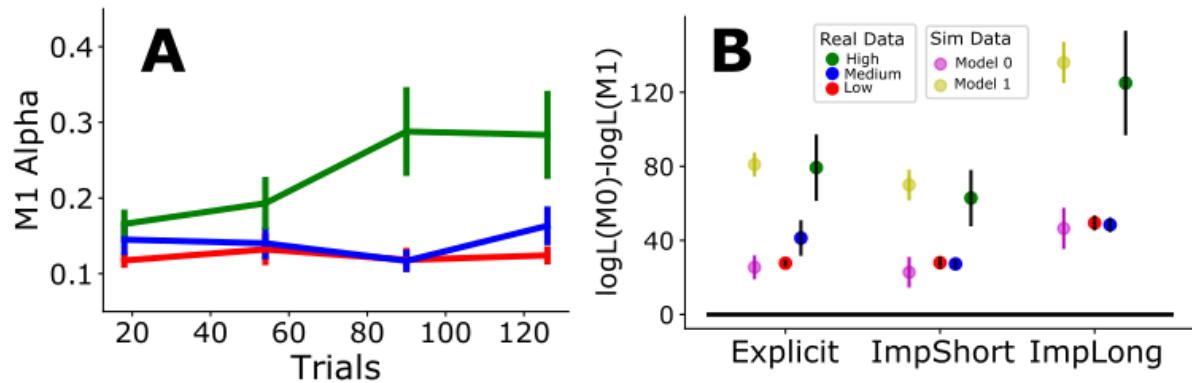


Figure 2.5 Model based assessment of pair-structure influence on visual exploration. **A)** Parameter alpha over time grouped by learning performance (color code as before) in Exp 1. We can see that visual exploration data of high learners over-time shows influence by the pair structure. **B)** Model fit: Advantage of fitting M₁ for Exps 1-3 on simulated and real data. Real data is separated by learning performance as before, simulated data is shown for 2*40 simulated participants for each experiment (40 with no pair influence M₀ simulation and 40 pair influence M₁ simulation). High learners are better fitted by M₁ to an extent very similar to M₁ simulated data, confirming a true pair influence. The benefit in log(L) for low learners is very similar to M₀ simulated data, showing that it is a consequence of fitting a more complex model and does not represent true influence by the statistical structure. (Y axis is the difference in negative log likelihood under M₀ and under M₁. Implicit Long data has larger values because -log(L) is summed across trials).

Model Based Analysis of Overall Statistical Influence

We have found that the α value of M₁, representing increased overall looking within pairs, could predict familiarity test performance ($r_{38}=.4865$, $p=.0015$, $p_{\text{exact}}=.0008$, BF=26.1). This effect emerged over time (**Fig 2.5A**), since there was no predictive power of M₁ α at the beginning first half of the experiment, and strong links in the second half end (Correlation with learning in temporal quarters Q1: $r_{38}=.25$, $p=.1247$, $p_{\text{exact}}=.1264$, BF=0.62; Q2: $r_{38}=.23$, $p=.1545$, $p_{\text{exact}}=.1492$ BF=0.52; Q3: $r_{38}=.53$, $p=.0004$, $p_{\text{exact}}=.0002$ BF=84.18; Q4: $r_{38}=.48$, $p=.0015$ $p_{\text{exact}}=.0006$ BF=24.91; Corrected $\alpha=.0125$). As **Fig 2.5B** shows the advantage of fitting M₁ was very similar in high learners (green) to data which was simulated by using the pair influence model (yellow). We used the Likelihood Ratio test, to decide participant-wise whether M₁ fitted the data better than M₀ (**Fig 2.5B**). We have found that 10

participants were better fitted by M_1 ⁹. Out of these 10 participants, 7 were high learners (7/11 high learners) and 3 were medium learners (3/10), further confirming that better learners are more influenced by statistical structure (See **Fig 2.5B**), while none of low-learners' (N=19) exploration patterns showed an influence of the stimulus statistics according to the model selection¹⁰.

↓Exp	Predictor:	β N Cells	β N Returns	β PairExp. R.	β Pair Ret. R.	β $M_1 \alpha$
Exp1 Explicit						
Lasso ($\lambda=0.093$)		-0.	0	-0.	0.6087	0.
Least Square		0.0804	-0.0515	-0.2119	0.7545	0.1497
Exp2 Implicit Short						
Lasso ($\lambda=0.095$)		0.	0.043	0.	0.033	0.2709
Least Square		0.2546	0.0418	-0.0719	0.0592	0.5728
Exp3 Implicit Long						
Lasso ($\lambda=0.103$)		0.	0.115	0.2849	0.1952	0.0214
Least Square		0.1953	0.1537	0.2659	0.167	0.2465

Table 2.2: Combined prediction of Learning. Lasso (top in each cell) and Least Square regression (bottom) weights for predicting test performance for five measures of visual exploration for experiments 1-3. Grey marks the strongest predictor for each experiment. In column 1, λ is the regularization parameter selected by cross-validation. The first two predictors are descriptives of visual exploration: N Cells =average number of visited cells per trial, N Returns=average number of visited cells per trial. The remaining three predictors are based on pair structure: Pair Exp R. = Pair Exploration Rate as defined above, Pair Ret. R.= Pair Return Rate as defined above, $M_1 \alpha$ is the average value of the pair influence parameter.

Combined Prediction of Learning

Although the eye-movement-based measures above showed a strong correlation with learning, inevitably, they were also correlated with each other to some extent. Therefore, it is important to analyze their respective contributions. We used five potentially relevant predictors: Number of Visited Cells, Number of Returns, Pair Exploration Rate, Pair Return Rate, and the α parameter from our M_1 model, with familiarity test performance as the dependent variable. These five predictors could explain 50.1% of variance in test performance. To account for potential over-fitting and select

⁹ Figure on model selection results in Appendix Fig. A.4

the sufficient predictors, we calculated the cross-validated Lasso regression on the same dataset. This analysis revealed that after regularization, 48.34% of variance could still be explained. All that predictive power could be accounted for by Pair Return Rate ($\beta=0.6087$), other predictors had zero weights assigned (See **Table 2.2.** above) showing that they could not explain additional variance reliably. The spatial-statistics-related measure of pair return rate was a better predictor of test performance than the general descriptive measures of exploration behavior (such as number of visited cells or returns), showing that good learners did not necessarily explore more or return more often to already visited cells, but instead, used the structure of the scenes to guide their eye-movements.

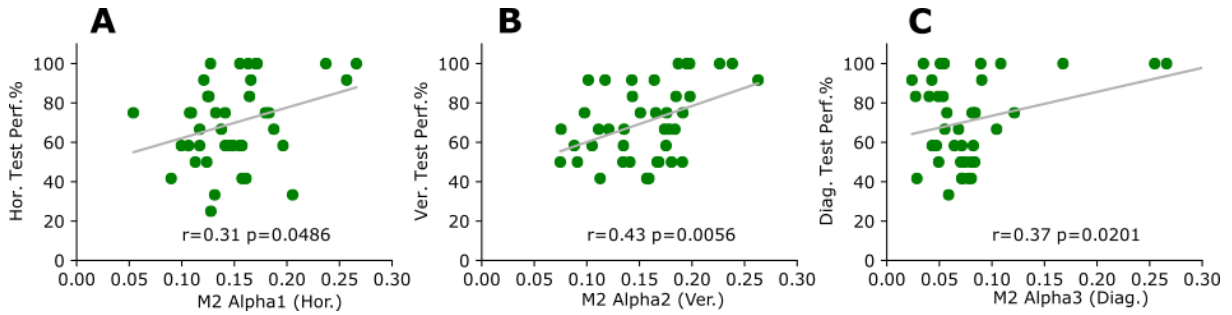


Figure 2.6. Predicting the content of learning from eye-movements in Experiment 1. This figure shows the link between pair influence and specific test performance for horizontal (A), vertical (B) and diagonal (C) pairs. This figure demonstrates, that the more eye-movements are affected by the pair structure for a given orientation (x-axis), the better they perform on test items from *that* orientation (y-axis). Dots are individual participants, with the least square line and the value of the Pearson correlation. (BayesFactor values for the three correlations: $BF_{hor}=1.33$ $BF_{ver}= 10.21$ $BF_{diag}= 3.075$)

Model Based Analysis of Specific Statistical Influence

In order to test whether eye-movements could predict what is learned in a more specific manner, we fitted M_2 with parameters α_{1-3} representing different orientations and tested whether it could predict familiarity test-performance on the three pair orientations. We found that the fitted values of α_{1-3} could predict performance in an orientation specific manner for all three orientations (**Fig 2.6**). The relationship was strongest for vertical pairs (**Fig 2.6B**), demonstrating that the more eye-movements were influenced by the vertical pair structure, the better performance was on test trials where the vertical pairs were queried. These effects were temporally emerging, as the slope of these

measures¹¹ was also predictive of direction specific learning outcome for the vertical ($r_{38} = .55$, $p = .0003$) and diagonal orientations ($r_{38} = .34$, $p = .0295$), but not for the horizontal ($r_{38} = .06$, $p = .7057$) pairs, where the pair specific influence seemed to emerge early on (see Appendix Fig A.5).

Discussion

In an active statistical learning experiment with explicit instructions, we have shown that people can learn the underlying statistical structure of the scenes, despite only being exposed to far less information at any given moment than in previous spatial statistical learning studies. We have also shown, that visual exploration patterns are influenced by statistical regularities between novel shapes. Furthermore, we have shown that this effect emerges quickly over time and is highly successful in predicting individual learning outcome. The best predictor of test performance was based on confirmatory looks, which is most likely the consequence of explicit hypothesis testing. Of course, it is unclear from these results, whether observers learned the structure of the scenes and then used their explicit knowledge for the confirmatory looks (high correlation of performance and looks is a consequence of learning) or observers used their hypothesis testing to learn the pairs by looking back-and-forth (high correlation is a consequence of the process of active learning using confirmatory looks). Most probably, these two aspects are difficult to separate. In any case, these findings demonstrate that active statistical learning based on eye movement measures is a good paradigm to track the process of learning spatial regularities in visual exploration patterns. Furthermore, the link between the orientation specific influence (Model 2) and performance on different pair orientations suggests that eye-movements can not only be used as an overall predictor of test success, but even as a specific indicator of what is learned during statistical learning.

¹¹ Linear Regression, using trial-number as predictor and the fitted values of α_{1-3} as dependent variable.

This experiment is somewhat similar to previous studies of human active information search, in a sense that participants had clear instructions regarding what the task was (look for pairs in our case), but no instructions about how to solve it. For example, previous studies have shown that children can efficiently ask questions to narrow a search space (Nelson et al., 2014), or adults can move their eyes effectively to reduce their uncertainty about a visual scene (Yang, Lengyel, et al., 2016). Our findings can be interpreted as a combination of these previous studies, since learning in our case not only requires eye-movements to collect information about scenes (as in Yang, Lengyel, et al., 2016), but also integrating information across trials (as in Nelson et al., 2014). Our task requires participants to integrate spatial relations between novel shapes both within and across trials, therefore, the statistical information can only be acquired across many trials. Hence the present results demonstrated that people use stimulus statistics during visual exploration in a situation that is more complex than in previous related work exploring human active information search.

Experiment 2 and 3.: Implicit Active Statistical Learning

Introduction

In Exp 1. above, we demonstrated that our paradigm can track learning of arbitrary spatial relationships between novel shapes. Earlier evidence that people are capable of visual active sensing came from an experiment, where observers could guide their visual search by integrating information on a given trial to perform an explicitly defined discrimination judgement in a well-defined task (Yang, Lengyel, et al., 2016). In contrast, visual exploration in our experiment was influenced by abstract spatial relationships that could only be acquired across many trials. Nevertheless, this effect still emerged due to a relatively well-defined task, therefore, it could not reveal much about the mechanisms of classical statistical learning, which is usually considered to be an implicit and automatic mechanism. Even though Experiment 1 used the typical stimulus set of

spatial statistical learning, it is possible that observers used explicit top-down mechanisms to learn that were different from the implicit mechanism required for discover the underlying structure of spatial regularities without external instruction. To address this possibility, we conducted Experiments 2 and 3 similarly to previous statistical learning experiments by not disclosing the regularities before the start of the experiment. Finding effects similar to those in Exp 1 without explicit instructions would confirm that people use statistical properties of visual scenes to guide exploration automatically, without a well-defined task. We expected such an agreement between explicit and implicit learning based on theories of active learning and curiosity reporting that people are inherently curious (Kang et al., 2009), and active learners (Gottlieb, Hayhoe, Hikosaka, & Rangel, 2014), therefore, presumably they search for regularities in the input, even if they are not told to do so.

We ran two experiments to investigate the process of active information search during implicit statistical learning and to test how learning performance and patterns of visual exploration are related. These two experiments differed only in the length of the training sequence: Exp 2 was had the same length as Exp 1, while Exp 3 was twice as long. Beyond this, these experiments had only one more difference compared to Experiment 1: the instructions did not uncover the underlying rule of the pair structures in the stimuli.

Using two different lengths allowed us to test the causal relationship between the amount of experience, the emerging internal representation of the statistical structure, and the ability of this information to guide visual exploration as measured by the change in patterns of visual fixations. Our first hypothesis was that, during the process of learning, visual search patterns would become increasingly influenced by the statistical structure of the task, and that this influence will be stronger in those participants who learned more about the structure of the stimulus stream. Our second hypothesis was that the eye-movement measures of learning we established in Experiment 1 could also track implicit learning, and the patterns of implicit learning would converge over-time to those in Experiment 1.

Methods

Participants:

80 students were recruited via a local student organization or via an online system. Participants received monetary or food voucher compensation of approximately 1500 HUF (~5 euros). The participants gave informed consent before the experiments, were naive to the purpose of the study and had never participated in other statistical learning experiments before. Half of the participants (Age: 22.1 +/- 2.8 years, 38 right handed, 13 male) completed Experiment 2 (the Short Experiment), the other half (Age: 23 +/- 5.5 years, 36 right handed, 10 male) were tested in Experiment 3 (the Long experiment). One additional participant was excluded as upon completing Exp 2 revealed not being naïve about visual statistical learning.

Procedure:

In the Short Experiment (Exp 2), stimuli and the procedure were identical to those in Experiment 1 using 144 trials of unique scenes displayed in a gaze-contingent manner for 6 seconds (see the details above). The learning phase in the Long Experiment (Exp 3) was twice as long as in Exp 1-2, with a brief break in the middle: during each half, each unique scene was presented once in a different random order. The total duration of the learning phase was approximately 16 minutes in the Short and 32 mins in the Long experiment.

The main difference from Experiment 1 was the instructions the participants received: before the learning phase, they were instructed to pay attention, explore the scenes and try to remember what they saw. They were also told that after the first part, they would have to answer questions about what they saw, but they were not told anything about pairs or possible statistical regularities. During the short break in the middle of the Long Experiment, participants were kindly asked to continue paying attention, and were reminded again that they would be quizzed after the learning phase.

Data Analysis

The same analyses were performed as the ones used in Exp 1 and detailed in the General Methods. For visualizing temporal changes, we used 36 trial long-bins, yielding 4 bins in the Short and 8 bins in the Long Experiment.

Results

Familiarity Test Performance

Learning was measured on the 36 2-IFC familiarity trials. In the Implicit Short Experiment, mean learning performance was 65.9 +/- 14.59%, which was significantly higher than chance of 50% ($t_{39}=6.8055$ $p<.0001$) (See Fig 2.7). In the Implicit Long Experiment, mean learning performance was 69.65 +/- 16.19%, which was also significantly above chance ($t_{39}=7.581$ $p<.0001$). The average performance was not significantly different between the two experiments ($t_{78}=1.0744$, $p=.2859$, $BF=.38$). The distribution of performance did not deviate from normal (Shapiro-Wills: Short = .989 $p=.959$; Long= 0.961 $p=.18$) in any of the experiments.




Figure 2.7: Familiarity Test Performance distribution in Exp 2 & 3. **A)** Exp2-Short Implicit **B)** Exp3-Long Implicit Mean Performance (Solid Black) Median performance (Solid Gray Line) is very similar, but the Long Experiment had a higher proportion of good learners, similarly to Exp 1 on Fig 2.3A. Dashed black line: chance. Colors separate participants by familiarity test performance, with red, blue, green, as low, medium and high learners

Visual Exploration Behavior Descriptives

Again, these measures do not address our main questions of interest, but since we use a novel paradigm where participants were free to choose their strategy without a well-defined task, we wanted to give an overview of behavior.

Number of visited cells

In the Implicit Short Experiment, participants visited 6.627 \pm 1.546 cells per trial out of the possible 9, while in the Implicit Long this average was 6.061 \pm 1.558. The number of visited cells /trial did not differ significantly across the two experiments ($t_{78}=1.612$, $p=.111$ BF=.71). The average number of shapes seen was necessarily lower (Implicit Short = 4.447 \pm 1.027, Implicit Long = 4.071 \pm 1.04) than the number of visited cells. We used Linear Regression for each participant separately to predict the number of visited cells over time by the trial number. The mean of the individually obtained slopes were significantly above zero for both experiments (Implicit Short $t_{39}= 2.3963$, $p=.0215$; Implicit Long $t_{39}= 3.0706$ $p=.0039$), confirming that the average number of visited cells was increasing over-time regardless whether the raw or the binned data was analyzed.

Looking Times

The mean looking time at the gaze contingent middle region of a cell in the Implicit Short experiment was 656 \pm 239 ms while in the Implicit Long it was 736 \pm 277ms. The difference in mean looking times was not significantly different between the two experiments ($t_{78}= 0.8651$, $p=.3896$, BF=.3216), but they were significantly longer than in the Explicit experiment above (Exp - Imp short: $t_{78}=2.811$, $p=.0062$, Exp-Imp Long: $t_{78}=4.108$ $p<.0001$). Looking times became shorter over time: the slope was significantly below zero in both experiments (Short $t_{39}=3.0291$, $p=.004$; Long $t_{39}=4.005$, $p=.0003$). Thus, participants looked less to each shape and, at the same time, fixated on more shapes as the training session progressed in both experiment without any noticeable difference between the short and long training. More than 70% of the overall looking time was spent within the gaze-contingent

mid-regions despite the fact that it constitutes only approx. 1/3rd of the total area of the grid (Short=74.05 +/- 6.6%, Long=72.05 +/- 13%).

Cell Transitions

The mean number of transitions/trial between cells in Implicit Short it was 7.36 +/- 1.71, while in the Implicit Long it was 6.43 +/- 1.86. The vast majority of transitions was made to adjacent cells (Short: 93.6 +/- 3.9%; Long 93.8 +/- 3% of transitions). The most common transition orientation in both experiments was horizontal (Num. horizontal transition per trial: Short= 3.68 +/- 1.05, Long= 2.98 +/- 1.15) followed by vertical (Num. vertical trans. per trial Short = 2.35 +/- 0.84, Long = 2.25 +/- 0.62) and diagonal (Num. diag. trans. per trial Short= 0.85 +/- 0.45, Long= 0.8 +/- 0.41). The total number of transition events over the course of the whole experiment for a participant was 1059 +/- 246 and 1851 +/- 535 for the short and long experiments, respectively.

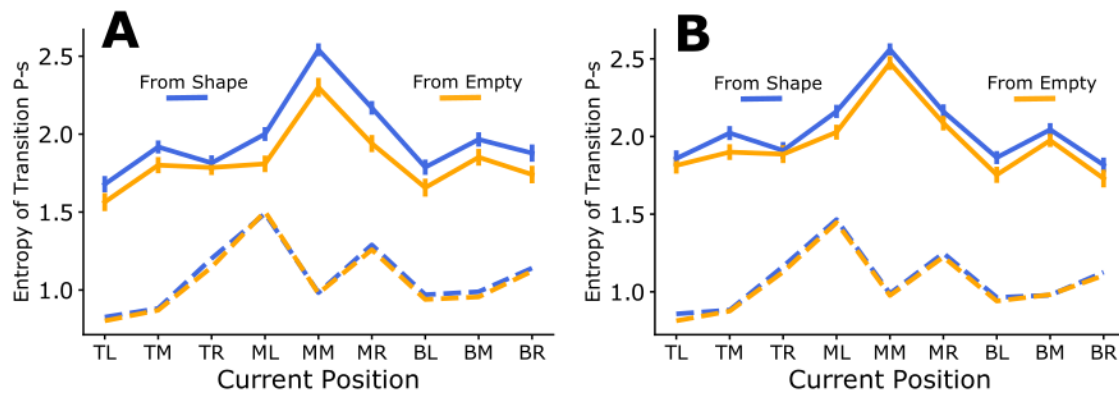


Figure 2.8: Entropy of Transition Distributions for Exp2 (A) and Exp3 (B): From Each Position on the Grid (x-axis), we calculated the entropy of transition distributions, separated by whether the cells contained an object (blue) or is empty (orange). Solid lines are the actual data averaged across participants, dashed lines are results from simulated scanner model. Error bars are SEM (not visible for scanner model). We can see that people are more unpredictable than what would be expected from scanning. Unlike the scanner model, whether the cell contains something influences the predictability of behaviour: how people continue exploration is more predictable if a cell is empty. (X-axis: cells from Top-Left (TL) to Bottom-Right (BR), T: Top, B: Bottom, L: Left, R: Right, M: Middle)

Predictability of Transitions

We found that, in both experiments, the entropy of transitions was higher if the cells contained a shape compared to cases with empty cells (Implicit Short $t_{39}=6.9333$, $p<.0001$; Implicit Long: $t_{39}=6.4809$, $p<.0001$) (**Fig 2.8**). This confirmed that the information content of the cells influenced exploration: eye movement transitions were more unpredictable if the cell contained a shape and more stereotypical if it did not. Hence, the visual input influences where people direct their next fixation.

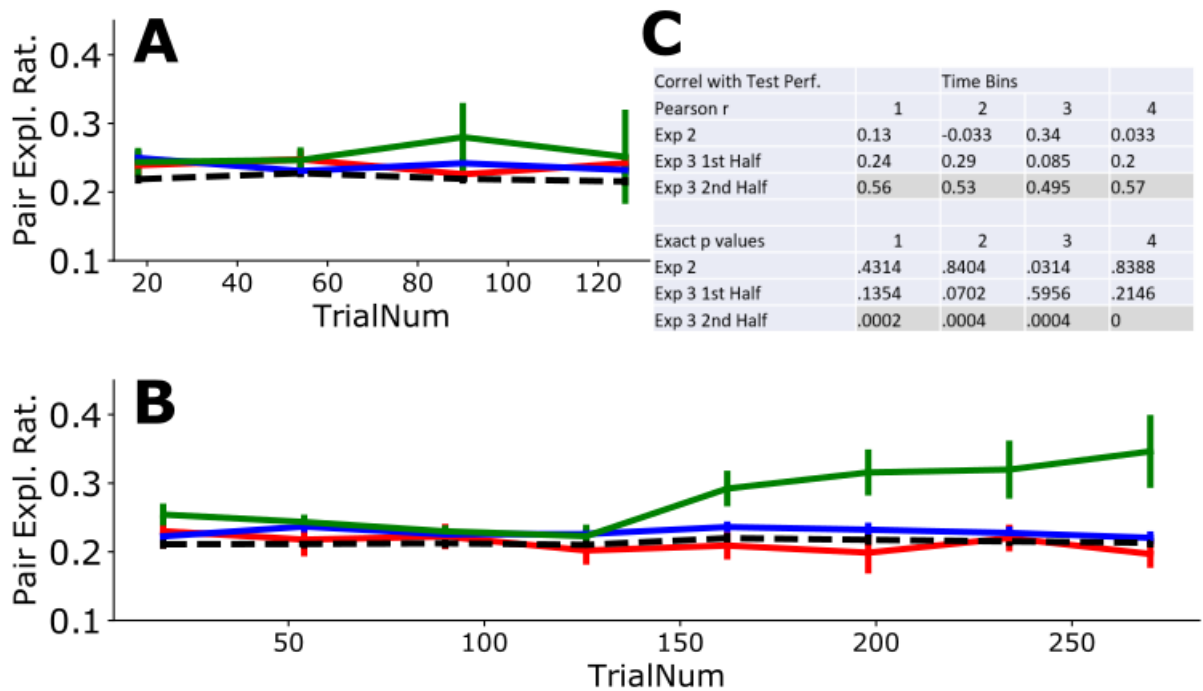


Figure 2.9: Pair Exploration Rate over time in the Short -Exp2 (A) and Long - Exp3 (B) experiments. High learners in the Long experiment gradually started to use statistical information to guide visual exploration. Participants are grouped by performance with red, blue and green as low, medium and high learners. Dashed Line represents shuffled chance. C) Correlation of Pair Exploration Rate with familiarity test performance for each time bin for the two experiments shows that only in the 2nd half of Exp 3 there is a strong relationship between looking patterns and learning. The top table contains the r values of the correlation, the bottom table the exact p values obtained by permutation. Dark gray background: significant correlations after Bonferroni correction (Exp 2 $p_{crit}=.0125$, Exp3, $p_{crit}=.00625$) (Bayes Factor Values for the correlation at each temporal bin in Exp 2: $BF_{Q1}=0.27$, $BF_{Q2}=0.2$, $BF_{Q3}=1.8$, $BF_{Q4}=0.2$. In Exp 3: $BF_{E1}=0.6$, $BF_{E2}=0.95$, $BF_{E3}=0.23$, $BF_{E4}=0.21$, $BF_{E5}=195.73$, $BF_{E6}=79.47$, $BF_{E7}=32.01$, $BF_{E8}=215.97$)

Learning Related Measures

Pair Exploration Rate

Our main measure of interest was whether and how people used the information about pairs to guide their visual exploration. We found that the Pair Exploration Rate correlated with the amount of learning in the Long Experiment ($r_{38}=.5513$, $p=.0002$, $p_{exact}=.0004$, $BF=138.5$ Fig 2.9B), but not in the Short Experiment ($r_{38}=.1654$, $p=.3079$, $p_{exact}=.31$, $BF=0.325$ Fig 2.9A). As the first half of the Long Experiment is identical to the entire learning phase of the Short Experiment, predictive relationships in the first half of the Long experiment should resemble the ones in the Short experiment. To test this, we looked at the same correlation measure in 36 trial long temporal bins for both experiments. None of the time-bins of the Implicit Short experiment were significantly predictive of learning, and similarly, none of the bins from the first half of the Long experiment show a significant correlation

(See Table **Fig 2.9 C**). On the other hand, in the second half of the Long Experiment, all time-bins were predictive of learning outcome (See Table on **Fig 2.9 C**), despite using a conservative Bonferroni corrected alpha level. To confirm that there was a link between the change in the within-pair exploration rate and learning, we calculated the slope of this rate, as this measure did not require binning the data. The slope of Pair Exploration Rate was strongly correlated with learning in the Long Experiment ($r_{38}=.4737$, $p=.002$, $p_{\text{exact}}=.0008$), but not in the Short Experiment ($r_{38}=0.0333$, $p=.8383$, $p_{\text{exact}}=.8416$). The Long result suggests that we can track learning through the eye-movements since they can predict learning performance on the subsequent familiarity test. Interestingly, the Short correlational result remained non-significant even though the magnitude of learning after the Short and Long training was virtually indistinguishable (65.9% vs. 69.6%). This indicates that the interaction between eye movements and the amount of learning is not a direct one and that the indirect effect

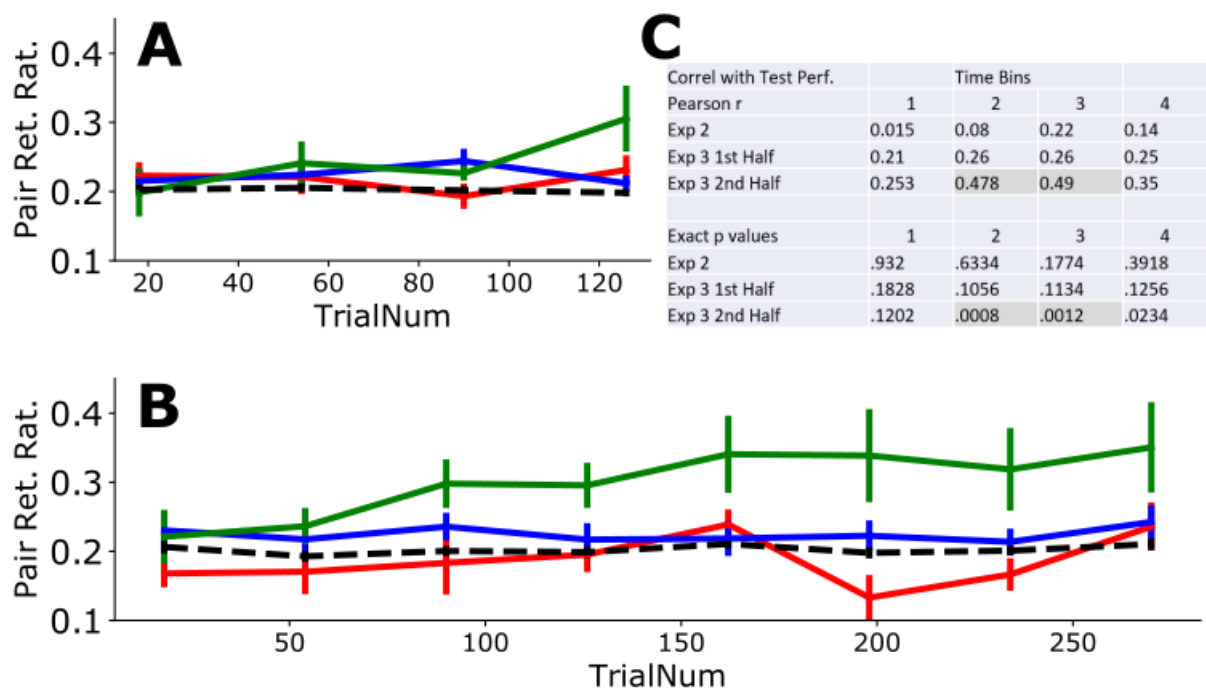


Figure 2.10: Pair Return Rate over time in experiments 2 & 3. **A)** In the Short Experiment there was no significant learning related change in this measure, while in the Long Experiment (**B**), it became highly predictive of familiarity test performance in the second half. Each bin contains 36 consecutive trials. Participants are grouped by performance as low, medium and high learners denoted by red, blue and green colors, respectively. Dashed line represents shuffled chance level. **C)** Correlation of Pair Return Rate with familiarity test performance for each time bin for the two experiments. The top table are the r values of the correlation, the bottom the exact p values obtained by permutation. Dark gray background : significant correlations (Exp 2 $p_{\text{crit}}=.0125$, Exp 3 $p_{\text{crit}}=.00625$) (Bayes Factor Values for the correlation at each bin Exp 2: $BF_{Q1}=0.19$, $BF_{Q2}=0.22$, $BF_{Q3}=0.48$, $BF_{Q4}=0.28$, Exp 3: $BF_{E1}=0.46$, $BF_{E2}=0.68$, $BF_{E3}=0.68$, $BF_{E4}=0.6$, $BF_{E5}=0.65$, $BF_{E6}=21.34$, $BF_{E7}=26.19$, $BF_{E8}=2.25$)

only emerges after certain amount of implicit learning has already taken place.

Confirmatory Looks: Pair Return Rate

We calculated a different measure, the proportion of returns made within pairs, which only includes looks that are made to cells that had already been visited on a trial. Similarly, to Pair Exploration Rate, we found a correlation with learning in the Long experiment ($r_{38} = .5433$, $p = .0003$, $p_{\text{exact}} = .0004$, $\text{BF} = 110.3$) that was stronger in the second half (**Fig 2.10B**). In the Short experiment there was no correlation with learning ($r_{38} = 0.1824$, $p = .26$, $p_{\text{exact}} = .2528$, $\text{BF} = 0.36$ **Fig 2.10A**). Looking at the temporal patterns confirms these results, with no predictive influence in any of the bins in the Short or in the first half of the Long Experiment (**Fig 2.10C**). The bins of the Long experiment became predictive of learning in the second half (**Fig 2.10C**). Similarly, to Exp 1, the slope of this measure was not predictive of learning in either experiment (Short $r_{38} = .0793$, $p = .6266$, $p_{\text{exact}} = .617$; Long $r_{38} = .2761$, $p = .0846$, $p_{\text{exact}} = .0834$).

Model Based analysis of Statistical Influence

Overall, the α parameter of the M1 model was correlated with learning performance in the Implicit Short experiment ($r_{38} = 0.383$, $p = .0147$, $p_{\text{exact}} = .01299$, $\text{BF} = 3.473$). Analyzing this effect separately in subsequent temporal bins showed that the influence of pairs on the exploration data emerges over-time (See **Fig 2.11 A,C**). The correlation with learning performance was not significant for the Long Experiment overall ($r_{38} = 0.252$, $p = .1171$, $p_{\text{exact}} = .1192$, $\text{BF} = 0.644$) due to the complete absence of such effect in the first half of training, but a strong relationship emerged in the second half as can be seen from the binned analysis (See **Fig 2.11 B,C**). The model selection showed that 8 and 9 participants in the Short and Long Experiment, respectively, are better fitted by M_1 than M_0 ¹². This classification was highly dependent on learning performance as the group of participants classified better by M_1 included 75% of high learners for both experiments (Short: 3/4, Long: 6/8), less than

¹² For individual model selection results, see Appendix Fig A.4

1/5th of medium learners (Short: 3/16; Long: 2/17), and below 10% of non-learners (Short: 2/20, Long: 1/15).

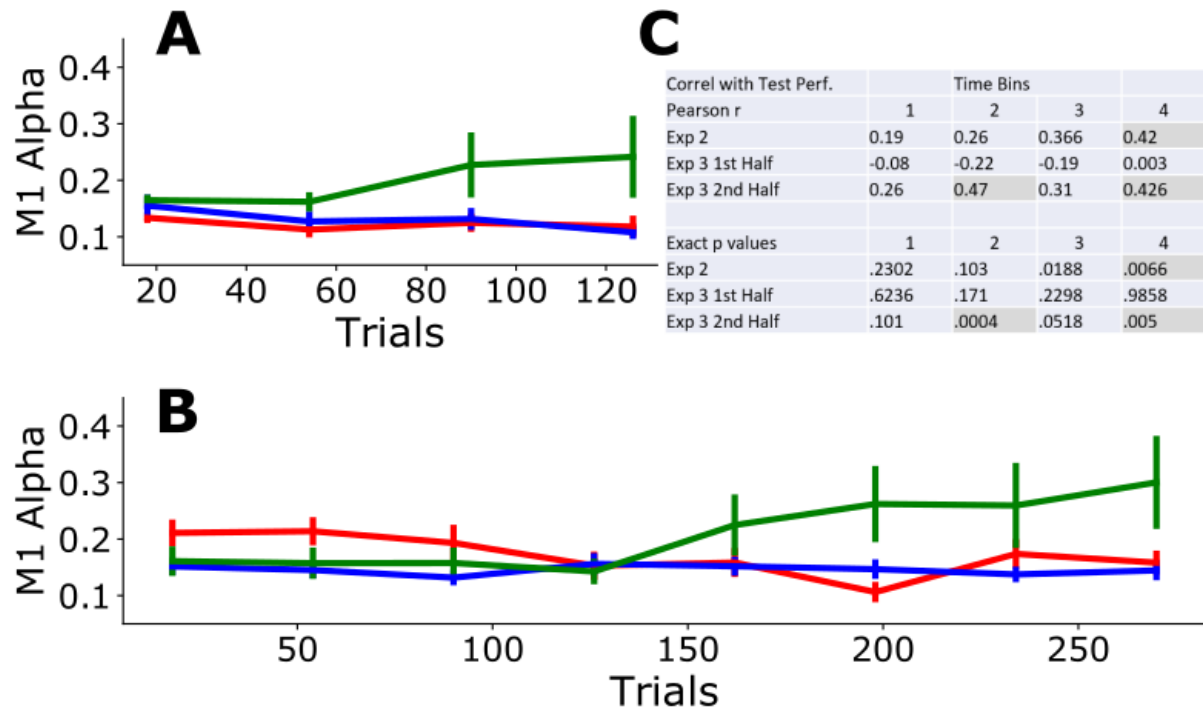


Figure 2.11. Model based assessment of Pair influence. Alpha parameter of M_1 over time in Exp 2 (A) and Exp 3 (B) Over time, high learners' (green) exploration patterns became more influenced by the pair structure of the task in both Experiments. Unlike the previous measures on Figs 2.8-2.9 this effect was also highly significant by the end in the Short experiment. There was no overall influence for medium (blue) and low (red) learners. C) Correlation of M_1 alpha with familiarity test performance for each time bin for the two experiments. The top table are the r values of the correlation, the bottom the exact p values obtained by permutation. Dark gray background : significant correlations, after bonferroni correction (Exp 2 $p_{crit}=0.0125$, Exp3 $p_{crit}=0.00625$) (Bayes Factor Values for the correlation at each bin Exp 2: $BF_{Q1}=0.39$, $BF_{Q2}=0.73$, $BF_{Q3}=2.68$, $BF_{Q4}=6.92$, Exp 3: $BF_{E1}=0.22$, $BF_{E2}=0.49$, $BF_{E3}=0.39$, $BF_{E4}=0.2$, $BF_{E5}=0.71$, $BF_{E6}=17.179$, $BF_{E7}=1.22$, $BF_{E8}=7.34$)

Combined Prediction of Learning

Since there are several different behavioral measures that could be related to learning, we wanted to see their respective contributions in predicting familiarity test performance. To achieve this, we first analyzed test performance with a Multiple Linear Regression with the same five predictors (including descriptive and pair structure related measures of visual exploration) as in Experiment 1. With a linear-regression based approach, these five predictors could explain 39.3% of variance in the Long Experiment, while only 17.89 % in the Short Experiment. To select the relevant ones of these features

and to prevent over-fitting, we analyzed the same data with cross-validated Lasso as in Experiment 1. This analysis found that 36.3% of variance could still be explained in the Long-Experiment, and 15.84% of variance was explainable in the Short Experiment (See **Table 2.2.**).

In the Long experiment, the most important predictor was Pair Exploration Ratio ($\beta=0.2849$), followed by Pair Return Ratio ($\beta=.1952$), with the other predictors assigned smaller, but non-zero weights (see **Table 2.2.**). In the Short Experiment, $M_1 \alpha$ was by far the most relevant predictor ($\beta=.2709$), with the other predictors largely irrelevant. Overall, this analysis shows that visual exploration measures related to the statistical structure of the task successfully quantify learning after about half an hour of experience (Long), but the same measures are unable to predict the robust learning performance after 10-15 minutes (Short), where only the model parameter α is successful in explaining some of the variance. This shows, that similarly to Experiment 1, spatial-statistics-related influences on eye-movements are better predictors of learning than descriptive measures of visual exploration: successful learners use the structure of the scenes to guide eye-

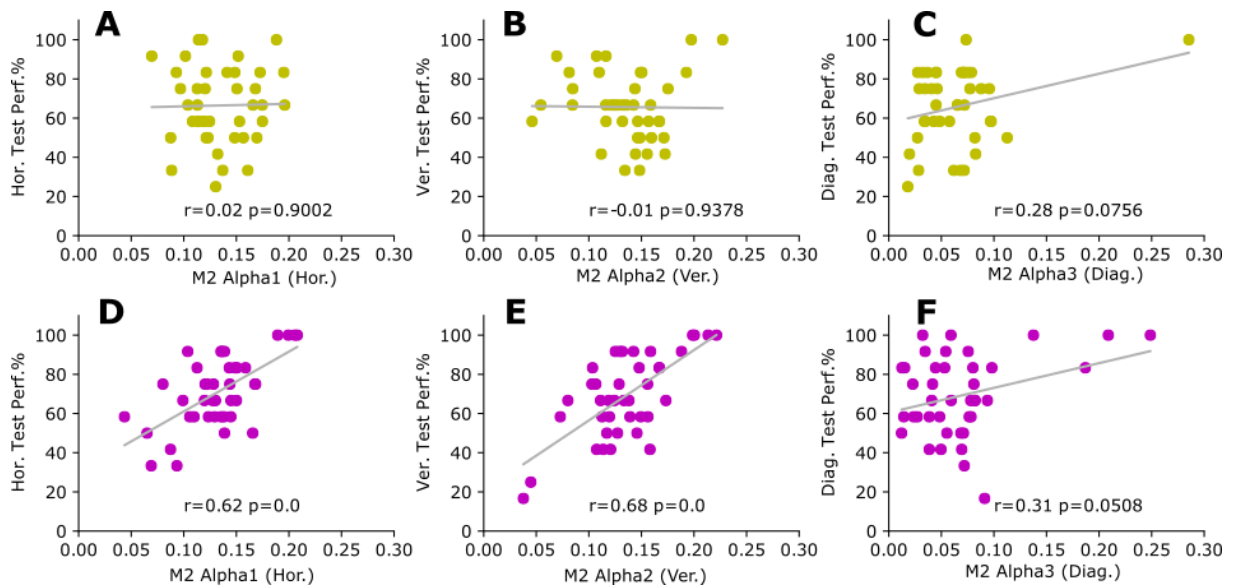


Figure 2.12: Eye-movement based prediction of specific learning outcome in Exp 2 & 3. Direction specific relationship between eye-movements ($M2 \alpha_{1-3}$ x-axis) and familiarity test performance (y-axis) for the horizontal/vertical/diagonal orientations (as Hor./Ver./Diag.) for the Implicit Short (A-C) and Long (D-F) Experiments. Least square regression line and Pearson correlation is shown on each figure. Vertical and Horizontal test performance had a tight link with vertical/horizontal eye-movements in Exp 3 ($BF_{hor} = 919.7$, $BF_{ver} = 16877$) but no relationship in Exp 2 ($BF_{hor} = 0.24$, $BF_{ver} = 0.21$). For diagonal pairs, this relationship was marginally significant for both Experiments, due to a few strong outliers, who demonstrated both strong diagonal pair influence on the eye-movements and performed well on diagonal test pairs (Exp2 $BF_{diag} = 1.01$, Exp3 $BF_{diag} = 1.9$).

movements, without necessarily visiting- or returning to more cells.

Specific Model Based Prediction of Learning

To test whether the content of learning has a specific effect on the eye-movements, we separated statistical influence into three parameters in M_2 (α_{1-3} corresponding to horizontal/vertical/diagonal orientations) and used it to predict test performance on the three orientations. We found that in the Short Experiment, eye-movements had no predictive power for any of the orientations (**Fig 2.12 A-C**) (apart for one strong outlier for diagonal orientation). In contrast, in the Long Experiment, both for horizontal and vertical pairs, there was a tight relationship (Pearson $r > .6$) between statistical influence and performance: a stronger influence of horizontal/vertical pair structure on visual exploration (as measured by α_1 & α_2) predicts better familiarity test responses on horizontal/vertical pairs (stats on **Fig 2.12 D-E**). This effect was also marginally significant for the diagonal pairs (**Fig 2.12 F**), with some high performing outliers showing a tight relationship, but no link for most participants. This connection in the Long Experiment was temporally emerging as, all three orientations replicating findings from the Explicit Experiment, the slopes of α_2 and α_3 were predictive of test performance on vertical ($r_{38} = .33$, $p = .0403$) and diagonal pairs ($r_{38} = .41$, $p = .0088$) respectively, but not the slope of α_1 for horizontal pairs ($r_{38} = -.01$, $p = .9734$), where the influence seemed to emerge early on (see also Appendix Fig A.5). This relationship between the increase in pair influence and test performance was absent for all three orientations in the Short Experiment (horizontal: $r_{38} = .26$, $p = .1013$; vertical: $r_{38} = .11$, $p = .5131$; diagonal: $r_{38} = .185$, $p = .253$). Taken together, the results using M_2 demonstrate that given sufficient time, the content of learning about statistical scenes can be predicted from eye-movements.

Discussion

Using two implicit active visual statistical learning experiments, we showed that despite robust statistical learning performance, after ~15 minutes of learning, most participants' visual exploration patterns were not influenced by this acquired regularity. However, given sufficient time, not only good learning performance but also the content of learning could be predicted from visual exploration patterns. These findings demonstrate that knowledge emerging from implicit statistical learning about the spatial relationships between novel objects can guide visual exploration.

It is interesting to contrast the findings suggested by our different measures. The raw behavioral data, based on selecting potentially learning related behaviors, suggests no influence of learning related changes in the Implicit Short experiment. Nevertheless, analyzing the same data with a more sensitive model-based approach, we could find learning related changes. Interestingly, while the same model was also predictive of learning in the Implicit Long and Explicit experiments, in those cases, the behavioral measures were dominant, better predictors of test performance suggesting that such a more sensitive model-based measure is beneficial early in the learning. A caveat is that M_1 did not show any learning effects in the first half of the Implicit Long experiment. This is somewhat puzzling, given that the first half of the Long experiment was identical to the Short experiment, apart from the fact that the participants knew that the experiment will be longer. Beyond random fluctuation across subject pools, we suspect that in the Long Experiment, participants did not focus as much in the first half as participants in the Short Experiment did, since they knew they had plenty of time was ahead. Nevertheless, participants benefited significantly from the longer learning phase in Exp 3 as shown by a higher proportion of good learners (20% of participants in Exp 3 vs 10% in Exp 2). Of course, we cannot be sure whether these good learners in Exp 3 would have been the best performers after completing the first half, but based on test performance in Exp 2, we can assume, that a subset of participants had already acquired an implicit representation of the statistical structure that was good enough to perform almost as well as at the

end of the long trial. This implicit representation might have contributed to the emergence of eye movement patterns in the second half of exploration trials that were reminiscent to those in the Explicit experiment, and unlike in the Short-Experiment, could specifically predict test performance on the different orientations (using M_2).

To our knowledge, this is the first study to reveal information about the process of spatial statistical learning. Previous attempts to explain the time course of visual statistical knowledge were confined to temporal statistical learning (Bertels, Boursain, Destrebecqz, & Gaillard, 2014; Karuza et al., 2014). Our Short Experiment shows that there is a strong implicit knowledge already after ~15 minutes, but this knowledge had no influence on the eye-movements for the vast majority of participants, apart from few very successful learners. Our Long Experiment revealed that after this Implicit representation has been built, in the second half of the learning phase, it started to influence visual exploration. Therefore, these two experiments together suggest a sequence of “statistical knowledge first, influence on visual exploration later” as a plausible relationship between statistical learning and eye-movements. What our study could not reveal was the nature of the emerging statistical knowledge for Implicit learners, since for most participants we found no influence on eye-movements, despite the reasonably good test performance. In line with previous studies (Turk-Browne et al., 2005), a parsimonious explanation of our finding is that a statistical representation emerges early on from paying attention to the stimuli, without any influence on eye movement behavior, and this pattern initially neither depends on nor influences patterns of active information search. However, given sufficient time for good learners, the two processes seem to become intimately related, where a gradual influence by the pair structure on looking patterns boosts successful learning, which in turn could make influence on looking patterns even stronger in a loop-like manner.

General Discussion

In three statistical learning experiments combined with eye tracking, we investigated the relationship between the acquisition of statistical information and visual exploration. The first experiment introduced a novel a gaze-contingent spatial statistical learning paradigm and showed that our measures can successfully track and predict the explicit learning of spatial regularities in the eye-movement data. The second and third experiments investigated implicit statistical learning using the same stimuli and two different learning-phase lengths. First, we extended our findings from Experiment 1 by showing that the same measures can track the transfer from implicit to explicit knowledge and predict individual learning success. Second, we found, that given extended exposure (Exp 3) eye-movements can predict what is learned about complex statistical scenes, suggesting a tight link between the development of internal statistical representations and looking patterns. Finally, we found that a strong statistical representation can be learned without any influence on eye-movements (Exp 2), suggesting that statistical learning relies on a multitude of statistical information provided by the scenes, many of which is not intimately linked to eye movements initially.

Our results across the three experiments show some remarkable patterns. There are several similarities between the patterns of the Explicit (Exp1) and of the second half of the Long Implicit (Exp3) experiments. First, we confirmed our claim on another large sample that the four different exploration measures (pair exploration rate, pair return rate, overall model-based pair influence/ M_1 , specific model based-based pair influence/ M_2) are reliable eye-movement-based predictors of learning, with very similar temporal emergence pattern between the Explicit and the second half of the Long Implicit Experiments. Second, we found a high proportion of very good learners (85%+ test performance) with about $1/5$ - $1/4^{\text{th}}$ of participants falling into this category both in the Explicit and Long Implicit experiments, while only $1/10^{\text{th}}$ of participants reached that criterion in the Short

Implicit experiment. Taken together, these measures suggest that both the exploration behavior and the acquired statistical knowledge are very similar in the second half of the Implicit Long and in the Explicit Experiments.

We found this intricate relationship between the emergence of statistical knowledge, and its influence on visual exploration using several measures between these two experiments, while we sidestepped the question of what the nature of the representation that emerges through this active learning process is. It is possible that despite the implicit task, after some learning has taken place, observers realized that there were pairs which could facilitate the emergence of statistics related patterns in exploration behavior, especially in the Long Experiment. Thus their performance measured in the final familiarity test might not be the outcome of a pure implicit process, echoing some recent findings in the literature (Bertels et al., 2014; Bertels, Franco, & Destrebecqz, 2012), but contradicting some other earlier reports (Kim et al., 2009; Turk-Browne et al., 2005). Nevertheless, based on the high performance and the above described similarities between the Implicit Long and Explicit experiments, we think that what we see in the Implicit Long Experiment is the transfer from implicit to Explicit representation.

Previous studies on memory-related eye-movements show that eye-movements can reflect memories about the relative spatial locations of objects. For example, it has been found that the relative but not the absolute locations of objects determine whether the patterns of eye-movements reflect previously seen configurations (Ryan & Villate, 2009). However, in these previous experiments, memory was probed immediately after stimulus presentation, and there was no need to integrate information across multiple trials, suggesting that those experiments have only measured visual short-term memory. Our study is completely novel in the sense that it shows that eye-movements can be used to probe spatial relationships between shapes that can only be acquired over many trials.

Our model-based method proved to be a more sensitive measure of statistical influences on exploration than the raw behavioral measures due to its relative sensitivity to within-pair transitions at a given location compared to average behavior. Furthermore, we were also able to link variability in what is learned to the directional influences on eye-movements using Model 2. One obvious shortcoming, however, is that we fitted a three-parameter pair-influence model (M_2) and used it to predict performance on each orientation, despite having six different pairs (two for each orientation). In the current data set, it is difficult to link looking patterns to success in identifying certain unique pairs at test. The main reason for this difficulty is that high learners responded to almost all test trial correctly, making the success rate at the different pairs equally 100%. This ceiling effect makes the prediction of learning success for a specific pair, given the number of times it was seen, impossible. Separation of pairs based on performance for medium learners is possible, however, most medium learners showed only weak overall structure related effects. Therefore, eye-movement based within subject prediction of learning success on unique statistical learning pairs remains to be a challenge for future research.

Eye-movements and statistical learning

To our best knowledge, our results are among the first to provide evidence about the role of eye-movements during statistical learning. One exception analyzed eye-movements in an infant statistical word-referent learning paradigm, finding notable differences in the eye-movement patterns of good and bad learners' (C. Yu & Smith, 2011). However, due to the limitations of infant studies, both familiarization, and test of this work was based on eye-movements, and unlike our study, it could not link an eye-movement-independent measure to eye-movements. Therefore, it could not show that eye-movements can reflect generalizable statistical knowledge. A similar paradigm was tested with adults using explicit instructions, offering a simple 'propose then verify' model as an explanation for explicit word-referent learning (Trueswell, Medina, Hafri, & Gleitman, 2013), suggesting that learning that proceeds by only keeping in mind one hypothesis at a time. It is possible that participants use similar strategies in our Explicit paradigm, where they explicitly hypothesize certain combinations,

which they can confirm or falsify later. An indication of such a hypothesis testing strategy could be returns to pairs, which were by far the best predictor of learning in Explicit experiment. In experiments 2-3, explorative transitions (Exp 2) and overall pair influence measured by our model (Exp 1) had tighter link with learning, suggesting that different strategies were dominating learning in the Implicit experiments. Previously, in the implicit statistical learning literature, there were attempts to link visual attention to statistical learning using reaction times. However, that line of research could only show that the mere presence of regularities can attract attention without being able to establish a link with individual learning outcome (R. Q. Yu & Zhao, 2015; J. Zhao et al., 2013). Therefore, our study is the first to link learning to eye-movements in implicit statistical learning.

It is interesting to scrutinize our findings in relation to the proposed inverse *u*-shape relationship between stimulus predictability and visual attention (Kidd, Piantadosi, et al., 2012). According to our results, better learning arises from and/or results in an increased search for predictable structures. According to Kidd et al.'s logic, if a fully predictive relationship is acquired, that should result in a weaker attentional bias. Indeed, we see this pattern in our data of a few participants in Experiment 1 (**Fig 2.4D**), but our task is complex enough so that most participants do not reach the level of knowledge that would render the relationships between pairs of shapes as completely predictable. Therefore, the weakening effect implied by Kidd et al.'s (2012) proposal, is probably not strong enough to dominate the overall proportion-of-returns to pairs, as the measurement for the majority of the participants is still increasing across the entire duration despite the drop toward the end for a few observers. Given some even longer learning period, it is possible that the drop-offs by the end of learning would dominate the observers' behavior and they would explore less across pairs after acquiring full knowledge of the relationship between the shapes. On the other hand, due to the heterogeneity of learning patterns and test performance across participants, finding such an inverse *u*-shape pattern in the exploration data over time could prove to be a difficult challenge for future research.

Future Directions

Our methods could not detect learning related changes in eye-movements for most implicit learners. Although a likely reason is that there are no such effects, it is also possible that such effects are present, but they do not affect exploration patterns, but other measures of eye-movements. It would be interesting to consider different eye-movement related measures and see if they could show pair related learning effects in implicit learners. One possibility is looking times, which following the logic of Karuza et al (2014), could become shorter for predictable shapes once they are learned. However, the challenge is that the opposite prediction can also be expected: during learning of predictive relationships, looking at a pair could elicit longer looking times initially, while the time-course of these effect could vary across participants.

Our data analysis models proved to be a sensitive measure of pair structure influences in the eye-movement data, however, by design they were not learning models that could provide insight into the process of information acquisition. In order to model the learning process itself, something along the lines of a Bayesian Chunk Learner (Orban, Fiser, Aslin, & Lengyel, 2008) could be used. The Bayesian Chunk Learner has been successful in explaining spatial statistical learning in different paradigms, but only in scenarios where the full statistical scenes were presented at once to the observers. It remains to be seen whether such a model could be constrained by the exploration data of human observers and predict individual learning outcome.

Conclusions

We showed that eye-movements are a sensitive measure, which correlate with and, therefore, could be used for predicting the learning of spatial regularities. We also provided some evidence that this method presumably tracks the transfer from implicit to a more explicit internal representation. However, statistical learning is an automatic process, which can be robust without any effect on visual exploration. These findings also speak to studies of human information search, showing that people can use arbitrary spatial relationships of novel stimuli to guide visual exploration. However,

such influences were mostly confined to successful learners, leaving the problem whether fully implicit statistical learning influences eye-movements as an open question for future research.

Chapter 3.

The link between Statistical Learning, Eye-movements and Working Memory

Summary

In this chapter, we investigated eye-movements in three visual statistical learning experiments, while participants performed a working memory task in parallel. In Experiments 4-5, we used a gaze-contingent spatial statistical learning paradigm to show that while participants used the global structure of the scenes to guide visual exploration, they were not affected by the local pair structure of the scenes. Furthermore, we also showed that working memory and statistical learning were linked suggesting that working memory capacity could be used as a continuous indicator of statistical learning. Using a temporal statistical learning paradigm in Experiment 6, we demonstrated that eye-movements represented a sensitive measure of statistical regularities even when observers showed no evidence of learning during a subsequent familiarity test.

Introduction

In many statistical learning studies (Fiser & Aslin, 2001, 2002; Saffran, Johnson, Aslin, & Newport, 1999), just as in our Experiment 2-3 of Chapter 2, learning emerges without participants receiving specific instructions on what to do exactly with the stimuli. We used this approach successfully in the previous chapter to establish learning related influences on gaze contingent exploration patterns. However, we found a large individual variability both in terms of learning performance (as measured by the familiarity test) and in the statistical influences on eye-movements. An important reason behind such variability could be the open-ended nature of the task, which allowed participants to freely choose how exactly they engage with the stimuli. One potential way to reduce this variability is to include a well-defined task that the observer has to perform, which is unrelated to the statistical structure of interest, but ensures that the observer's attention is maintained on the stimuli. Similar approaches have been used before successfully to measure the learning of temporal- (Turk-Browne et al., 2005; J. Zhao et al., 2013) and spatial regularities (Chun & Jiang, 1998). A second advantage of having an independent task, while unbeknownst to the participants, the stimuli have regularities, is that responses related to the task might be used to track learning itself (Chun & Jiang, 1998; Howard Jr & Howard, 1997; Karuza et al., 2014). A third advantage is that this method can provide insights about the interaction between the explicit task and statistical learning. Interestingly, these advantages were not explored extensively before. For example, despite fact that some of the measures in the studies mentioned above relied on visual search, little is known about whether those effect are mediated via overt attention, since eye-movements were not tracked in most experiments (R. Q. Yu & Zhao, 2015; J. Zhao et al., 2013).

A good candidate for a parallel task performed during implicit statistical learning is the one-back task testing working memory performance (Owen, McMillan, Laird, & Bullmore, 2005). First, a simple one-back memory task requires continuous engagement with the stimuli, but only occasional

interruptions of the stimulus presentation stream. Second and more importantly, abilities to perform in working-memory tasks have been shown before to correlate with learning performance. Specifically, several studies reported that when associations between sensory elements are learned in an implicit manner, observers demonstrate an increased working memory capacity for compound stimuli containing the associated pairs, presumably because the co-occurring elements are stored in the memory more efficiently as chunks (Brady, Konkle, & Alvarez, 2009; Brady, Störmer, & Alvarez, 2016). Therefore, a working memory task running in parallel while eye-movements are measured during statistical learning could not only keep people engaged with the stimuli, but it could also be used as a measure of the chunk learning process itself. However, at present time it is underexplored what type of regularities working memory capacity can benefit from. Furthermore, it is unknown whether adding such a task would impact statistical influences on eye-movements. To explore these questions, we ran the next three visual statistical learning experiments while the participants performed a parallel working memory task at the time of exposure.

Experiments 4-5: Active Statistical Learning and Working Memory

Introduction

Previous studies used a parallel visual search task to measure implicit learning of the global arrangement of identical elements predicting a single target (Chun & Jiang, 1999) or predictive temporal relationships between novel shapes (J. Zhao et al., 2013). Meanwhile, a different study showed that working memory capacity is sensitive to learned regularities (Brady et al., 2009). However, the regularities in this latter study were easily distinguishable as they did not require segmenting complex scenes or sequences of stimuli in an unsupervised manner as is necessary in typical visual statistical learning studies. There are only a few published papers linking these two research directions by exploring whether the kind of associations that are regularly used in statistical

learning studies would also be able to enhance working memory capacity by enabling chunked representations (Nassar, Helmers, & Frank, 2018). Moreover, none of these explored the link of these phenomena to eye-movements.

We had three main goals with combining active statistical learning with a working memory task: 1) To confirm that people can learn about the spatial statistical structure despite the potentially taxing parallel task. 2) To test whether working memory performance can indicate the progression of statistical learning with more complex stimuli than the one used by Brady et al (2009), which did not require unsupervised segmentation across trials. 3) To see what kind of statistical information is learned and whether this learning influences the patterns of visual exploration.

In order to achieve the first goal in Exp 4, we combined the active statistical learning paradigm of Chapter 2 with interleaved memory probes. In order to have a clear baseline for the 2nd and 3rd goals, we ran a control experiment (Exp 5), in which we kept the global structure of the scenes intact but shuffled the shapes within the overall silhouette. This manipulation allowed us to test two levels of statistical complexity (pair structure vs global structure) with respect to their respective influence on visual exploration and working memory performance.

Methods

Participants

Participants gave informed consent before the start of the experiment and received 1500 HUF worth of food vouchers as compensation. 41 students completed Exp 4, from which one participant was excluded because of not exploring the scenes and had an average looking time to central cell regions over 3 seconds (Group Mean \pm SD = 539 \pm 114 msec), leaving a final sample of 40 (13 male, 35 Right Handed). 37 students (9 male, 34 Right Handed) completed Exp 5 without any exclusion.

Stimuli

In Experiment 4, the stimuli were the same as in Experiments 1-3, with 144 unique scenes (**Fig 3.1 A-B**). In Experiment 5, the stimuli were the same, but the arrangement within the scenes was randomly selected: the same scenes were used as in Experiment 4, however the position of the shapes within the scenes was shuffled on each trial scenes (**Fig 3.1 A,C**). This way, the outline of the scenes was identical to those in Exp 4, but the position of shapes within the scenes was unpredictable. Importantly, this manipulation also left the joint appearance probability of the shapes unchanged:




Figure 3.1. Structure of Exp 4 & 5 **A**) Example Statistical Learning Pair Structure. **B**) A statistical learning scene assembled from three pairs in A, as used in Exp 4. (and Exps 1-3.) **C**) Example scene of Exp 5. Locations of shapes from the same three pairs are shuffled while keeping the original overall silhouette. This way, the global structure of the scene is intact, without any spatial predictive power for individual shapes.

shapes forming pairs were always present in the same scenes, but they were not predictive of each other's relative spatial location.

Working Memory Task

The stream of active exploration trials was occasionally interrupted by working memory probes. On these probes, all 12 shapes were simultaneously presented at the two sides (6-6 on each side) of the presentation grid. The order of the 12 shapes on the working memory probes was randomly selected for each participant but did not change within the experiment. Participants had to use the mouse to select the six objects that had been present on the immediately preceding exploration trial. They had unlimited time to perform this selection and received no feedback on their choices. There were 14 working memory trials pseudo-randomly interleaved within the 144 exploration trials. The same

pseudo-randomization was used for each participant, to avoid unnecessary variability across participants. Two additional memory trials were included: after the last exploration trial and after the exploration trial following the first memory trial (6th trial), resulting in a total number of 16 memory trials. The goal of including two subsequent memory probes early on was to encourage participants to pay attention even immediately after memory trials. Later memory trials did not follow subsequent learning trials.

Familiarity Test

In both experiments we used the same 36 trial two-interval forced choice familiarity test as in Exps 1-3, with the 6 real pairs tested against 6 foil pairs. This was an especially challenging test in Exp 5, since participants could not use the spatial arrangement of the real pairs as a cue. In Exp 5, the only information to learn which shapes form pairs was that pairs of shapes were always present in the same scenes albeit at unpredictable locations, while other shapes were co-present on maximum 50% of scenes. However, this regularity in Exp 5 was very hard to notice since the shapes forming a pair would not be adjacent on most of the trials, and thus they would not be visible together due to the gaze-contingent set-up.

Procedure

After calibrating the eye-tracker, participants performed 15 practice trials, during which they familiarized themselves with the exploration of the scenes in a gaze-contingent manner by revealing randomly selected images of dogs on the 3 by 3 presentation grid (as in Exps 1-3). The practice exploration was interrupted three times with practice working memory probes, providing feedback on the number of correct choices after each memory trial. After the practice, the calibration of the eye-tracker was double-checked and if necessary recalibrated, and next, the main experiment started. Participants were told to pay attention, explore the scenes and perform the working memory task as well as they could. They were not told about any regularities, nor that they will have to answer some additional test questions after the exposure presentations.

Data Analysis

On the visual exploration data of Exp 4, we have performed the same analyses that were described in the General Methods section of Chapter 2. Although the stimuli of Exp 5 had no spatially predictive pairs, we calculated the same measures as for the other experiments, as if the pairs were in their original positions (before the shuffling). This way, Exp 5 gave us a baseline for pair transition measures: the proportion of looks within positions that contain pairs in Exp 4, but without an actual pair structure.

In order to assess the knowledge of Global Statistical Structure, we introduce a novel measure, based on three types of transition from a given shape:

X_1 : Transition to the other shape of the pair the shape belonged to

X_2 : Transition to another shape (that is not the other shape of the pair)

X_3 : Transition to an empty cell

We defined our new measure, the Shape Exploration Rate as: $(X_1 + X_2) / (X_1 + X_2 + X_3)$. This measure is sensitive to the global structure, without being sensitive to the internal pair structure.

To obtain a chance level, we randomly paired the visual exploration data with the presented stimuli 100 times for each participant and calculated the above-defined measure for each shuffled data set. Finally, we averaged over the 100 simulated values to get an individual chance level.

To compare measures between groups with largely different sample size, we used the permutation/randomization test (Craig & Fisher, 1936). First, the overall dataset was randomly divided into groups two (same as the original sample-sizes, without replacement) 5000 times. The difference between the groups was calculated for each permutation. Afterwards, we determined where the actual measured difference fell within the distribution of permuted differences, to obtain a p value.




Figure 3.2 Familiarity Test Performance Distribution Exp 4 & 5 **A)** Exp 4 **B)** Exp 5. Participants are grouped into Low, Medium and High Learners as Red/Blue/Green. Vertical Solid Black/Grey Line: Mean/Median, Vertical Dashed Line: Chance. Performance was highly above change in Exp 4, with a significantly better performance than Exp 5, where the performance was still slightly above chance (stats in main text).

Results

Familiarity Test Performance

On the 2-IFC familiarity test in Exp 4, participants demonstrated a robust learning (63.47 +/- 17.6% Correct; Difference from chance: $t_{39}=4.781$, $p<.0001$, $BF=866.15$) (**Fig 3.2A**). Performance on experiment 5 was slightly but significantly above chance (55.86 +/- 14.01% Correct, Difference from Chance $t_{36}=2.499$, $p=.0171$, $BF=2.66$) (**Fig 3.2B**). Performance in Exp 4 was significantly better than in Exp 5 ($t_{75}=2.0605$, $p=.0428$, $BF=1.44$).

Working Memory Performance

In both experiments, participants performed the working memory task significantly better than chance (Exp 4 $t_{39}=10.9014$, $p<.0001$; Exp 5 $t_{37}=11.598$, $p<.0001$). On average, they selected 4.02 +/- 0.59 correct images in Exp 4. and 3.94 +/- .48 in Exp 5, thus their performance was not different in the two experiments ($t_{75}=.677$, $p=.5007$, $BF=.29$) (**Fig 3.3A**). To test whether participants' performance was changing over time, we fitted a least square regression line to the individual working memory performance for each subject with trial number as the predictor. The mean slope of the regression line was positive in Exp 4 ($t_{39}=3.535$, $p=.0011$, $BF=28.79$), indicating that participants' performance was improving over-time (**Fig 3.3B**). In Exp 5, the overall trend was also positive, but it fell short of significance ($t_{36}=1.7457$, $p=.0894$, $BF=0.7$). The slopes between the two

experiments were not significantly different ($t_{75}=1.6124$, $p= .1111$, $BF= .722$), with Bayes Factor suggesting that data is insensitive. The reasons for this could be that some participants improved over-time in the working memory task in Exp 5, for reasons unrelated to the pair structure.

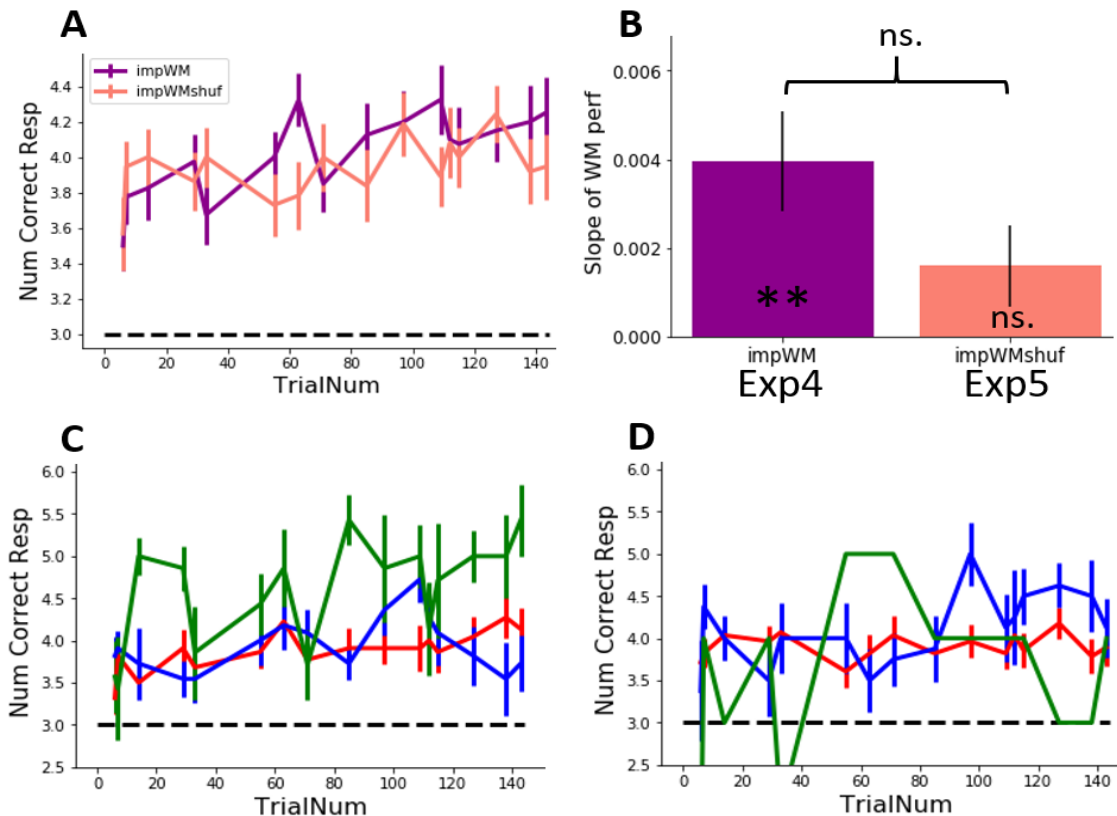


Figure 3.3. Working Memory Task Performance in Exp 4 & 5 **A)** WM task perf. over time (16 trials interleaved into the 144 exploration trials) for Exp 4-5 (impWM-impWMshuf). Overall performance was not different between the experiments. **B)** Slope of WM performance was significantly above chance in Exp 4 but not in Exp 5, though the difference was not significantly different. **C)** Working memory performance in Exp 4, grouped by familiarity test performance. Though initially everyone was at the same level, good statistical learners improved in the working memory task over time **D)**. No relationship between familiarity test performance and working memory task was found in Exp 5. Colors represent Low Medium and High learners as in Fig 3.2. Errorbars: SEM

Relationship of Working Memory and Statistical Learning performance

In Exp 4, good learners on the familiarity test, were also better in the working memory task ($r_{38}=.423$, $p=.0066$, $p_{\text{exact}}=.0046$, $BF=6.938$). All groups started from the same level of working memory performance (**Fig 3.3C**), suggesting that the relationship is in fact related to learning, and not just a general influence of attention or effort. This is further supported by the fact that in Exp 5, we found no relationship between statistical learning and working memory performance ($r_{35}=-.011$, $p=.9484$,

$p_{\text{exact}}=.9492$, $\text{BF}=0.205$, **Fig 3.3D**), suggesting that for this correlation the presence of the pair structure necessary.

Descriptives of Visual Exploration

Average looking time at the gaze contingent mid-region of cells was 539 +/- 115 msec in Exp 4 and 559 +/- 107 msec in Exp 5. Notably, these looking times were shorter and consequently the number of visited cells was larger than those reported for the Implicit Experiments (Exp 2 & 3: 656 and 736 ms), but very similar to the Explicit Experiment (Exp1: 531ms) of Chapter 2. This suggests that giving participants a well-defined task encouraged them to explore the scenes faster. Similarly, the number of transitions was not different between the two experiments. Participants made 8.39 +/- 1.32 transitions between cells per trial in Exp 4, and 8.18 +/- 1.42 transitions in Exp 5, adding up roughly to about 1200 transition events over the course of the Experiments (Total Number of Transition Events: Exp 4= 1207.4 +/- 190.1, Exp 5=1177.9 +/- 205.0). Importantly, as in Experiments 1-3, entropy-based analysis showed that transitions were more unpredictable if they were initiated from cells that contained an object at any given trial as opposed to empty cells (Exp 4= $t_{39}=9.6859$ $p<.0001$, Exp 5 $t_{36}=9.056$ $p<.0001$).

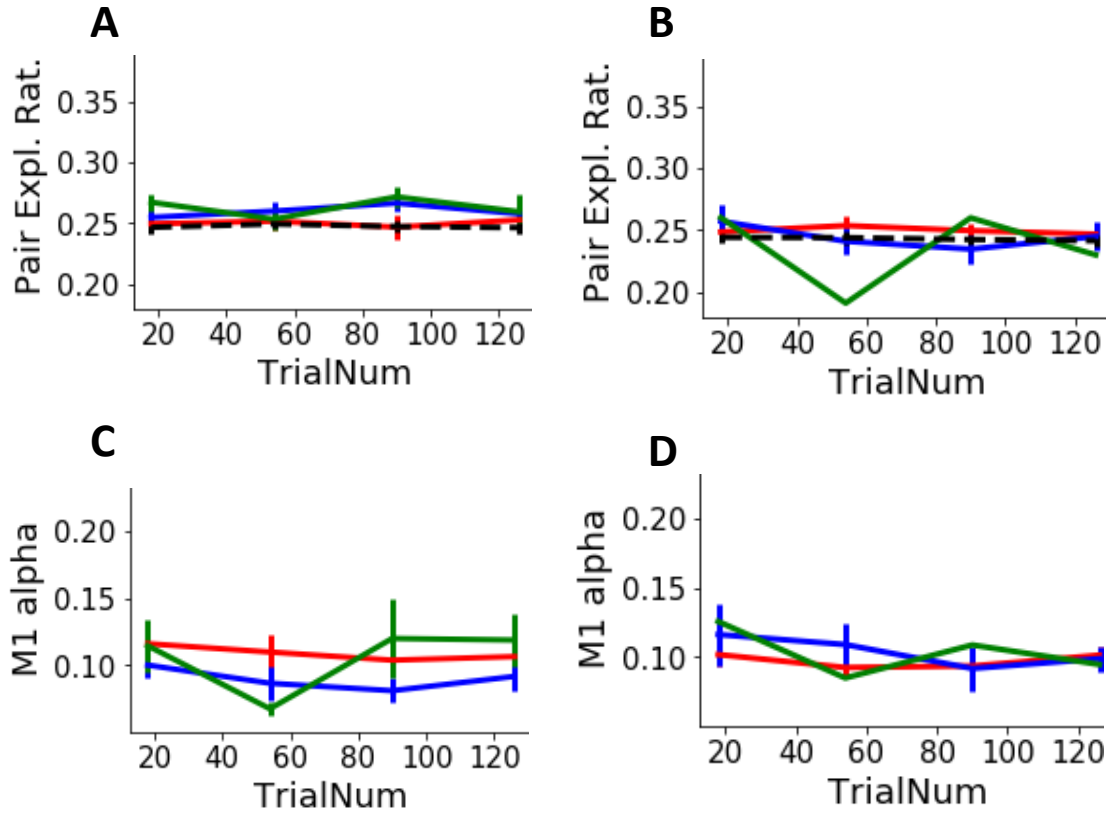


Figure 3.4. Pair Structure influence in Exps 4 & 5. There was no learning influence of the pair structure in either of the experiments. **A)** Pair Exploration Rate in Exp 4 **B)** Pair Exploration Rate in Exp 5. **C)** Model based pair influence in Exp 4. **D)** Model based pair influence in Exp 5. Colors: familiarity test performance as before. Error bars: SEM, Green line on B & D is a single participant. Dashed Black line on A-B: shuffled chance.

Pair Structure effects on Visual Exploration

Pair Transition Measures

To measure the effect of pair learning on the eye-movements, we calculated Pair Exploration Rate (**Fig 3.4A-B**) and Pair Return Rate¹³. Neither measure showed pair-structure-related changes in Exps 4 or 5, nor could they predict performance on the familiarity test (Correlation test performance: Pair Expl. Rat: Exp 4: $r_{38}=.1632$, $p=.3144$, $p_{\text{exact}}=.3270$, $\text{BF}=0.32$; Exp 5: $r_{35}=-0.1513$, $p=.3715$, $p_{\text{exact}}=.366$, $\text{BF}=0.3$; Pair Ret. Rat: Exp 4: $r_{38}=0.138$, $p=.3953$, $p_{\text{exact}}=.3954$, $\text{BF}=0.28$; Exp 5: $r_{35}=.043$, $p=.7988$, $p_{\text{exact}}=.808$, $\text{BF}=0.21$). Furthermore, the measures in Exp 4 were not significantly different from Exp 5 (Pair Exploration Rate: $t_{75}=1.1306$, $p=.2618$, $\text{BF}=0.41$; Pair Return Rate: $t_{75}=1.8195$, $p=.0728$, $\text{BF}=0.98$), suggesting that the values in Exp 4 (**Fig 3.4A**) are a consequence of exploring the silhouettes and do not reflect any influence by the pairs.

¹³ For definition of measures see General Methods in Chapter 2, for descriptive stats Appendix Table A.1.

Model based Results

We have fitted the pair-influence model (M1) to the exploration data of Exp 4, which was used in Experiments 1-3 in Chapter 2, and found no relationship between learning and the value of the pair structure sensitivity parameter α ($r_{38}=-0.113$, $p=.4891$, $p_{\text{exact}}=.4964$ BF=0.25, **Fig. 3.4C**). We did not find any correlation between α and learning in Experiment 5 either ($r_{35}=0.004$, $p=.9796$, $p_{\text{exact}}=.9794$, BF=0.2; **Fig. 3.4D**), but this was expected as there was no spatial pair structure in the stimuli of this experiment. For this reason, whatever fitted values we found, it could only be a consequence of over-fitting and cannot reflect a true influence of a statistical structure. The fitted values of α were not different between the two experiments ($t_{75}=0.5944$, $p=.554$ BF= .28), despite the prominent contrast between the pair-based structure of Experiment 4 and the “pair-less” baseline structure of Exp 5. This is a further confirmation that the pair structure in Experiment 4 had no or minimal influence on participants’ explorations.

↓ Exp.	Predictor:	β N Cells	β N Returns	β PairExp. R.	β Pair Ret. R.	β M ₁ α
WM - Exp 4	Lasso ($\lambda=.292$)	0.	0.	0.	0.	-0.
	Least Square	0.111	0.2483	0.1486	0.0198	-0.0669
WMshuffled - Exp 5	Lasso ($\lambda=.266$)	0.	0.	-0.	0.	-0.
	Least Square	0.0817	0.2179	-0.1198	-0.0016	-0.0434

Table 3.1. Combined prediction of learning. Lasso (top in each cell) and Least square (bottom) regression weights for five predictors of visual exploration (as described in main text). The regularized Lasso regression shows, that unlike in Exp 1-3, none of the 5 visual exploration-based predictors could reliably predict familiarity test performance in Exps 4-5. The regularization parameter λ was selected with cross-validation. Values are rounded to four decimals.

Combined Prediction of Statistical Learning

Using a linear regression approach, with five measures of exploration behavior as predictors (Num of Visited Cell, Num of Return, Pair Exploration Rate, Pair Return Rate, M₁ α), we have found that about 10% of variance familiarity performance could be predicted (Exp 4=11.82%, Exp5= 8.137%). However, the cross-validated Lasso regression showed, that after regularization none of the variance in

familiarity test performance can be predicted with these measures, with zero regularized weights for all five predictors for both experiments (**Table 3.1**).

Evidence of Global Structure Learning

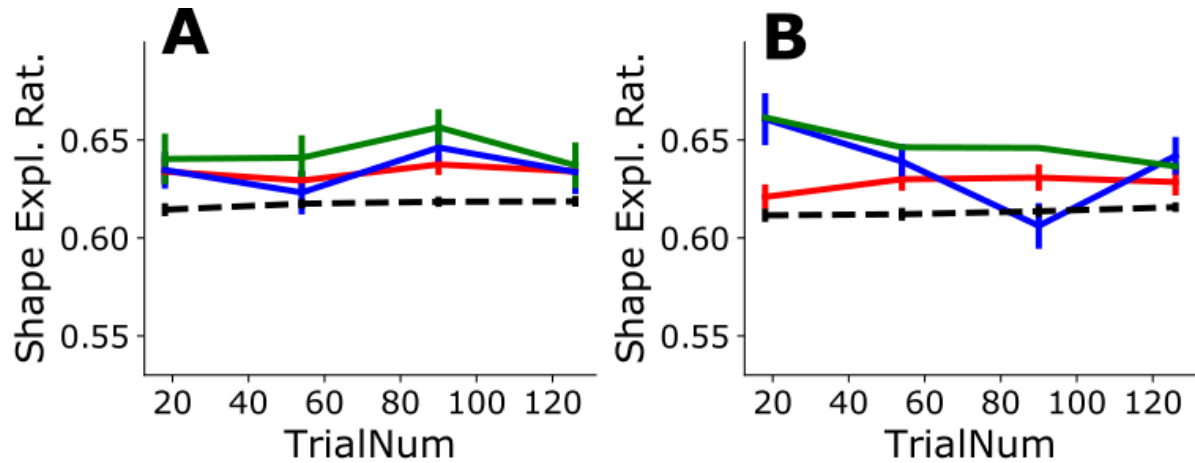


Figure 3.5. Global Structure influence in Exp 4 & 5. **A)** Shape Exploration Rate in Exp 4 **B)** Shape Exploration Rate in Exp 5. In both Experiments participants made more transitions within shapes than what would be expected by the shuffled chance. This suggests that they used the silhouette structure, to find more shapes. Note that the sample size in the blue group in B is modest ($N=8$), which could underlie the large drop in the third temporal bin. Errorbars: SEM

In order to measure the effect of knowledge about the global silhouette of scenes, we calculated the proportion of explorative transitions from a shape to another shape (Shape Exploration Rate as defined in the Data Analysis Section). We have found that in both experiments, explorative transition within shapes were more frequent than what would be expected by chance (Exp 4 $t_{39}=6.6321$, $p<.0001$; Exp 5 $t_{35}=4.7539$, $p<.0001$, **Fig 3.5**), showing that participants use the global structure of the scenes in order to find the stimuli more effectively (this is also supported by the analysis in the next section). This effect was not linked to individual learning success, as shown by the lack of significant correlation with performance on the final familiarity test (Exp 4 $r_{38}=.2229$, $p=.1668$, $p_{\text{exact}}=.1698$, $\text{BF}=0.5$; Exp5 $r_{35}=0.1115$, $p=.5113$, $p_{\text{exact}}=.5162$, $\text{BF}=0.25$). Meanwhile, the same measure was highly predictive of working memory task performance in Exp 5 ($r_{35}=0.4388$, $p=.0066$, $p_{\text{exact}}=.0074$, Bonferroni corrected $p_{\text{crit}}=.0125$, $\text{BF}=7.14$), but not significantly in Exp 4 ($r_{38}=0.2151$, $p=.1826$, $p_{\text{exact}}=.1807$, $\text{BF}=0.46$). This suggests that better usage of the global structure of scenes

during exploration helped performing well in the working memory task in Exp 5, while in Exp 4 working memory performance benefited from the pair regularities as described above.

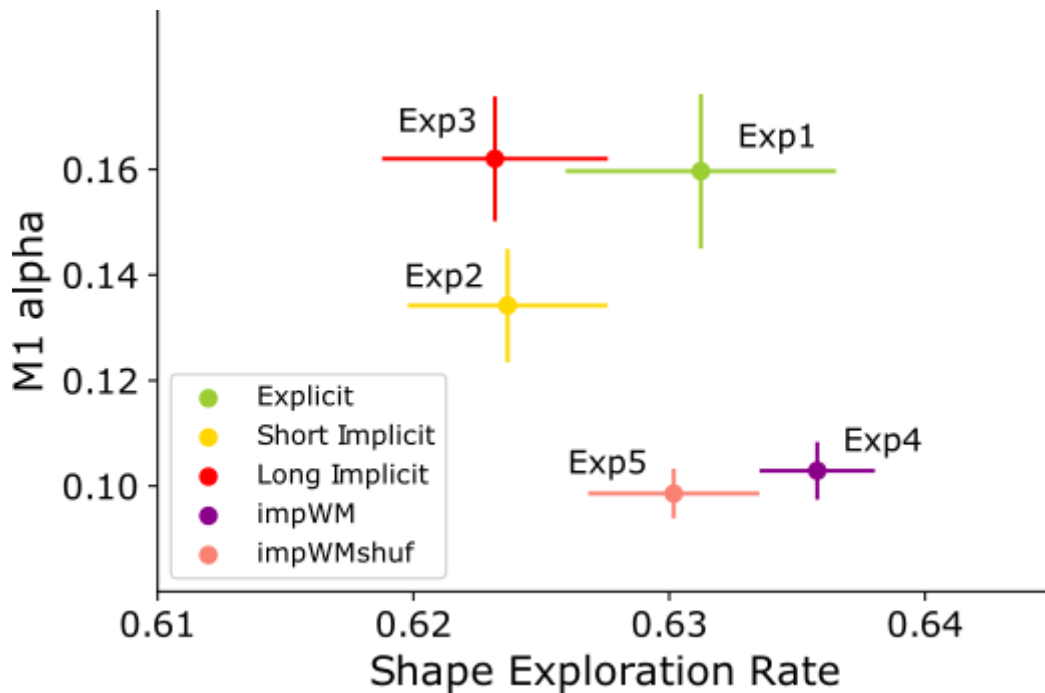


Figure 3.6: Pair vs. Global structure. Contrasting Measures Across Experiments 1-5. X axis: Shape Exploration Rate: reflecting global knowledge of scene statistics. Y-axis: M1 alpha representing pair influence on exploration behaviour. The dots are averaged across participants in Experiments 1-5. While pair influence was the lowest in the working memory experiments, the influence of global structure was the strongest. Errorbars: SEM for the two measures.

Contrasting Measures Across Experiments

The combination of measures used for quantifying the influence of the pair and the global structure show very different patterns across our five active statistical learning experiments (**Fig 3.6**). The global structure of the scenes, measured as looking proportion within shapes (regardless of pair structure), had the strongest effect in the Implicit Working Memory Experiment (Exp 4), while the influence of pairs was the weakest in the same experiment comparable to that of the control Exp 5, which had no pair structure to rely upon. Meanwhile, the pair structure had the highest effect on the Explicit pair test (Exp 1) and influenced participants equally strongly in the Long Implicit setup (Exp 3) with a notably less effect of global structure in the latter. Participants in the Short Implicit

experiment (Exp 2) relied somewhat less on the pair structure than in the Long Implicit experiment (Exp 3), but equally little on global silhouette information.

To test whether these are meaningful differences, we focused our analysis on the three experiments with pair structure and implicit instructions (i.e.: instructions not revealing the pair structure): Exps 2,3 and 4. We compared the results of Exp 4 with a working memory task to the combined data set of the other two Implicit Experiments having no working memory tasks. We found that while in Exp 4 the influence of global structure was significantly higher ($t_{118}=2.7811$ $p=.0063$, Permutation Test $p=.0064$, $BF=6.13$), the influence of pairs ($M_1 \alpha$) was significantly below the other two experiments ($t_{118}=3.7406$, $p=.0003$, Permutation Test $p=0$, $BF=86.65$). This confirms that while in the working memory experiment, participants were highly engaged with the scenes, they focused significantly more on the overall global structure of the stimuli than in Exps 2-3 resulting in high looking proportions between shapes. At the same time, these looking patterns were barely influenced by the pairs presented on any given trial, in a striking contrast with the strategies applied in Exps 2 and 3.

Discussion

We found that people can acquire spatial statistical regularities during active exploration with a parallel working memory task. We also found that statistics can be learned and utilized at different levels. In our task, two distinct types of statistical information were present: the general structure of overall silhouettes of scenes (Exp 4 & 5) and the particular pair structures of the components of the scenes (Exp 4 only). Each of these was uniquely linked to the parallel working memory task. Generic silhouette learning helped in Exp 5 to explore shapes above chance, boosting working memory performance. The pair structure provided an additional benefit for WM performance in Experiment 4, as suggested by two different measures: 1. temporal improvement in memory performance 2. better working memory predicted better statistical learning on the familiarity test. Nevertheless, in contrast to Exps 1-3, in Exp 4, we found no influence of the pair-structure learning on the visual

exploration patterns, only an effect of the global statistical structure of the presented scene silhouettes.

These results indicate an intriguing dichotomy: while learning of all useful statistics during training can be reliably confirmed via measurements of performance and correlations, the eye movements are influenced by these statistics only to the extent to which they are relevant for efficient execution of the task at hand. Specifically, both global shape and pair structures helped in remembering elements of the scenes, therefore, they both were acquired. However, the task of remembering as many individual shapes as possible did not require enhanced exploration within pairs. Once the pairs had already been learned, upon looking at one of the shapes, the location and identity of the pair was already known, hence the memory task could be solved without looking systematically to both shapes within pairs. In this case, uncertainty about the scenes would be reduced to a greater extent by looking at other locations, thus the within-pair attraction did not emerge in the eye movement patterns. The intriguing question of what is needed to find eye-movement-related effects within an unrelated parallel task will be addressed by the next experiment.

In principle, the increasing working memory performance in Exp 4 (**Fig. 3.3B**), could be explained either by the fact that memory capacity benefits from the statistical structure (Brady et al., 2009), or by a simple learning effect that people get better in the task over time. If the improvement were solely due to general learning with time (getting better at the task), then participants of Exp 5 would also be expected to show a significant benefit. The lack of such a significant improvement in the performances in Exp 5 (**Fig 3.3.B**) suggests that the pair structure was a crucial requirement for reliable improvement in the working memory task. This link between statistical learning and working memory performance is further supported by the significant correlation between working memory and familiarity test performance in Exp 4. The lack of such relationship in Exp 5 rules out the option that attention or motivation by itself could establish such a relationship between the two measures. Instead, a more parsimonious explanation of this correlation is that it is a consequence of statistical learning that, as a common underlying factor, boosts both working memory and familiarity test

performance. Therefore, our findings extend those described by Brady et al (2009), by showing that even when discovery of statistical structure requires segmenting complex novel scenes implicitly, working memory still benefits from learning about those regularities.

Experiment 6.: Eye-movements are attracted by temporal regularities

In Chapter 2, we found two conditions when learning statistical structures could influence eye movements: either there was an explicit task that directed the observer's attention and hence eye movements to the statistical structure in question (Exp 1), or there was no task, so after an extensive exposure, observers begin to benefit from the underlying structure (Exp 3 and to some extent Exp2). In Experiment 4 of the present chapter, we used an unrelated task, and found that while the pair structure was learned, it did not influence eye movements. What is the explanation of these opposite results?

One possibility is that such eye movement effects can only manifest themselves when participants are either informed explicitly about the structure or they are not controlled by any parallel task. An alternative possibility is that structure-related eye-movement effects could still be present even in a working memory paradigm, but, the setup of Exp 4 hindered their manifestation. It is conceivable that with simpler statistical structure and measurements of eye movement changes (e. g. increasing looking duration to certain areas), it would be possible to demonstrate structure-related eye movements even in the presence of an unrelated parallel task. This could be achieved with some parts of a scene being more random, while other parts containing more regularity: a statistical influence in such a paradigm would not require specific patterns of eye movements, only more looking at the area with regularities- as in preferential looking paradigms.

Demonstrating an effect of statistical structure on eye movements in such a setup would be significant because, instead of a rigid top-down relationship in which it is always the task that defines eye-movements, it would suggest a more dynamic balance between the complexity of the environmental statistics and the malleability of the active learning system. Depending on how much statistical structure there is in the environment, how developed the internal representation of this

structure and how complex the task at hand are, eye movement might be controlled to a different extent by the different factors, of which the task itself is only one.

A suitable method to test this has been described recently (J. Zhao et al., 2013). This study found that attention was implicitly attracted by visual event sequences containing transitional regularities, but the experimenters only assessed participants' response times and did not measure their eye-movements during the trials. Specifically, participants passively observed 4 continuously running parallel streams of visual shapes in four quadrants of the screen. The streams were frequently interrupted by a visual search array, appearing at the locations of the stimulus streams. The observer's task was to find a target and decide whether it was left- or right-oriented. Unbeknownst to the participants, images in three out of four visual streams were presented in a random order, while the fourth stream contained statistical regularities in the transitions between shapes. The regularities were similar to previous temporal statistical learning experiments (Barakat et al., 2013; Fiser & Aslin, 2002). The main finding of Zhao et al (2013) was that participants were faster to respond to the target if it appeared at the location with statistical structure. This was the first study to show that attention was biased by the mere presence of regularities in the input. However, an alternative explanation for the results of Zhao et al (2013) could be that participants did not attend more to the location of the structured stream, but the search targets were more easily distinguishable at the structured locations: if a sequence is more predictable, the search target could more pop-out at the regular location regardless of attention. The previously reported effect of contextual cuing relies on very similar measures – and very different statistical information- and it has been shown to influence eye-movements (M. S. Peterson & Kramer, 2001). However, that study has also been criticized on the basis of not really reflecting attentional guidance only faster response selection (Kunar et al., 2007).

To rule out such a pop-out based explanation and test if attention is biased by regularities, it would be informative to have a continuous measure of attention in a similar paradigm. An obvious candidate is eye-movements, as they are a common measure of overt attention orienting. Therefore,

if attention is influenced by regularities, it is expected that this effect will show up in overt attention and will be measurable by gaze location. Interestingly, statistical influences on eye-movements had only been shown with simpler manipulations, such as high probability of target appearance in one region of screen (Y. V Jiang et al., 2014) and not with complex transitional regularities.

The assumption underlying the attentional bias to regularities is that this effect subserves learning: for a sequence to attract attention, it should be learnable but not too easily learnable (Kidd, Palmeri, & Aslin, 2012). The low learning performance (around 54% correct responses on the familiarity test) reported by Zhao et al (2013) suggests that the sequences used in this paradigm are not too simple, so they likely attract attention for a long time. Given that a structured sequence is clearly less complex than a random sequence and that separating a structured sequence from a random one takes time, this paradigm is a good candidate to demonstrate an eye-movement-based attentional bias to regularities. If bias to regularities is related to learning, it should be possible to link individual biases in attention to learning. On the other hand, if statistics of the stimulus sequence is fully acquired by the observer, attentional biases might not persist any longer. One study investigated the complex temporal relationship between exposure time to statistical regularities and the emerging selective bias to attend more those regularities (R. Q. Yu & Zhao, 2015). This study used a modified version of the paradigm in (J. Zhao et al., 2013) by manipulating the location or presence of statistical regularities over time. The main finding of the study was that the attentional bias towards locations containing regularities persisted even after those regularities already vanished (R. Q. Yu & Zhao, 2015). However, the exact link between attention and learning was not established in this research, therefore, it is not known whether there is any direct relationship between the temporal bias to regularities and learning. One problem with the design of the Yu & Zhao (2015) study is that each participant might had a different learning curve, therefore, looking at the averaged learning across blocks of trials and participants might have hid many of the potentially learning related temporal variability in behavior.

To investigate whether the looking location preferred by the observer is influenced by the regularities at that location, we adapted the paradigm of Zhao et al. (2013). Because the complexity of the design and the low learning performance in the original paradigm, we decreased the number

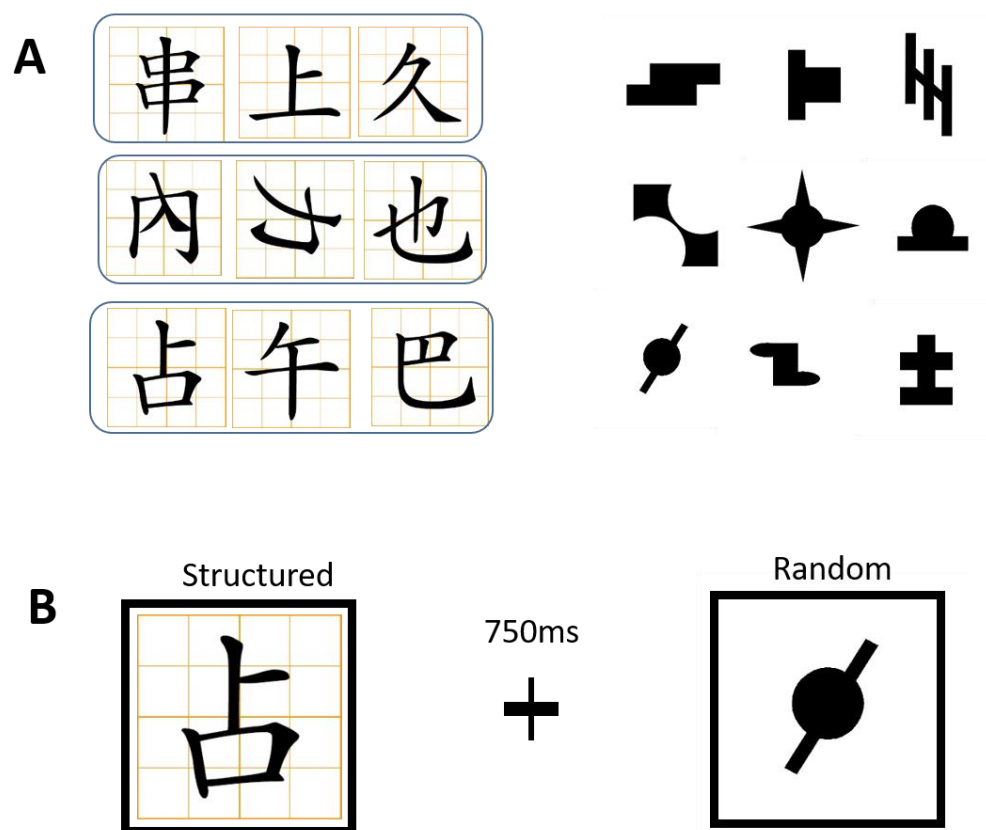


Figure 3.7. Stimuli and Structure of Exp 6 **A)** 2 set of 9 stimuli (1 set of shapes -1 set of Chinese characters, background grid was not present in the experiment) were used for the experiments. One of the sets was arranged into triplets, (top left). Stimuli within a triplet always followed each other. The other set of shapes appeared in a random order. **B)** The streams were presented in parallel at the two sides of the screen (shape presentation time and ITI both 750ms). During the intertrial interval, the squares were visible with a blank inside. The actual distance of the two squares relative to the size of the squares was larger than on this figure, covering an area of 22.4 visual angles.

of stimulus streams to one structured and one random. Additionally, we changed the parallel task, to one that encourages participants to pay attention to the stimuli. We expected that people will look more to the stimulus stream containing regularities, and we also expected that the strength of this effect will predict learning success on the subsequent familiarity test.

Methods

Participants

32 students (11 Male, 3 left handed, mean age 22.5, range 18-40 y) participated in the experiment. None of the participants spoke or studied Chinese. The selected sample size was more than 25% larger than sample size in previous similar paradigms (N=22-25 in R. Q. Yu & Zhao, 2015; J. Zhao et al., 2013) to gain more power.

Stimuli

Two sets of 9 shapes were used for the experiment. One set of shapes was selected from simple Chinese characters, the others were abstract shapes commonly used in visual statistical learning experiments (**Fig 3.7A**). For each participant, one set of shapes was always divided into 3 triplets of three shapes. We used a 15" screen 60Hz screen for stimulus presentation with a resolution of 1280*1024. Two streams of shapes were presented at the left/right sides of the screen in parallel on a gray background (**Fig 3.7B**). The center of the two streams was 23.75 cm apart, which covered 22.4 degrees of visual angle, since the viewing distance was 60 cm.

Procedure

Before the start of the experiment participants practiced the task (described in detail below) with different stimuli (drawings of common objects), which were presented in a random order at the left/right side of the screen. After the practice, the eye-tracker (Tobii 50 Hz) was calibrated with a 15-point calibration. Before the experiment started, participants were instructed to pay attention and perform the working memory task, but they were not given any specific instructions about where to look. They were not told about possible regularities in the stimuli. Participants used a chin-rest during the course of the experiment.

A random stream of shapes on one side and structured at the other was presented for 450 trials. The side (L/R) and the stimulus type (abstract shape vs Chinese character), which contained the regularities was counterbalanced across participants. The images were presented at 400 pixels left or right from center, and were surrounded by a black square, which was constantly visible and had a size of 160*160 pixels. Following the presentation times used by Zhao et al (2013), both stimulus presentation time and the inter-trial interval was 750ms.

In one of the streams, the shapes had a random order with a constraint that individual shapes could never repeat immediately and the number of times each shape was presented was balanced overall (9 shapes repeated 150 times). The structured stream consisted of triplets, meaning shapes within a triplet always followed each other in fully predictable order. Triplets could never repeat twice in a row immediately. We balanced the number of times a triplet was presented (overall 150 repetitions/triplet).

The presentation stream was interrupted 36 times with memory probes: these probes were pseudo-randomly interspersed after 8-12 trials. At each memory probe, a shape was presented centrally, and participants had to decide whether it was the immediately preceding shape at one of the sides (answer: left/right arrow for the two sides) or it was not present on the previous trial (answer: up arrow key). Each shape was used as a memory probe twice over the course of the experiment, once immediately after it was presented as the member of the stream (correct target) and once as a “distractor” when it was not present on the screen immediately before the memory probe (incorrect target). Since we were interested in eye-movements while participants observed the streams of uninterrupted stimuli, we decreased the number of the interleaved probes relative to the Zhao et al. (2013) study. Therefore, our working memory probe trials were not aimed at replicating the reaction time findings of that study. The primary goal of the memory probes was to ensure that participants’ attention was maintained at the two streams of shapes. Potentially, these memory probes could also be used to assess influences of statistical structure on working memory performance.

After the 450 trials of stimulus presentation and the memory probes were over, participants had to perform a two-interval forced choice familiarity test. Participants did not know in advance that there would be such a test. In this test, the three real triplets were tested against three foil triplets. The foil triplets were assembled from the stimuli of the structured stream, but the shapes were arranged in an order that never occurred during training. Shapes kept their temporal position within the foil triplets (eg.: if a shape was the first element in a true triplet during presentation, it was also a first element in the foil triplets, combined with shapes from different true triplets). Stimuli during the familiarity test were presented centrally and had the same presentation time and ITI as the stimuli during training (750 ms). Each triplet was tested against each foil twice resulting in 18 familiarity test trials. Half of test trials started with the true triplet, the other half with the foil. No feedback was given.

Data Analysis

No participants' data was excluded from the analysis of the study, but gaze data was only used from samples, for which the eye-tracker had a valid measurement of both eyes, as assessed by the software toolbox Talk2Tobii. This criterion meant that 92 +/- 6.6% of collected gaze data could be analyzed during stimulus presentation. If a trial had less than 25% valid data, it was excluded from further analysis, this resulted in the exclusion of an average of 2.5 trials per participant (.057% of trials).

There exist a large number of eye-movement-based measures of memory in the literature: number of fixations, fixation duration, proportion of fixations in area of interest, proportion of viewing time etc. (for a comprehensive list, see: Hannula, 2010). Since we had no clear hypothesis about which of these measures would be affected by our manipulation, we were seeking for a single measure that could assess in an integrative manner the different ways people might engage more with one stream of the stimuli than with the other one. Moreover, we used the recorded eye positions, instead of a

fixation-based approach, because defining fixations is problematic and thus introduces an additional subjective factor to the analysis. We chose calculating the mean looking location at each trial and averaged this measure over all trials for each participant separately to have an individual measure of overall visual preference.

Statistics were calculated using Matlab, SPSS and the JASP package (Wagenmakers et al., 2018).

Results

Familiarity Test Performance

Overall participants responded correctly to 52.43 +/- 12.4% of trials (**Fig 3.8A**), which is not significantly different from the chance of 50% ($t_{31}=1.11$, $p=.2752$ BF= 0.3316). Since the distribution of correct responses was not normal (Shapiro Wilk=.92, $p=.0206$), we used bootstrapping to obtain a 95% confidence interval, which includes chance: 48.1-56.8%. This confidence interval, however, also contains the performance reported by a previous similar study (54 +/- 7.9% test performance in Zhao et al 2013).

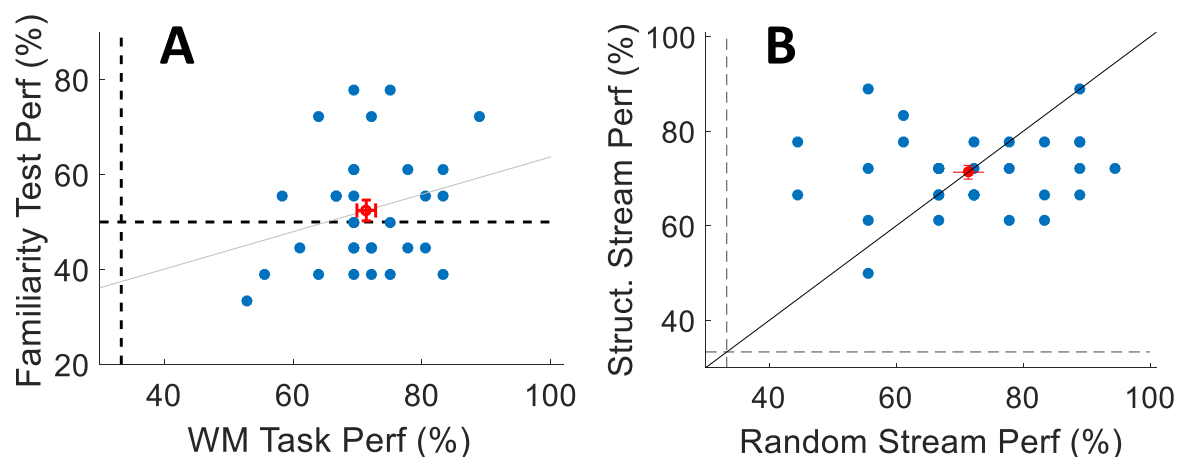


Figure 3.8. Task Performance in Exp 6. **A)** Working Memory (WM) and Familiarity Test Performance: The 1-back memory task (x-axis) and the final familiarity test (y-axis). Dashed lines, chance for the two tasks. Blue dots: individual data. Red: Mean+/-SEM of performance for both tasks. The performance on the WM task was highly above chance, the performance on the familiarity test was not different from chance. The relationship between the two tasks (grey line: least squares) was numerically positive, though not significant. **B)** Working memory performance depending on whether the stimulus was from the random (x-axis) or the structured stream (y-axis). The average performance on the two tasks was the same, but with a large individual variability in whether they do better with one type of stimulus or the other. (black solid line=identity)

Working Memory Task Performance

Participants' mean performance on the memory task was 71.35 +/- 8.21 % (**Fig 3.8A**) highly above the chance of 1/3 ($t_{31}= 26.1936$ $p=0$). Although knowing the statistical structure of the triplets could make the working memory task easier, there was no significant correlation between performance on the memory task and on the familiarity test ($r_{30}=.26$ $p=.1484$, $BF=.598$). Performance on the stimuli coming from the random or the structured stream was identical on average (**Fig 3.8B**, mean difference equal=0, $t_{31}=0$, $p=1$, $BF= 0.19$). The difference in performance between the structured and random stream shapes was also not affected by the side ($F_{1,28}=0.432$, $p=.5166$ $\eta^2=.015$) or whether the Chinese or the abstract shape containing stream contained the regularities ($F_{1,28}=0.276$, $p=.6033$ $\eta^2=.01$). This shows that in the current paradigm the performance in the working memory task was not affected by representations of statistical regularity.

Descriptives of Eye-Movement Data

The trial length of 750ms allowed participants to saccade to both shapes at the two sides during a trial, but they could also look at only one of the stimuli or stay at the fixation cross while inspecting stimuli at the two sides peripherally. Since participants were free to pick between these strategies, there was considerable individual variability in their behavior. The most common strategy was to look at only one of the stimuli during a trial (43.26 +/- 17.73 % of trials), closely followed by looking at both stimuli (42.49 +/- 21.84 % of trials). Participants did not look at any of the two cells containing the stimuli on 13.47 +/- 21.31% of the trials.

Statistical Structure Influence on Eye-movement data

To measure the influence of the stimulus stream's statistical structure on average looking location, we used a balanced 2*2 ANOVA, where the two factors were statistical structure (structured vs random) and stimulus type (abstract shape vs Chinese character). We have found that both the

stimulus type and statistical structure had a significant main effect (Structure: $F_{1,28}=4.5$, $p=.0429$ $\eta^2=.138$; Stimulus Type: $F_{1,28}=6.82$, $p=.0143$ $\eta^2=.196$) with no significant interaction ($F_{1,28}=.02$, $p=.8896$, $\eta^2=.001$) (See **Fig 3.9**). This indicates that position of looking at each trial was attracted by the statistical structure of stream. To see whether this overall effect was present from the beginning, we repeated the same analysis for the first 30 trials of the data and found no significant effects for either factor (Structure: $F_{1,28}=1.407$, $p=.2450$ $\eta^2=.048$; Stimulus Type: $F_{1,28}=.53$, $p=.4722$ $\eta^2=.019$, Interaction: $F_{1,28}=1.692$, $p=.2043$ $\eta^2=.057$). The absence of the effect at the beginning of the experiment suggests that the attraction of gaze emerged over time due to learning and it could not be explained by an unrelated bias.




Figure 3.9: Statistical Influence on Eye-movements. Average looking location was influenced by both statistical structure and looking location with no interaction between the two factors. This demonstrates that preferential looking location was attracted by statistical structure. (X- axis: St - Statistical Structure, Ch=Chinese Shapes, L-Left, R-Right) (Error bars: SEM)

Unrelated to our main research questions, the average looking location was closer to the left side of the screen (**Fig 3.9**, Mean Looking Location from Center: -28.01 ± 81.61 pixels, $F_{1,28}=4.794$, $p=.037$, $\eta^2=.146$). This left bias in the data is not surprising, since the participants could look at the shapes presented at the two sides in the order they preferred. Starting the inspection on the left more often is consistent with a well-known tendency to proceed with visual search from left to right (Spalek & Hammad, 2005).

Discussion

Using a temporal statistical learning paradigm with different areas of the screen containing structured and random streams of stimuli, we have shown that people prefer to look more to regions that contain statistical regularities. The lack of learning in the subsequent familiarity test suggests that looking location bias can be a sensitive measure that indicates an influence of visual stimulus statistics even before that knowledge could be detected in a familiarity test. This is somewhat surprising as familiarity test is usually considered to be a sensitive measure that can show memory in situations where explicit recall would be at chance (Yonelinas, 2001). In contrast, infant studies have used preferential looking to assess otherwise hard-to-measure knowledge for a long time (Spelke, 1985).

Our findings expand the findings of the previous study of (J. Zhao et al., 2013) suggesting that the attentional bias towards regularities they reported can be detected in eye-movement behavior. This confirms that previously reported reaction time results indicating attentional attraction were not due to a pop-out mechanism because of the violation of a learned sequence or easier response selection (Kunar et al., 2007), but due to an overt attentional bias that can be measured by eye-movements.

Since eye-movements represent a continuous measure of oriented attention, they could be used for measuring potential changes in attentional biases over-time (R. Q. Yu & Zhao, 2015). Our original goal was to link the individual variability in the temporal changes in data to individual learning outcome. Unfortunately, the low familiarity test performance did not allow such an analysis. The lack of learning might seem surprising given the low number and highly predictive stimulus order in the structured stream. Moreover, the participants were clearly engaged with the stimuli enough to perform a memory task well above chance. However, the presentation of a parallel random stream could completely overshadow the discovery of the temporal regularities in the structured stream. This detrimental effect of mixing random and structured stimuli has been shown in previous studies (Jungé, Scholl, & Chun, 2007), albeit in a set-up where the random stimuli preceded the presentation

of structured input (as if predisposing people not to look for structure). It appears that parallel presentation of structured and random input has a similar effect.

We used a parallel working memory task, but we did not find an influence of stimulus statistics on working memory performance based on two measures. First, performance on the 1-back task was identical regardless whether the stimulus was from the structured or the random stream. This is surprising since it shows that the spatial attention bias, unlike the reaction time advantage shown by Zhao et al (2013), does not influence task performance across the spatial locations. Second, unlike in Experiment 4 with the spatial paradigm, we found no correlation between working memory and familiarity test performance, probably because of the low overall learning performance. These results have to take with caution though as the effects were extremely weak, and stronger learning and/or more sensitive measurement of the interaction might change the picture.

We have shown, that the presence of regularities attracts overt attention, even in the absence of clear evidence of learning those regularities. This confirms that eye-movement measures are a good candidate for investigating implicit learning processes, which can be more sensitive than a parallel working memory task. The low overall performance in the familiarity test, however, suggests that this specific paradigm is not suitable to investigate individual learning patterns without further improvement.

General Discussion

Across a set of visual statistical learning experiments with a parallel working memory task, we have investigated the relationship between eye-movements, working memory performance and statistical learning. In Exp 4, we showed that implicit knowledge of statistical structure can emerge in parallel with improvements in a parallel working memory task that is related to that statistical structure, even without any detectable effect of the pair structure on eye movements. In contrast, in Exp 6, we have found that presence of regularities can affect eye-movements in the absence of recallable learning at the subsequent familiarity test. This supports earlier findings that eye-movements could detect memory processes that otherwise could not be recalled (Hannula, 2010; Hollingworth et al., 2001).

There are profound differences between the paradigms we used that could underlie these opposing results. While in both experiments participants were free to look wherever they wanted within the screen, in Exp 6, such a variation in looking position had no further functional relevance beyond side preference, while in Exp 4 it determined stimulus presentation gaze-contingently. It is possible that eye-movements are a more sensitive measurement of memory if they are freely guided and not tied to the stimulus presentation. This would contradict a previous study, that showed probabilistic influences on eye-movements only manifest if target presentation is gaze-contingent (Paeye et al., 2016).

To test if the presence/lack of gaze-contingent manipulation is partially responsible for the different results across experiments, the manipulation could be reversed: Exp 6 could be tested in a gaze-contingent set-up, while Exps 4-5 could be tested with a non-gaze contingent presentation. In theory, it is possible that such a manipulation would reverse the findings. However, since there was no “recallable” learning in Exp 6, making that paradigm gaze-contingent could further harm learning effects. Another option would be to test Exps 4-5 without a gaze contingent manipulation, with the

whole scene visible at once. The problem with such an approach is that seeing shapes of pairs together by utilizing the visual periphery could completely abolish the learning-related influences on the eye-movements we observed in Chapter 2 as well as the global structure influences described in the current chapter. Resolving the contradiction above requires the future development of a non-gaze contingent eye-movements measuring method that can be used to track spatial statistical learning.

To sum up, we have shown that cognitive processes affected by the acquisition of statistical information from visual scenes depend on both the type of regularities in the scenes and the requirements of the task executed in parallel. Preferential looking location can indicate the discovery of statistical regularities during performing a working memory task, even before that knowledge could be recalled on a subsequent familiarity test, suggesting that eye-movements can be a sensitive measure of learning. On the other hand, we also found an opposite outcome, in which the learning of spatial regularities did not influence eye-movements, while it was linked to working memory through increased capacity that indicated better learning performance.

Chapter 4

Effects of Past Probabilities on Perceptual Decision Making: Exps 1-4

Summary

In a set of four visual discrimination experiments we investigated how past stimulus appearance probabilities influence perceptual decision making. We found that human perceptual discrimination was influenced not only by short-term, but independently and equally strongly by long-term changes in stimulus probabilities. The long-term effects could easily overwrite the short-term ones and could elicit a counterintuitive bias against making the locally more rational choice. In Experiments 1-4, we show that same local stimulus statistics across many dozens of trials can elicit very different preferences depending on the long-term experience, and more specifically, these effects depend on the relative values of short- and long-term summary statistics capturing stimulus probabilities.

Introduction

How do we integrate uncertain sensory input with our expectations from the statistics of long-term experience? An intriguing proposal, which treats perception as an unconscious Bayesian inference (Kersten et al., 2004; Weiss, Simoncelli, & Adelson, 2002) posits that the past serves as a prior basis, which is combined with the likelihood of uncertain sensory evidence to reach a rational interpretation of the current input. This approach has been rapidly gaining influence, as a number of seminal papers demonstrated that human behavior can be quantitatively captured within the Bayesian framework. Specifically, people have been shown to integrate uncertain sensory evidence with probabilistic expectations in a manner that can be described as near Bayes optimal (Ernst & Banks, 2002; Körding & Wolpert, 2004). However, the literature is still divided on this issue as there are reported instances of human behavior falling short of Bayes-optimal integration (Ackermann & Landy, 2014). This discrepancy led to proposals that extended the benchmark for optimality by imposing realistic computational limitations (Vul, Goodman, Griffiths, & Tenenbaum, 2014) or suggested to switch focus away from optimality altogether (Rahnev & Denison, 2018).

There exists a different line of research that goes beyond the canonical approach of combining long-term statistics and current sensory input and demonstrates an interaction between long-term experience and recent statistics, which together influence momentary perception (Chopin & Mamassian, 2012). The main idea of this approach is that the difference between recent input and long-term experience will determine current expectations about the momentary input, as current statistics should resemble prior models. While the proposal is intriguing, it has proven to be controversial (Maus et al., 2013). For instance, Bohil and Wismer (2014) have shown that while momentary decisions are biased by the base-rate of long-term experience, there is no interaction between past and momentary statistics. Using reaction times, an additive interaction between long-term and recent experience has been described: people are influenced by both long-term and recent experience and they are the fastest if recent statistics resemble long-term expectations (Wilder et al.,

2013). While in the last few years there has been a surge of interest in short-term serial influences of past stimuli and decisions (Cicchini et al., 2017; Fischer & Whitney, 2014; Fritsche et al., 2017; Liberman et al., 2014), there are virtually no studies exploring whether these effects have anything to do with the above described long-term influences.

The impact of long-term experience likely depends on whether conditions are stable: in volatile environments people should care less about long-term experience, since they are less predictive of the momentary conditions. Indeed, using a reinforcement learning framework, several studies have shown that people can adjust their reliance on past experience, depending on whether they perceive the environment as changing (Behrens et al., 2007; Nassar et al., 2010). Furthermore, during sequential perceptual decision making, people are more prone to adjust decision boundaries if conditions are volatile (Norton et al., 2017). However, there is no agreement about the best model to capture human behavior in such perceptual decision-making tasks in volatile environments, as very different models emerge as best predictors of behavior, even within the same participant (C. Summerfield et al., 2011). The reason why sometimes a simple working memory model captures the experimental data better, while in other cases a reinforcement learning or a Bayesian model prevails, remains to be resolved (Norton et al., 2017; C. Summerfield et al., 2011).

This brief literature review demonstrates that while both long-term and recent stimulus statistics are important factors in perceptual decisions, their exact relationship is yet unclear. To address this knowledge gap in the literature we conducted a series of visual discrimination experiments. The general logic of the experiments was to provide a training block with feedback to form long-term probabilistic expectations, followed by a test block (feedback absent) that both set up the observer's recent experience and assessed his/her decision-making performance. Across a series of three experiments, the test session had the same unbalanced structure (one of two alternatives appeared with a higher probability than the other), but we manipulated the training probabilities (i.e. the appearance probability of the objects during training). In a fourth experiment we assessed the role of long-term probabilities when the momentary probability conditions were balanced. Our design

allowed us to separate long-term and current statistical influences and to validate our main hypothesis: long-term expectations from the training block should strongly influence the participants' decision during the test session and the extent of this effect should depend on the change between appearance probabilities during training and test.

Methods

Stimuli and Procedure

Two shapes out of a set of 11 were randomly selected for each participant to serve as the discrimination stimuli. On each trial, the stimulus of size of 204*204 pixels (≈ 4.7 visual angle) was presented centrally (circa 4.7 visual angles) on an iMac 27" (2560*1440) using Psychophysics Matlab toolbox. Participants watched a screen in a dimly lit room at a viewing distance of 60 cm and used the left/right buttons of the keyboard to provide responses as to which shape they saw on the given trial. Instructions emphasized accuracy over reaction times but did ask for timely responses as there was an upper limit of four seconds to respond. The instructions did not mention stimulus probabilities or changes in the task structure. Trials were presented on a grey background display within a blue "box", a 256*256 pixels large blue square, spanning approximately 5.7 visual angles. On each trial, one of the two shapes was presented embedded in Gaussian noise for 200 ms, while the thin frame of the box (12 pixels wide) remained visible (**Fig 4.1A**). After the stimulus disappeared, the center of the box turned white until the participant responded. After the response, the box reverted to blue until the next stimulus. The interval from the response to the next stimulus (RSI) was sampled randomly from a normal distribution with mean=1100 ms and SD=100 ms.

During the training block, negative feedback was given after each mistake (in the form of red exclamation marks) and no positive feedback after correct responses. During the test block there was only feedback if participants made a mistake on the 1/8th lowest noise trials in order to maintain attention (performance was over 90% in these trials, totaling to approximately 1% of test trials with feedback).

At each trial, varying levels of Gaussian noise were added to the grayscale stimulus (**Fig 4.1B**). The training started at a low noise level, and the variance of the Gaussian noise was gradually increased with a “2up/2down” adaptive staircase procedure. The training lasted for 200 trials, to have an estimate of discrimination threshold. After 15-30 seconds of break, a 300-trial long test phase followed, where the Gaussian noise was sampled uniformly-randomly between low noise and the threshold reached during training.

The appearance probability (AP) of the two shapes was manipulated across the four experiments. Three experiments (Exps. 1-3) had a biased 65% AP during test, meaning that one of the shapes appeared 65% of the time during the trials. These three experiments differed only in their AP during training phase (Exp. 1=50, Exp. 2=65, Exp. 3=75%, **Fig 4.2A**). A fourth experiment used a biased training with 75% AP during training but a balanced 50-50% AP during test.

Participants

80 Hungarian students (18-30 year old) participated in Exps. 1-4 (20 in each experiment) and received monetary compensation. The participants gave informed consent before the start of the experiment and were unaware as to the purpose of the study. Four additional students completed one of the experiments but were excluded: three participants (1-1-1 from Exp. 1,2,4) due to chance level performance (below 60% on the easiest 1/3rd of trials), suggesting that they either did not pay attention or did not understand the task. One additional participant was excluded from Exp 1 because of a very strong bias to give the same response all the time (3 SDs above the mean).

Data Analysis

Psychometric Curves

To generate psychometric curves, we fitted a sigmoid function to the response proportions divided into three noise bins, according to the equation below:

$$\Psi(x; \alpha, \beta, \lambda) = \lambda + (1 - 2\lambda)F(x; \alpha, \beta)$$

Here, Ψ represents the proportion of "frequent" responses, and x is the noisy stimulus, encoded as (-3,-2,-1) for rare stimuli with low, medium, and high noise, and (1,2,3) for frequent stimuli with high, medium, and low noise. F is the cumulative normal distribution with mean= α and SD= β . The parameter α is the threshold (or PSE, point of subjective equivalence), representing the bias of the observer. The parameter β is the slope, which reflects perceptual sensitivity. These two parameters were fitted to the individual data, using Matlab and the *fminsearch* function. We fixed the stimulus independent lapse rate (λ) at .01, similarly to previous related work (Fritsche et al., 2017).

Logistic Regression Modelling

Since the basic psychometric curve ignores inter-trial influences (Fründ, Wichmann, & Macke, 2014), we turned to logistic regression models including past events as predictors. This can be interpreted as an extension of the psychometric curve with inter-trial influences, used in many previous papers investigating past influences on perceptual decisions (eg: Braun et al., 2018; Norton et al., 2017). We fitted the logistic regression individually, to investigate the key factors influencing each participant's responses. This analysis was performed in Python using the *scikit-learn* library (Pedregosa et al., 2011). Since the stimulus presented on a given trial has the strongest influence on the response, we used a model containing the current stimulus only as a null model. Stimulus noise was encoded in the model on a linear scale¹⁴ between -1 and 1 and not binned as for the psychometric curves. We compared the null model based on the present stimulus to more complex models that use various combinations of the bias term representing all long-term effects, together with the previous response and previous stimulus in one or two of the previous trials as predictors. Calculated values of log likelihood, cross-validated likelihood, Bayesian Information Criterion (BIC), Akaike Information Criterion (AIC), and the small sample corrected AIC (AICc) were used to compare goodness of fit

¹⁴ see additional analysis on justifying the linear encoding of noise in Appendix Text B.1.

across models. For cross validation, we repeatedly fitted the model on batches of 95% of trials that were randomly selected and tested it on the remaining 5%.

We performed a simulation analysis in order to determine which measure of goodness-of-fit described above was most likely to distinguish true generating model from the alternative models. For this, we simulated responses of the 12 candidate models to the stimulus sequence presented to each participant during Exps. 1-7 (see details of Exps. 5-7 in Chapter 5) and evaluated the models' performance. The 12 models included all combinations of the following regressors: Present Stimulus (null: 0), Bias (B), Previous Decision (D1), Previous two Decisions (D2), Previous Stimulus (S), and the full model with all five parameters (F). During the simulations, first, parameter values for each simulated participant were sampled from a normal distribution with mean and SD of the model fitted to the empirical data set. Next, the 12 candidate models were fitted to each data sequence and the four measures of model fit were calculated for each measure (AIC, AICcorr, BIC, Cross-validated likelihood-CVLL, See Appendix Figure B.1.). The entire fitting process was repeated 20 times. For each model fit measure, based on all simulated participants' data, we calculated the percentage (from 20*12 runs) with which the best fitted model was the true generating model: the AIC/AICc measures based either on total or on the mean of individual likelihoods performed at 100%, while CVLL and individually-fitted BIC both performed at $\approx 71\%$, and BIC fitted on total likelihood performed at $\approx 33\%$ (chance=8.33%). Thus, our analysis based on recovering parameters of simulated data showed that AIC/AICc was the most accurate measure for finding the true model on our dataset (Suppl Fig 4). Therefore, we applied this measure for performing the actual model selections in this study by fitting each candidate model individually to the full dataset and selecting the best model according to AIC. We note that the best model based on the cross-validated likelihood measure was very similar to that obtained with AIC.

To compare the relative significance of the various predictors in any given model, we standardized the logistic regression weights by multiplying the absolute value of the weights with the SD of the predictors (Menard, 2004). This ensured that the beta weights did not depend on the scale of

encoding the predictors and their predictive power could be directly compared. This method was used instead of the customary way of standardizing the data to z-scores before fitting the model, because standardizing would remove the bias, which was one of our main measures of interest. We validated our standardized method by using a range of different values to encode the different predictors and confirming that the relative beta values remained the same. The regression weights across studies were compared by standard one-way ANOVA significance test. For descriptive statistics on the fitted model parameters for Exps. 1-7 see Appendix Table B.1.

Recent Past Influences

To analyze the interaction of the recent past and long-term probabilities in the raw data, we needed a measure of the recent probabilities. To this end, we binned the response of each trial into one of four bins depending on the number of times the frequent shape was selected as the response in the 3 trials preceding the current one (0,1,2 or 3 frequent responses during the last 3 trials). To perform this analysis, we used a mixed ANOVA with recent past as the repeated factor and the experiments as a between subject factor. Two participants in Exp 3 were excluded from the analysis due to a missing data point in one of the bins (because they did not have a sequence of 3 rare responses throughout the 300 trial test session).

Confidence Intervals

Confidence intervals were obtained by bootstrapping across participants 5000 times, calculating the mean of each bootstrapped sample, and taking the 2.5% and 97.5% percentile values of the obtained distribution of bootstrapped means.

Reaction Time Analysis

To test our measure of interest, we compared median reaction times for frequent vs. rare responses with t-tests. When the data presented strong outliers in the data (**Fig 4.4D**), we also report nonparametric tests (Wilcoxon signed rank test). Descriptives for all experiments alongside both *t*-test and Wilcoxon signed rank test can be found in Appendix Table B.5. for Exps. 1-7.

Bayes-Factors

To obtain Bayes factors (BF) for between group comparisons we used an uniform distribution of possible differences between groups (Dienes, 2014). For comparing decision bias across groups, we used a maximum difference of 1, since the group values change in the range of $[-0.5 +0.5]$. However, as the value of the BF depends on the maximum possible difference, we re-calculated the Bayes factor with smaller (.7) and higher value (1.5) for maximum difference, finding support for the same conclusions. The reaction time analysis of differences across groups used a 250 ms maximum difference across groups. To calculate BF for 2-sampled comparisons, we used the BayesFactor package with a non-informative Jeffrey-Zellner-Siow prior on the value of effect size (Rouder et al., 2009). We used the convention that BFs smaller than 1/3 provide evidence for the null, between 1/3-3 are inconclusive, while BFs larger than 3 provide evidence for the alternative hypothesis.

allowed us to quantify the relative contribution of each factor on perceptual decision making. As it should, the current stimulus was the most important determinant of momentary choices (**Fig 4.1D**, Appendix Table B.1). From the other predictors, the bias term had the largest influence on choices, followed by the previous two decisions. In what follows we will demonstrate how this strong decision bias can be explained by long-term stimulus probability changes of the training and test block.




Figure 4.2: Experiments 1-3. **A)** Structure of Experiments 1-3, experiments differed in training AP and had identical test AP. **B)** Correct Response Proportions as a function of noise (Color code as in A, Errorbar: SEM) P of correct answers was highly noise dependent, falling from over 80% at low noise, to about 60% at high noise. **C)** Response probabilities (dots+SEM) and Psychometric curves during test were strongly manipulated by training probability in Exps. 1-3 (color code as in A). X-axis represents the rare and frequent stimulus at different noise levels. Y-axis, the probability of choosing the frequent shape. Dashed line: unbiased prediction (with average empirical sensitivity).

Descriptive statistics of Behavior

On average, participants responded correctly on about 70% of trials (Mean +/- SD Exp. 1 = 70.01 +/- 6.45%; Exp. 2 = 70.96 +/- 6.58 %; Exp. 3 = 72.58 +/- 8 %). Performance was highly noise dependent changing from about 80% performance at the lowest 1/3th lowest Noise level to about 60% at 1/3 highest noise trials (**Fig 4.2B**).

Sensitivity, as measured by the slope of the Psychometric function, was not different across experiments (**Fig 4.2C**, ANOVA $F_{2,57}=0.235$ $p=0.791$ $\eta^2=0.008$). Median response times from stimulus onset were 796 +/- 156 ms in Exp. 1, 709 +/- 159 ms in Exp. 2 and 784 +/- 191 ms in Exp. 3.

Strong long-term influences on decisions

We have found that long-term probabilities strongly influenced decision preference: the bias of the psychometric curve (PSE) was significantly modulated by the summary statistics of the training block ($F_{2,57}=11.213$, $p<.0001$, $\eta^2 = .282$, **Fig4.2C**). Analyzing the effect experiment-wise, we have found that after balanced training conditions, increasing the AP to 65% (Experiment 1) resulted in a counterintuitive decision preference towards the rare element (PSE shift: $t_{19}= 3.4228$, $p= .0029^{15}$). The effect cannot be a consequence of perceptual adaptation or an alternation bias, as it is also highly significant in our logistic model, which accounts for short-term influences (Logistic Model Bias: $t_{19}=4.431$, $p =.0003$, $BF=109.3^{16}$) (**Fig 4.4A**). This effect emerged over time since it was absent in the initial 30 trials ($t_{19}=0.3105$, $p=.7595$, $BF=.24$, see also **Fig 5.4A**) and was most likely the consequence of the difference from long-term experience. This is further supported by Exp 2, as when the AP-s did not change between training and test, there was no overall decision bias (PSE shift: $t_{19}=1.4668$ $p= .1588$, Model Bias: $t_{19}=0.6227$, $p =.5409$, $BF=0.28$). Furthermore, if the change from training was in the opposite direction (in Exp. 3), this resulted in the opposite effect, a highly significant preference for the frequent element (PSE shift: $t_{19}=3.3142$, $p= .0036$; Model Bias: $t_{19}=2.7457$, $p = .0129$, $BF=4.16$). These long-term effects were stable, with no significant change in bias between the first and second half of test in any of the three experiments, demonstrating a lasting effect of long-term probabilities (Exp. 1 $t_{19}= 0.302$, $p=.7659$ $BF= 0.24$, Exp. 2: $t_{19}=1.5843$, $p=.1296$ $BF= 0.68$, Exp. 3: $t_{19}= 0.9889$,

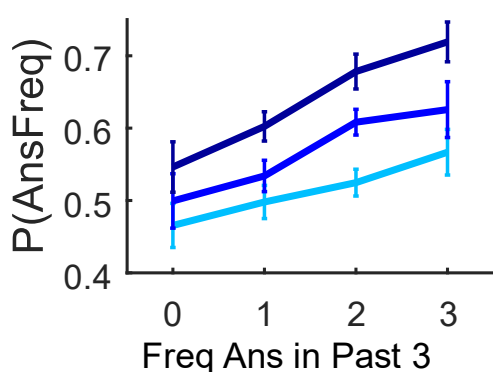


Figure 4.3. Immediate past and long-term probabilities had an additive influence on decisions. The more often a shape was chosen in the recent past (last 3 trials-x-axis) the higher the probability of choosing it on the current trial (y-axis). This effect was very similar in Exps 1-3, but was shifted by the long-term probabilities, without an interaction between these two factors, showing that short-term influences and long-term biases influence perceptual decisions independently. Color code: Experiments 1,2,3=light, medium, dark blue. Error bars (SEM)

¹⁵ Descriptives on PSE in Appendix Table B.2.

¹⁶ Descriptives on Model Bias in Appendix Table B.3

$p=.3351$, $BF=.36$).

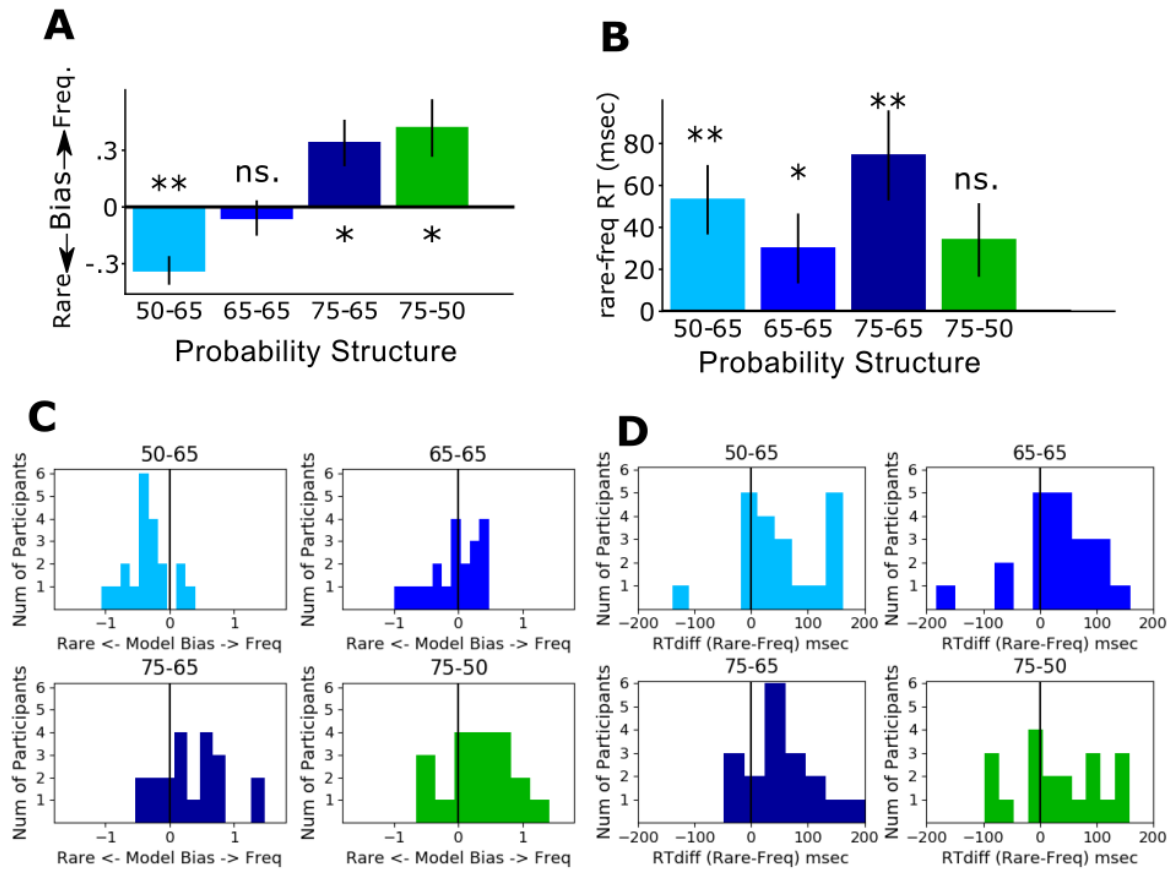


Figure 4.4. Contrast of Decision Bias and Response Times. **A)** Mean model Bias from Logistic Regression for Exp 1-4 (x-axis). The change from training to test in appearance probabilities was the main determinant of choice bias **B)** RT difference (rare-freq.), in all experiments median “frequent” responses were faster. Notice the contrast for Exp 1 (50-65) between the RT and decision bias. The RT effect was the weakest in Exp 4, where the momentary p-s were balanced, and “frequent” was defined based on training. **C)** Model bias distribution for Exps. 1-4. **D)** RT difference distributions for Exps. 1-4. Because of the outliers in RT, we report non-parametric test results as well. (Error bars=SEM)

Independence of immediate past and long-term influences

Decisions were also influenced by recent choices, with a tendency to repeat recently frequent answers. We analyzed the probability of a frequent response, given the occurrence rate of frequent answers in the immediately preceding three trials, finding very similar patterns across experiments. Analyzing this effect with a mixed ANOVA showed that both recent past probability (within subject $F_{3,165}=14.883$, $p<.0001$, $\eta^2=.213$), and long-term probabilities (between subject $F_{2,55}=10.375$, $p=.0002$,

$\eta^2=.274$) had a significant influence on response probabilities with no interaction between the two factors ($F_{6,165}=4.52$, $p=.8426$, $\eta^2=.016$) (**Fig 4.3**). The logistic model further confirmed the independence of recent and long-term influences: weights of the immediately previous responses were not different across experiments 1-3 (One-way ANOVA, $F_{2,57}=1.0424$, $p=.3592$, $\eta^2=0.035$).

Comparing Long Term influences on Decisions and Response Times

Interestingly, reaction times showed a very different pattern from the decision preferences described above. In experiments 1-3 with 65% test AP, median reaction times were faster to the currently more frequent element, suggesting higher sensitivity to momentary probabilities and weaker long-term influences (**Fig 4.4B**, Exp. 1 RT difference Frequent-Rare: $Z = 3.0239$, $p=.0025$, $BF= 9.8$ Exp 2: $Z= 1.9786$, $p=.0479$, $BF=0.9$ Exp. 3: $t_{19}=3.4559$, $p=.0026$, $BF=15.8^{17}$).




Figure 4.5. Experiment 4. **A**) Probability structure: the experiment had a balanced test, but a biased training (75%). **B**) Probability of Correct Responses as a function of Noise. **C**) Response proportions by noise and stimulus type (dots) and Psychometric curve reveal a strong bias toward the stimulus that had been frequent during training. Dashed line: unbiased prediction with average sensitivity. Exps. 1-3 in opaque in the background as comparison.

Balanced Test Sanity Check- Exp 4

In an additional experiment (Exp. 4), we assessed the role of past probabilities, when the conditions during test were balanced (**Fig 4.5**). In Experiment 4, there was 75% AP training, which was identical to Exp 3. This was followed by a balanced test where both stimuli had equal probability of

¹⁷ Response Time descriptives in Appendix Table B.5.

occurrence. Otherwise, the paradigm was identical to Experiments 1-3. During the test we assessed whether participants prefer the stimulus that had been frequent during the preceding training session. We have found a strong preference toward the previously frequent element, as shown by both the Psychometric Curve ($t_{19}=4.8758$, $p=.0001$, **Fig 4.5C**) and the logistic model bias ($t_{19}= 2.7406$, $p=.013$, $BF=4.12$ **Fig 4.4A**). In line with Exps. 1-3, this effect was stable with no significant change between the first and second half of the test session ($t_{19}=0.4865$, $p=.6322$ $BF=.26$).

Assessing reaction times in Exp 4, we found that participants were faster in giving “frequent” responses, but not significantly so (**Fig 4.4B**, $Z= 1.6053$, $p= .1084$, $BF=1.01$). This shows that despite the fact that participants demonstrated the strongest decision bias in Exp. 4, among Exps. 1-4, this was not sufficient to elicit a robust reaction time difference due to the balanced momentary probabilities. This result supports the claim that, in comparison to decisions, influence of past probabilities is weaker on reaction times.

Discussion

Using a set of perceptual discrimination experiments, we demonstrated that past stimulus probabilities strongly influence perceptual decision making: the exact same base-rate appearance probability could elicit an opposing decision bias, depending on differences in long-term experience.

While base rate influences have been previously used in perceptual categorization studies, only an attractive influence of past probabilities has been described: participants were biased toward choosing the previously more frequent category (Bohil & Wismer, 2014). Our Experiment 3 and 4 replicates that unsurprising finding, but also extends it by showing that depending on the direction of probability change, a preference for even the rare category can emerge. The underlying cause of the complex pattern of results we observed could be the way in which humans integrate supervised and unsupervised information, an aspect that is largely ignored in most studies of perceptual decision making (an exception: Gibson, Rogers, & Zhu, 2013). A semi-supervised sequence of information

could lead to the surprising interactions between long-term and current statistics. The supervised parts could help forming a stable internal representation, which is updated with unsupervised recent experience in sometimes locally “irrational” manner. In the following, we further explore the novel insights gained by of our results.

The lagged logistic modeling approach allowed us to quantify short-term influences. Results confirmed previous findings assigning short-term positive serial influence to past decisions and a negative, perceptual adaptation like effect to previous stimuli (Fritsche et al., 2017). However, these complex short-term serial effects were unrelated to our main manipulation of long-term probabilities. Hence, our findings complement previous proposals focusing on short-term effects only, which suggested that short-term positive serial effects could underlie the stable perception of the environment (Fischer & Whitney, 2014). We found that long-term effects, even as simple as appearance rates, can influence perceptual decisions as strongly as short-term effects. Importantly, this effect was completely involuntary, since the observers were not asked to consider anything but the present trial and they were not aware of any of the appearance probability shifts occurring during the experiments. Thus, our overarching conclusion is that short- and long-term statistics automatically and continuously influence every momentary human decision, and the final effect will be inevitably shaped by the various contextual aspects of the task.

Perhaps the most surprising finding of the current results is that long-term influences can be so strong that within a particular context, they can bias people against choosing a locally more frequent stimulus. This causes us to examine previous perceptual decisions making studies under a new light as they have only shown a positive influence of the past, i.e. bias towards stimuli that that had been more frequent (Bohil & Wismer, 2014). Long-term influences were considered to be positive effects enhancing the preference toward the stimulus (Chopin & Mamassian, 2012) or pattern (Wilder et al., 2013) that had been more frequent in the past. Rather than being a simple positive baseline, our results suggest that at any time, there could be more than one such effect, they could emerge at

multiple time scales, and they could interact in complex manner producing opposing types of aggregate biases (positive or negative) depending on the relative strengths of the effects.

We found a dissociation between the influence of long-term and immediate past influences on perceptual decisions. Such a dissociation is corroborated by another study showing that explicitly cued probabilities affect the threshold, but recent trials the sensitivity of perceptual decisions (Kaneko & Sakai, 2015). One of the main novel aspect of our paradigm is that long-term influences are arising from a long implicit period of experiencing changing stimulus appearance probabilities and not from an explicit cue which signals which is the most likely stimulus as in most previous studies (Kaneko & Sakai, 2015; Mulder et al., 2012).

Another notable finding is that under particular circumstances response times and decisions are affected in opposing ways by stimulus statistics. Particularly, we have shown that long-term influences were specific to decisions, as opposed to reactions times, which are more sensitive to current probabilities. The differing effect of probabilities on reaction-times and decisions can be interpreted in the framework of the drift diffusion model (Ratcliff & Rouder, 1998). Specifically, it has been proposed that one can conceptualize long-term effects as a biased starting points in the diffusion process (Mulder et al., 2012), while short-term influences affect the process of evidence accumulation, manifesting in faster reaction times for recently frequent elements (Urai, Gee, & Donner, 2018). However, to obtain a similar dissociation to that in Exp 1 (decision bias opposing reaction time advantage), both a bias in the starting point and asymmetric evidence accumulation rates would be required (Appendix B.2 Fig & Text). Moreover, the different noise levels would also require across-trial variability in the drift-rate. Allowing all these parameters to vary would make drift-diffusion an overly flexible model, enabling it to fit almost any pattern of choice and reactions times (M. Jones & Dzhafarov, 2014). Furthermore, even though the model in which recent trials can influence the starting point of the drift rate/prior probabilities (i.e. bias) might look convincing, some authors suggested that prior probabilities might affect the drift rate itself (Hanks, Mazurek, Kiani,

Hopp, & Shadlen, 2011), suggesting that the effect of expectations on uncertain sensory evidence accumulation might be more integrated.

Our results from Experiment 1 are somewhat puzzling: why do participants prefer the rare stimulus? Experiments 2 and 3 show that the difference in long-term probabilities is a key factor. However, these experiments cannot explain why such irrational behavior arises in Exp 1 in the first place. In theory, participants could adapt to the new changed stimulus statistics resulting in a higher proportion of correct responses. Instead, what seems to happen is that people readjust their internal representation of the stimulus sequence in a highly irrational way. The cause of this readjustment is presumably a discrepancy between their expectations (balanced stimulus probabilities) and the observed unbalanced stimulus sequence at the beginning of the test phase. From the perspective of the observer, this discrepancy could arise for two different reasons: I. One of the stimuli is more frequent (change in prior). II. They are better able to detect one of the stimuli (change in likelihood), probably because it is less noisy¹⁸. If the participant followed explanation I., the internal representation should be adjusted in the opposite direction, and we would not have found a preference for the rare stimulus. This suggests that participants in fact followed explanation II by wrongly assuming a weaker ability to detect noisy stimuli from the rare category, resulting in a bias toward choosing it more often. This suggests that the withdrawal of feedback, combined with a shift in probabilities, could be a main underlying factor in the results of Exp. 1. We will test this hypothesis in Chapter 5.

To sum up, perceptual decisions are influenced by the difference between recent and long-term probabilities, independently from short-term serial influences. This pattern of results is specific to decisions, with reaction times more sensitive to current probabilities. People adjust their internal

¹⁸ A model-based approach to this explanation is detailed in the General Discussion Chapter.

models of stimulus distributions after changes in environmental probabilities, sometimes resulting in locally irrational biases.

Chapter 5

Change-dependent weighting of past probabilities: Exps. 5-7

Summary

In a set of three visual discrimination experiments we examined how changes in stimulus probabilities influence the way humans' update their reliance on past probabilistic information. We demonstrate that reliance on past statistics is highly dependent on the dynamics of the change in the latent parameter of appearance probability, not only and not even mostly on its slow cumulative statistics across individual trials. In particular, we show in Exp 5 that gradual changes in the AP at the transition from the training to test session do not elicit the striking integration of information about past probabilities found in Chapter 4. In contrast, a sudden transient apparent shift in stimulus statistics without any true long-term changes can elicit an equally strong and lasting bias on perceptual decision making to those evoked by true changes in Chapter 4 (Exps. 6-7). We also show that these effects are specific to decision, as response times remain more sensitive to momentary stimulus probabilities as in the previous experiments.

Introduction

The canonical approach to investigating human perceptual decision making is based on simple stimuli (random dot motions coherence, oriented Gabors, textures), steady state conditions (collecting responses under identical circumstances with repetitive trials, in which only one parameter is changing), the assumption of i.i.d. data (each trial can be investigated independently from any other trial, Green & Swets, 1966; Stanislaw & Todorov, 1999) and the idea that any adjustment in the process is implemented by gradual learning across many trials (Rescorla & Wagner, 1972). The underlying assumption of this approach is that perceptual decisions are essentially momentary-stimulus-driven as long as the task is simple and well-specified. The existence of adaptation effects (Thompson & Burr, 2009), the traditional priming effects (Treisman, 1992), and the more recently scrutinized serial decision making effects (Fischer & Whitney, 2014; Fritsche et al., 2017) has been acknowledged for a long time, but according to the common wisdom, these dynamic effects were of secondary importance when investigating the underlying process of truly perceptual decisions. While under the simplest conditions these assumptions might hold, it is becoming increasingly clear that even minute changes to the quality of the stimulus or the nature of the task can lead to unexpected violations, which in many cases calls into question the existence of purely stimulus driven perceptual decisions. Specifically, in Chapter 4 we showed that as long as the quality of the sensory information is degraded, long-term effects influence decisions well beyond what was expected from the results of the idealized setup. In the present chapter we will take a look at the non-dynamical assumption of the canonical approach and seek to answer whether the long-term summary statistics that can influence decisions so powerfully, indeed, emerge by gradual small-step integration based on past experiences.

To investigate this issue, one needs to manipulate the dynamics of changes in stimulus statistics and measure their effect on decision making. Previous results showed that people reacted to the

dynamics of changes in stimulus probabilities by adjusting their learning rate according to the perceived volatility of the environment (Behrens et al., 2007; Nassar et al., 2010). While these results speak to the dynamics of learning, they approach the issue from a particular direction, the general volatility of the stimulus generating process. Instead, we were interested in exploring the effect of changes in a more event-based manner by assessing the effects of individual changes spotted by the observer implicitly.

In the experiments of Chapter 4, there was one notable change point at the transition from the practice to the test session, when the appearance probabilities of the shapes changed, the noise applied to individual trials switched from staircase to random, and the observers stopped receiving continuous feedback. We found a strong effect of this one change point on the observers' decision-making behavior, suggesting that the change-point might be the crucial underlying factor of obtained results.

Therefore, in the current chapter we conducted three experiments, manipulating both stimulus probabilities and change dynamics in different manner and assessing the effect of these manipulations on the observers' behavior. Exp. 5 used the same probabilities as Exp. 1, but unlike in Exp. 1, where the change was abrupt, the change from 50 to 65% in Exp 5 was introduced gradually. Exp. 6 used balanced 50-50% probabilities both during training and test but had a sudden shift in stimulus probabilities at the beginning of test (similarly to Exp. 1), from which the AP gradually returned to baseline (unlike Exp. 1). Exp. 7 replicated Exp 6 using slightly different parameter settings in order to ensure the generalizability of the effect.

Eliminating long-term effects by gradual changes: Experiment 5

The rationale of Experiment 5 was to test whether changing the APs between the training and test conditions abruptly vs. gradually made any difference in the observers' decision making. Being able to influence the observer's behavior by simply changing the dynamics by which the steady state of

the test condition was achieved would cast serious doubt on whether the simplified canonical approach to decision making could help identify the nature of the underlying process. We expected to find that people adjust their internal representation of the task/stimulus structure if they experience a gradual change to a smaller extent relative what we found with a sudden change.

Methods

Experiment 5 followed closely the procedure of Exps. 1-4 (**Fig 4.1**). Specifically, the first 200 trials (training) were presented with feedback, and an adaptive staircase procedure was used. Next, the test session followed with 400 trials instead of 300 used in Exps. 1-4. The additional trials were introduced because the change in stimulus appearance probabilities occurred throughout the first 80 trials (**Fig 5.1A**), therefore, we needed a longer test session to have the same number of trials after the change.

Gradual stimulus appearance probability changes were pseudo-randomly controlled in four 20-trial-long periods: out of the first 20 trials, 10-10 showed the “frequent” vs. “infrequent” shape in a randomized order, and the ratio changed gradually to 11-9, 12-8, and finally to 13-7 in the next of 60 trials. The 13-7 ratio corresponds to 65% appearance probability for the frequent element. After the initial 80 trials (during which the change had occurred), the appearance probability remained at 65% without a pseudo-random control of the sequence. In all other respects, Experiment 5 was identical to the first four experiments.

Participants

20 students (18-30 years old) completed the experiment after giving informed consent and received monetary compensation.

Analysis

To test our main question of interest, we used 300 test trials for data analysis (as in Exps. 1-4) after discarding the first 80 trials of the test sequence, during which the gradual shift occurred.

Psychometric curves and the logistic regression weights including recent past influences were fitted individually the same way as for Exps. 1-4. Since model selection was performed on the combined dataset from Exps. 1-7, we used the same logistic model including an intercept, the current stimulus, and the previous two decisions.

Results

Descriptive statistics of Behavior

The overall performance in Exp 5 was very similar to those in the four base experiments: participants responded correctly on 71.36 \pm 6.3 % of trials (**Fig 5.1B**). As in the earlier experiments, the performance was highly noise-dependent with performance dropping from over 90% at the easiest 1/3rd of trials (83.05 \pm 6.75 %) to about 60% at highest noise level (62.74 \pm 6.3 %). The median response times were on average 753 \pm 76 ms. Sensitivity, as measured by the slope of the psychometric curve was not different from Exp 1 ($t_{38}=0.4636$, $p=.6456$, **5.1C**).

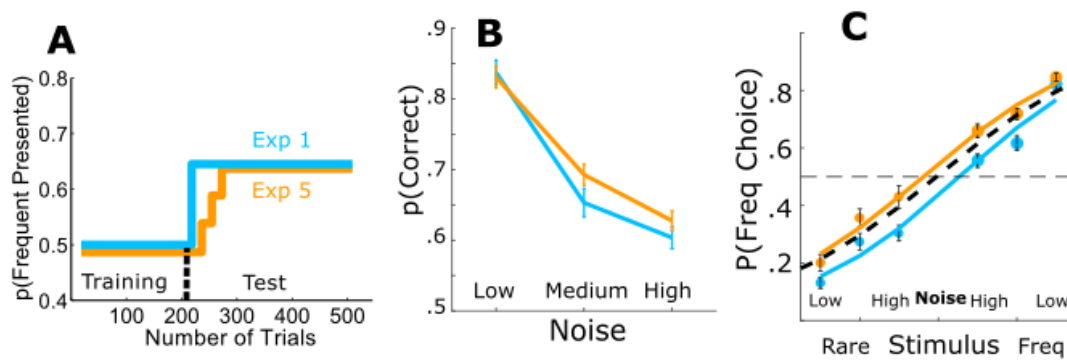


Figure 5.1. Experiment 5. **A)** Probability structure of Exp 5 (yellow) compared to Exp 1 (blue). Both experiments had a balanced training followed by a 65% test. However, unlike in Exp1, in Exp 5 the change took place gradually over the course of 80 trials. We used 300 trials after the change had happened in Exp 5. and compared it to the entire 300 trial test session of Exp 1. **B)** Performance in Exp 1 & 5. **C)** Response proportions as a function of stimulus and noise (dots) and the fitted psychometric curves in (lines) in Exp 1 & 5. Errorbars: SEM.

Long-term influences on decisions

In Exp 5, after the gradual change, there was no overall preference for either stimulus (**Figs 5.1C, 5.2A**) (Model Bias, $t_{19}=0.0073$, $p=.9943$, $BF=0.23$). This suggests that since there was no noticeable change in stimulus statistics, participants did not adjust decision criteria and just responded to the presented stimuli in an unbiased manner. This result is in stark contrast with the outcome of Exp 1, where the same training and steady probabilities during test elicited a strong compensation effect, leading to a significantly different bias from that in Exp 5 ($t_{38}=2.913$, $p=.006$, $BF=20.04$). This effect was stable during the test session, with no significant change between the first and second half of the test session ($t_{19}=1.122$, $p=.2758$, $BF=0.40$).

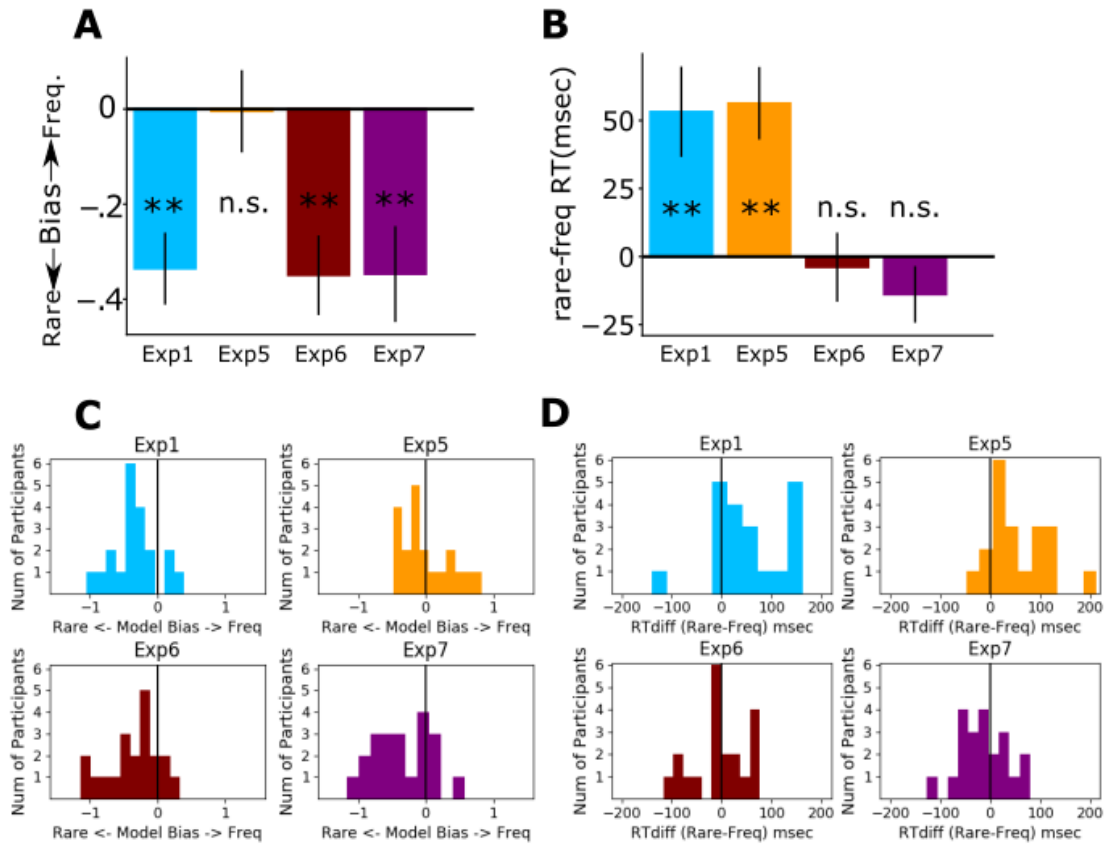


Figure 5.2 Reaction times and decisions are differently affected by change dynamics. **A)** Decision bias for Experiments 5,6, and 7 with Exp1 as comparison. Large majority of participants were strongly biased against choosing the locally more frequent shape during the stable period of Exp 6-7, while there was no overall preference in Exp 5. **B)** Reaction Time Differences. Rare-Frequent. When the momentary probabilities were unbalanced, participants were faster to respond with the frequent option (Exp 1, Exp 5), with no difference with balanced momentary probabilities (Exp 6, Exp 7). **C)** Distribution of biases in Exps 1,5,6, and 7. **D)** Distribution of RT differences in Exps 1,5,6, and 7. (Error bars: SEM)

Long Term influences on Response Times

In contrast to decision patterns, we found that after the gradual change, reaction times were faster for the frequent element in Exp 5 by 56 ms on average (**Fig 5.2B** RT Difference Frequent-Rare: $t_{19}=4.22$, $p=.0005$, $BF=71.6$), similarly to Experiments 1-3, where the mean speed-up values were 49 ms, 36 ms and 78 ms, respectively. Further, there was no statistical difference between Exp 1 and Exp 5 in terms of this reaction time speed-up for frequent responses ($t_{38}=0.1474$, $p=.883$, $BF=.1202$). This shows that the same momentary probabilities elicited a very similar reaction time pattern, regardless of the past change dynamics.

Inducing a long-term effect by a sudden change without shifting appearance probabilities

Experiment 6

In Exp 5, after eliminating the notable shift in the appearance probabilities of shapes at the boundary of the training and test sessions, we found that observers exhibited no evidence of adjusting their behavior despite the strong difference between the true appearance probabilities in the training and test periods. Although this indicates a paramount role for a clearly detectable change point in the observed sequence in how observers adapt to the statistics of their environment, the scope of this effect is not clear. Specifically, Exp. 5 showed that gradual build-up can *eliminate* the shift in the observers' bias. However, is it possible for a well-positioned sudden change to *force* the observer to adjust his/her bias despite having no true change in the appearance probabilities at all? How strong can this effect be? How long would it last? In Experiment 6, we provide an answer to these questions.

Methods

Exp 6 was almost identical to Exp 5 with one crucial difference: after the 200-trial training session with 50-50% appearance probabilities (APs), at the beginning of the test session, the APs jumped to 65-35%, and then gradually returned to the 50-50% level (**Fig 5.3A**). Specifically, in the first 20 trials of the test session, 13 vs. 7 “frequent” and “infrequent” shapes were shown, respectively, in a randomized order. As mentioned above, this ratio corresponds to 65% appearance probability for the frequent element, imitating a strong shift in the AP in favor of one of the shapes. However, in the following three sets of 20 trials, this ratio was gradually changed back to 12-8, 11-9, and finally 10-10

ratios, arriving to the same chance performance experienced during the training period. As before, the order of appearance within the 20 trials was pseudo-randomized. In other words, except for 60 trials at the beginning of the test session, the entire training and the entire test session provided strong cumulative evidence that the two shapes appeared with equal chance. Even within the 60 manipulated trials there was only a brief 20-trial long period when the two shapes complied with the 65-35%, after which there was a gradual decay to 50-50%. In all other aspects, Experiment 6 was identical to Experiments 1-5 (**Fig 4.1**).

Participants

22 students (18-30 year old) completed the experiment after giving informed consent, with 2 exclusions due to chance performance (below 53% correct overall) leaving a final sample of 20.

Analysis

Once again, we used 300 test trials for data analysis (as in Exp 5) after discarding the first 60 trials of the test sequence, where the gradual shift occurred. For those 300 trials, we followed the same procedure in fitting the psychometric curves and calculating the logistic regressions weights as in Exps 1-5.




Figure 5.3 Experiment 6. **A)** Structure of Exp 6 (brown) compared to Exp 1 (blue). In Exp 6, there was a sudden change to $AP=.65$ frequent at the beginning of the test block, which gradually disappeared over the course of 60 trials. **B)** Performance in Exp 6 was highly noise dependent, similarly to Exp 1 and the other previous experiments. **C)** Response Probabilities as a function of stimulus and noise, and psychometric curves in Exp 1 & 6. The fitted functions were almost completely of the top of each other. After the change happened, decision in Exp 6 showed a strong preference for the stimulus that had been rare during the training just as in Exp. 1. (Black dashed line indicates the no-bias condition. Error bars: SEM)

Results

Descriptive statistics of Behavior

Overall accuracy was 72.15 +/- 8.53%. Performance was highly noise dependent with performance dropping from about 85% at the easiest 1/3th of trials (84.15 +/- 10.37%) to about 60% at high noise (62.44 +/- 7.46 %) (**5.3B**). Median response time was 739 +/- 121ms on average. Sensitivity (as measured by the slope of psychometric function) was not significantly different either from Experiment 5 ($t_{38}=0.4578$, $p=.6497$), or from Exp 1 ($t_{38}=0.7257$, $p=.4725$).

Long-term influences on decisions

In Experiment 6, we found a significant preference *away* from stimulus that was frequent during the transient probability shift, after the change happened (**Figs 5.2B & 5.3C**, Model Bias: $t_{19}=4.1714$, $p=.0005$, $BF=64.96$). This effect emerged over time as observers still demonstrated a tendency of bias to the direction of the locally frequent shape during the initial exposure to the unbalanced sequence at the beginning of the test sequence, in the first 30 test trials ($t_{19}=1.9862$, $p=.0616$, $BF=$




Figure 5.4: Emergence and Temporal Stability of Long-term effects. In Exps 1,6,7(A, B, C) during first experiencing an unbalanced stimulus sequence (initial 30 trials at the beginning of test), observers were unbiased (Exp1-**A**) or demonstrated a tendency to prefer the frequent stimulus (Exp 6,7, **B,C**). Upon some experience with unbalanced probabilities, they started to prefer the rare stimulus, regardless if probabilities stayed unbalanced (Exp 1) or returned to 50-50 (Exp 6,7). This compensation like effect was stable, with a significant preference for rare stimuli after several minutes of balanced experience by the second half of test session (Exp 6,7). Init.30.: Initial 30 trials of test session. 1stH: 1st Half of test session, 2ndH: 2nd half of test session. (for Exp 1 this constitutes the halves of the whole test session, for Exp 6&7 it is the initial 30 trials from the beginning of test, and the two halves of the 300 trials after the change took place). (For descriptives on this data for Exps 1-7 see Appendix Table B.4.)

1.181, **Fig 5.4B**). After the shift to prefer the locally rare shape took place, the effect was stable during the entire test session, with no significant change between the first and second half of trials ($t_{19}=0.9304$, $p=.3638$, $BF=.3409$, **Fig 5.4B**). Overall, the bias of the observers was significantly different from that in Experiment 5 ($t_{38}=2.8965$, $p=.0062$; $BF=20.0075$) and virtually identical to that obtained by the original setup in Experiment 1 ($t_{38}=.1229$, $p=.9028$, $BF=.1564$).

Long Term influences on Response Times

In contrast to Experiment 1 and 5 we have found no overall difference in RT of the stable period between the previously frequent and rare responses (**Fig 5.2B** $t_{19}= 0.3058$, $p=.7631$, $BF=0.24^{19}$) despite the clear decision bias. This is also in line with our findings from Exp 4, confirming that a bias resulting from past probabilities is not sufficient to elicit a strong reaction time-difference, if the current probabilities are balanced.

Scaling of change-point-related effects: Experiment 7

In Exps 5 and 6, we found a very strong effect of information provided immediately after the point of a sudden change in conditions. Since such a dynamic effect influencing perceptual decision making has not been reported before, the nature of such an effect is unknown at present. We ran Experiment 7 with two purposes in mind. First, we wanted to replicate the results of Experiment 6 and thus to confirm the pivotal role of detected or assumed changes on perceptual decisions. Second, we wanted to assess how precise this process was: Would the immediately detected new local information be proportionally incorporated in the observer's decision behavior?

¹⁹ Appendix Table B.5. for descriptive stats

Methods

Experiment 7 was a slightly modified version of Experiment 6, in which the sudden change in appearance probabilities at the beginning of the test session was larger, jumping to 75-25% instead of 65-35%. This means that the subsequent return to the 50-50% baseline took 100 rather than 60 trials as in Experiment 6, starting with 15-5 ratio in the first set of 20 trials and eliminating the difference by steps of 1 in following sets of 20 trials (**Fig 5.5A**). In all other aspects, Experiment 7 was identical to Experiments 1-6 (**Fig 4.1**).

Participants

25 Participants (18-30 year old) completed the experiment after giving consent, out of which 4 were excluded leaving a final sample of 21. Three participants were excluded due to chance performance (below 60% at low noise), one because of a surprisingly high performance, over 90% correct overall and 100% correct at high noise, leaving very few trials to calculate our measures of interest, since bias measures are impossible to calculate if the answer is always veridical.²⁰

Analysis

Once again, we used 300 test trials for data analysis (as in Exps 5-6) after discounting the first 100 trials of the test sequence, where the gradual shift occurred. For those 300 trials, we followed the same procedure in fitting the psychometric curves and calculating the logistic regressions weights as in Experiments 1-6.




Figure 5.5. Experiment 7: **A**) Probability structure of Exp. 7, vs. Exps. 1 & 6. A strong initial change (to AP=.75) returned to a balanced probability over the course of 100 trials in Exp. 7. We compared the periods after the change had happened. **B**) Performance in Exp. 7 was highly noise dependent, with Exp. 6 and Exp. 1 as comparison (thin lines). **C**) The effect of this stronger change on response bias was very similar in Exp 7 to those in Exps. 6 & 1 as shown by the response proportions (dots) and the fitted psychometric curves (thin lines). Black dashed line indicates the no-bias. Error bars: SEM

Results

Descriptive statistics of Behavior

Performance was 71.18+/- 6.41% on average, and highly noise dependent with performance dropping from over 85% at the easiest 1/3rd of trials (87.02+/- 6.92%) to about 60% at highest noise level (61.87 +/-7.5 %) (**Fig 5.5B**). Median response time was on average 761 +/- 132 ms. Sensitivity was not significantly different from Exp. 6 (**Fig 5.5C** Slope of Exp 6 vs. 7: $t_{39}=1.3494$, $p= .185$).

Long-term influences on decisions

Replicating the main finding from Exp6, in Exp 7, we found a significant bias *away* from stimulus that was frequent during the transient probability shift. (**Fig 5.2B & 5.5C**, Model Bias: $t_{20}= 3.4377$, $p=.0026$, $BF=15.22$). The effect emerged over time since it was absent initially at the beginning of the test ($t_{20}=1.326$, $p=.1998$, $BF=0.49$, **Fig 5.4C**) but was stable throughout the balanced period of the test session after the change happened, with no systematic change between the first and second half of the test session ($t_{20}=1.1126$, $p=.279$, $BF=.393$). This long-term effect was very similar to that obtained in Exps. 6 and 1 (Model Bias Difference Exp7 vs Exp6: $t_{39}= 0.0186$, $p= .9852$ $BF= .168$; Exp7 vs Exp1 $t= .0899$, $p=.9288$ $BF= .171$), and significantly different from Exp 5 ($t_{39}= 2.5939$, $p=.0133$, $BF= 9.64$). This finding confirms that a transient shift in stimulus probabilities under high uncertainty is sufficient to bias subsequent perceptual decisions. The very similar overall magnitude of the effect found in Exps. 6 and 7 attests that this effect does not simply scale proportionally with the magnitude of the sudden perceived shift.

Long Term influences on Response Times

The RT results of Exp. 7 replicated the corresponding findings in Exp. 6 as well. We found no overall difference in RT of the stable period between the previously frequent and rare objects (**Fig 5.2B**, $t_{20}= 1.3357$, $p=.1966$, $BF=0.49$) despite the clear decision bias.

	Predictors' Weights			
↓Dependent Var	Training AP	Test AP	Overall AP Change	Sudden AP change
Decision Bias alpha= .09816	$\beta = 0.1076$	$\beta = 0.$	$\beta = -0.$	$\beta = -0.3331$
RT (rare-freq) alpha= .0169	$\beta = 0.$	$\beta = 0.2783$	$\beta = 0.$	$\beta = -0.1741$

Table 5.1: Predicting Decisions Bias and RT in Exps. 1-7. To select from correlated experimental parameters as predictors, we used Lasso regression separately to fit our main measures of interest. Training AP only had three possible values (.5, .65,.75) Test AP only had two values (.5,.65). Overall AP change was simply the difference Test AP-Train AP. Sudden AP change was encoded based on the sudden difference between training and test (For Exp 5,6,7: 0,.15,.25). For the Experiments 1-4, "Overall AP change" and "Sudden AP change" had equal values (since the only available value was "sudden"). We fitted the model using the sk-learn library. The value of the cross-validated regularization parameter alpha is shown in table for each dependent measure. The regularized model could explain almost 30% of variance in decision bias ($r^2=0.27$). The response time analysis could explain about 13.5% of variance. The proportion of variance explained was highly constrained (for both RT and Bias) by the ability of a categorical variable (experimental manipulations) to predict a measure on a continuous scale. (Gray background: best predictors)

Combined Prediction of RT and Choice in Exps. 1-7

We wanted to test the extent to which our different experimental manipulations predicted our measures in a combined model on the entire data set. To achieve this, Cross-Validated Lasso regression (Tibshirani, 1996) was used with the main experimental parameters as predictors (Training AP, Test AP, Overall Probability Change, Sudden Probability Change **Table 5.1**) to fit the individual decision bias and reaction time differences. The sudden probability change was by far the most important determinant of choice bias (regularized weight $\beta = -0.3331$) while the training probability had a small additional predictive power (**Table 5.1**). This shows that a sudden increase in probability between training and test predicts a stronger preference for the rare element. In contrast, using the same regressors to predict reaction times, we found that the test probability was the best predictor ($\beta = .2783$), followed by the sudden probability change ($\beta = -.1741$). This shows that the higher the test probability, the faster participants were in responding to frequent elements. This analysis confirms the dissociation between momentary probabilities and sudden probability changes in their respective

influence on response times and decisions, with reactions time being more sensitive to momentary probabilities, and decision bias to sudden probability changes.

Discussion

We have found that a sudden shift in stimulus probabilities is necessary (Exp. 1 vs. Exp. 5) and sufficient (Exps 6,7) to elicit a lasting preference for the locally rare element. The results of Exps. 6-7 are especially striking as stimulus probabilities were balanced for most of the experiments, still a transient shift in stimulus probabilities could elicit a lasting irrational bias toward the previously rare stimulus. Reaction times showed a very different pattern, with a strong influence of the current probabilities in Exp 5, and no clear patterns during the balanced period of Exps. 6 & 7.

It is interesting to contrast the response time and decision bias patterns across the whole set of experiments in Chapters 4 and 5. The difference is best demonstrated by Experiments 1 and 5, as both experiments had the same 65% test session with the change dynamics being the only difference, and yet in Exp 1 there was a strong decision bias towards the rare elements, with a clear response time advantage for the frequent element. In contrast, the decision bias was largely absent in Exp. 4, despite very similar reaction time patterns. While these experiments show that response times and decisions are sensitive to different manipulations, they are obviously not independent. Looking at response time differences and decision bias across the full dataset (Exps. 1-7) shows that there is a highly significant correlation between the two measures ($r_{139}=.401$, $p<.0001$). Therefore, our results show that despite the fact that these measures are correlated on the individual level, they are sensitive to stimulus statistics on different time-scales.

We have suggested that the explanation for the results of Exp. 1 could be a wrongly adjusted internal model of the statistical properties of the stimulus sequence after a change in stimulus probabilities. Experiments 5-7 confirm this assumption and provide a compelling explanation, namely that internal representations are readjusted as a consequence of the sudden discrepancy between long-term and current probabilities. Exp. 5 shows that a gradual shift is insufficient to elicit such an effect, since if

there is no sudden discrepancy between long-term expectations and current statistics, there is no need to readjust the internal model. Experiments 6 and 7 confirm the crucial importance of sudden stimulus probability shifts in readjusting the internal model: a sudden discrepancy in stimulus probabilities elicits a readjustment in how people interpret these noisy stimuli: instead of noticing the change in AP, participants change their model of their ability to detect these noisy stimuli²¹. This readjusted model is used afterwards, when unbeknownst to the participant, the experiment has returned to a balanced stimulus structure.

This explanation can be linked to findings on volatility influences in decision making (Behrens et al., 2007; Glaze et al., 2015). Those studies have shown that a discrepancy between past and current stimulus statistics trigger an increase in learning rate that facilitates the learning of the changed statistics. Our experiments also show that a shift in statistics promotes learning about the new conditions, however, the unsupervised nature of the task elicits a change in the internal model that is irrational, resulting in suboptimal performance. Therefore, the absence of feedback during the test is also likely to contribute to the novel pattern of results that is completely unlike the previous results of the reinforcement learning literature (Nassar et al., 2010).

Taken together, these findings reveal that change dynamics are a key factor in determining how people use their internal representation of stimulus statistics to bias perceptual decisions. Sudden changes elicit an adjustment of the internal model, while unnoticeable gradual changes do not elicit an update of the internal representation for a long time even when it would be warranted by long-term accumulation of local evidence.

²¹ See more on this explanation in the General Discussion Chapter

Chapter 6

General Discussion

Active learning as sequential decision making

In this thesis, we presented two main lines of research, one on eye-movement-based active statistical learning (Chapters 2-3), and the other one on sequential perceptual decision making (Chapters 4-5), in a largely separate manner. Indeed, the paradigms and measures applied in these two research projects were very distinctive, thus warranting their separate treatment to a certain extent. However, scrutinizing the underlying processes revealed that, in fact, they represent two intimately entangled aspects of an overarching phenomenon. Active exploration and sequential decision making are not only related, but active exploration can be interpreted as a special version of sequential decision making. In the following section I will elaborate the intuition behind relating these two research areas.

In an active experimental set-up, just as during natural visual behavior, each saccade is a decision in itself, which is influenced by both the current sensory input, but also by representations of past stimuli and past eye-movements (Hayhoe & Ballard, 2005; Posner, Rafal, Choate, & Vaughan, 1985). This is reminiscent of perceptual decision making, where the sequence of past decisions and stimuli influence how the momentary stimulus is interpreted (Fritsche et al., 2017; Maloney et al., 2005). However, in the active set-up, multiple layers of complexity are added to the process due to a recurrent processing loop, since each explorative decision on the next fixation influences the sensory input that will arrive in the next moment, as well as the future state of the decision-maker (Yang, Wolpert, & Lengyel, 2016). Importantly, this connection between past and future states in the active set-up is closer to the complexity of real life than that of a simple sequential perceptual decision-making process is. Outside of the lab, every deliberate or implicit decision about which part of a painting to look at, which road to take at the intersection, or which lunch menu to buy will affect both the available sensory input and future possible actions. In comparison, most laboratory sequential decision-making tasks are hugely simplified versions of the natural active set-up: decisions do not affect future stimuli, only the interpretation of those stimuli via internal states, and in most

cognitive psychology experiments, even the effect of internal states is ignored by assuming independent trials.

We are not the first to propose a link between these areas. For example, Ahmad and colleagues suggested that active sensing could be treated as a Bayesian sequential decision making problem (Ahmad, Jolla, & Yu, 2013). In this theoretical approach, the authors successfully formalized visual search as sequential decision-making using a Markov Decision Process model (Ahmad et al., 2013). Linking a Markov Decision Process to actual experimental data had also been attempted recently (Hoppe & Rothkopf, 2017). Hoppe & Rothkopf (2017) showed that people can plan multiple saccades ahead and select their fixation targets differently depending on the available time, as if they optimized the trade-off between immediate and later information.

Active learning as ecological decision making

While the above studies are remarkable steps in modeling the sequential nature of visual exploration, they are still confined to the limited scenario of searching for a single noisy target. This setup is simpler by orders of magnitude than our active statistical learning task (Chapter 2,3), since our task required integrating information both during and across individual trials. This difference can be conceptualized as the distinction between active sensing and active learning. Both the theoretical (Ahmad et al., 2013) and the experimental (Hoppe & Rothkopf, 2017; Najemnik & Geisler, 2005) approaches investigating visual search in the past used the framework of active sensing: information had to be integrated and exploited only within a single trial. In contrast, in our task, stimuli on any single trial cannot reveal anything about the underlying structure, thus, successful active exploration can only emerge in an active learning framework, i.e. if the observer can benefit from information integrated across multiple trials. In active sensing, the prior probabilities of the environment are assumed to be known (Yang, Wolpert, et al., 2016), which might be the case for single target visual search (Hoppe & Rothkopf, 2017; Najemnik & Geisler, 2005) or binary discrimination (Yang, Lengyel,

et al., 2016). In a more complex scenario, good decisions cannot be based on solely the momentary stimulus but must be learned via interactions with the environment. This suggests an interpretation whereby visual decisions have to rely on an internal model, which must be continuously learned and updated based on the sensory input. Thus, efficient decisions are only possible if expectations are well-tuned to the statistical properties of the environment (Fiser, Berkes, Orbán, & Lengyel, 2010). In an active set-up, this implies a loop-like interaction between exploration and learning, where knowledge acquired will influence where the observer looks next, but also the position where the observer looks will influence what can s/he will learn. This framework points to an integrated view in which learning and attention are in a continuous interaction (Chun & Turk-Browne, 2007).

In order to fully understand active learning, one must first address the preliminary question of how people can learn the relevant statistics of their surroundings. The simplest answer is that in most scenarios people receive feedback on whether their actions were successful or not. Accordingly, in many experimental designs, feedback is used to learn about the structure of the task (eg: Behrens et al., 2007; Glaze et al., 2015). In real life, however, many of our actions do not have an outcome that is immediately available and can be easily distinguished as rewarding or punishing. Mimicking such a natural scenario, after most decisions in the experiments of the current thesis, there is no immediate nor even delayed feedback. What can drive learning in unsupervised scenarios of this sort? From the psychological perspective the answer is curiosity (Kang et al., 2009; Kidd & Hayden, 2015), which might be implemented in the brain via the intrinsic reward for information that reduces uncertainty (Foley, Kelly, Mhatre, Lopes, & Gottlieb, 2017; Jepma et al., 2012). We expect that this rewarding influence of information could be important in motivating our participants to discover and learn the regularities of the stimuli even if they cannot achieve any immediate reward by doing so.

Active statistical learning in the light of active learning theory

The theory outlined above appears almost self-evident, people are active learners (Gottlieb, 2012), and visual attention and learning continuously influence each other (Chun & Turk-Browne, 2007). On a closer look, however, there is only limited evidence that human visual behavior is well described by the predictions of these theories. Very little is known about how environmental statistics influence visual attention, how such effects are related to learning those regularities, and how these interactions between learning and visual attention depend on task, supervision, and the types of regularities. This is why we conducted Study 1 using a basic spatial statistical learning paradigm (Fiser & Aslin, 2001) and embedded in a novel gaze-contingent active exploration set-up (Chapter 2) to make sure that eye-movements are tightly linked to sampled visual information.

Our main finding was that implicitly and explicitly learned spatial regularities guide visual exploration, showing that statistical representation effects on eye-movements emerge via learning even without explicit guidance. To our knowledge this is the first study to show such complex statistical learning effects on explorative visual decision. What can these findings tell about our theoretical understanding of humans as active learners? The fact that implicitly acquired spatial stimulus regularities influence eye-movement patterns shows that explorative visual decisions do, in fact, rely on statistical representations available only by learning over a longer period. Furthermore, we found that a learned representation of the environment could not only guide visual exploration, but also predicted learning on the subsequent test. This suggests that, once learned, stimulus statistics influence visual exploration, but in turn visual exploration enhances the learning those regularities. The complexity of this effect is far beyond previous results on active sensing (Hoppe & Rothkopf, 2017; Najemnik & Geisler, 2005; Yang, Lengyel, et al., 2016), since in our task these regularities could only be acquired across many trials, and still, stimuli appearing at different locations could bias the direction of visual exploration.

Our results also stress the relevance of the task on the manifestation of these effects. We showed that active exploration uses the hidden underlying statistical structure of the environment only when this structure is relevant for the task (Chapter 2: Exp 1- Explicit), or when learning is guided by an implicit, possibly curiosity driven, process (Chapter 2, Exps. 2-3, Implicit). However, when the task is unrelated to the regularities, learning effects prompted by the statistical structure are suppressed and do not dominate active exploration (Chapter 3, Exp 4-Working Memory), despite the fact that the regularities are learned nevertheless.

Further investigation of the link between learning and active exploration suggests that while initially learning influences visual exploration, it is only later that the two processes become intimately tied together. Our design using explicit instructions (Chapter2/Exp1) could start with a hypothesis testing process (Trueswell et al., 2013) and the experimentally observed direct gaze returns in this setup could signal testing of potential regularities consistent with the hypothesis. However, in any case, it is difficult to investigate the causality of these processes as learning can exert its influence on gaze directions only via looking behavior at the pair level, which makes it a “chicken or the egg” problem to some extent.

To investigate the causal link between active exploration and learning, previous studies played back the exploration data of previous participants to passive observers and analyzed whether seeing the same stimuli can result in similar performance as actively exploring them (Markant & Gureckis, 2014). A similar approach could be applied to test whether passively observing the same scenes vs. active exploration leads to a similar learning performance. Same performance by observational learning and active exploration would suggest that eye-movements are only a consequence- and not a prerequisite of learning. In contrast, the theory of active learning predicts that observing somebody else’s exploration data would hinder performance (especially in the explicit scenario). Intuitively, the exploration patterns used by a previous observer at any given trial should not be helpful for the learner since the previous observer was most probably at a different stage of acquiring environmental regularities. In addition, since passively observing the scenes is a very distinct kind of

behavior than active exploration (for example, it is likely less engaging), any resulting difference in performance may reflect motivational or attentional factors. This could make direct comparisons between learning via active exploration vs. “learning via observing another’s exploration” problematic.

Across all three experiments of Chapter 2, influences of environmental statistics on eye-movements became stronger over time, and such a tighter link between statistics and eye-movements predicted better learning performance on the subsequent familiarity test. Thus, our results confirm the general predictions of the theoretical account of active learning. Meanwhile, there are more specific predictions not addressed directly through the current experiments, which can nevertheless be speculated about in the light of our results. For example, the active learning account postulates that attention should be directed to stimuli in order to maximally reduce uncertainty about the environment based on what is informative at the current state of learning (Settles, 2010; Yang, Wolpert, et al., 2016). This predicts that once knowledge about a statistical relationship (e.g.: a spatial pair) is fully acquired, further looking at a given pair does not provide any extra information for the learner about this specific statistical rule. Hence, following our paradigm, after every pair is learned, attention should not be directed in accordance with the statistical structure. Our design was not aimed at testing primarily this “inverse u-shaped” dynamic between attention and learning (Kang et al., 2009; Kidd, Piantadosi, et al., 2012), but we still observed evidence indicative of this pattern in some explicit learners (Exp 1 in Chapter 2). It is very likely that most participants were still at the upward part of the theoretical “inverse u-shaped” curve, and thus did not reach a level of confidence in the statistical structure that would start weakening these influences. It is also true that our stimuli in Chapter 2 were built exclusively of pairs, thus once the pair structure was fully acquired, there was nothing else to learn from the presented stimuli. It would be interesting to include not fully predictive pairs with higher-order statistical structure in a similar statistical learning paradigm to test whether once the simple pair knowledge is learned, eye-movements would begin to explore these more complex relationships. This could pave the ground to research exploring whether depending on

the state of the internal model of the environment, different kind of information would attract attention at different times in accordance with the current need/stage of the learning process.

It also remains open for future research if the approaches successful in modeling active sensing (Ahmad et al., 2013; Hoppe & Rothkopf, 2017; Yang, Lengyel, et al., 2016) could also be applied to our active exploration data. Such an inquiry could reveal the extent to which explorative eye-movements correspond to the predictions of an optimal active learning model, which is trying to reduce uncertainty about the environmental structure over a longer period of time (across many trials). Establishing such a link could show that explorative actions are aimed at acquiring relevant statistical representations of the environment, which in turn could boost successful subsequent learning and actions.

Modeling the learning of environmental regularities

In the second half of the thesis (Chapters 4-5) we focused on sequential perceptual decision-making. In this area, sizeable effort has been directed recently at short-term inter-trial influences on the order of few seconds (Akaishi et al., 2014; Cicchini et al., 2017; Fischer & Whitney, 2014; Fritsche et al., 2017). However, our main interest was in the largely ignored past influences emerging over a longer period (5-15 minutes). We have shown that past stimulus probabilities strongly influence perceptual decisions, and that the extent of this influence depends on the presence of sudden shifts in stimulus statistics, sometimes resulting in locally irrational decision biases.

Similarly to some of the previous reinforcement learning studies (Behrens et al., 2007; Nassar et al., 2010), we have also found that changes in stimulus distributions play a crucial role in the update of internal models. However, participants' behavior around the change point was very different in our experiments due to the unsupervised nature of the task: our findings were often in direct opposition to what would be expected based on gradual learning of single environmental parameter with feedback (Behrens et al., 2007; Rescorla & Wagner, 1972). Specifically, we have found that sudden

shifts in stimulus probabilities can elicit a lasting and seemingly counter-intuitive preference toward the locally rare element (Chapters 4-5).

Superficially, our results are somewhat similar to those of Chopin & Mamassian (2012) since both studies describe an influence of both long- and short-term probabilistic experience. However, they report that the negative influence (adaptation) of the immediate past is modulated by differences from long-term experience, as if the current perceptual decisions were balancing the long- and short-term past (Chopin & Mamassian, 2012). In contrast, we showed that changing probabilities over time set up a long-term bias, which acts independently of the attractive influence of decisions from the immediate past. Our framework suggests that the momentary bias does not balance the long-term and immediate past, rather it is just a consequence of these two largely independent influences.

Since the previous models mentioned above cannot explain our findings, we propose a more principled explanation based on a Bayesian Observer Model below. To introduce the model, we start with an example demonstrating the omnipresent uncertainty that is ubiquitous in visual decision making. It is well known that the same visual input can result in very different interpretations depending on context (Gregory, 1970). Imagine walking in a foreign forest on a foggy night, while trying not to bump into anything. Suddenly, you realize that you see fewer trees ahead of you. The reason for this could be that there are, in fact, fewer trees and you are approaching the end of the forest or, alternatively, the fog might have become heavier hiding some of the trees from view. In other words, attributing the noticed changes to different parts of the observer's model, specifically either to the prior term (frequency of trees) or the likelihood term (more noise on perceiving trees), can account equally well the novel visual experience. Depending on which of the two hidden causes the observed change is attributed to, the adequate action can be the opposite (i.e.: either walk faster to get out quickly or slower to compensate for the increased difficulty). This example is analogous to our task, when a sudden change in the observed frequency of one stimulus type can be attributed to a change in appearance probability or alternatively, to a weaker ability to perceive stimuli from the rare category, which leads to opposing decision biases for higher performance.

It is often difficult to separate whether people rely on an altered prior or likelihood (or a mixture of the two) since, the data might be explained equally well by both type of changes (Laquitaine & Gardner, 2018). However, in our case, had participants changed their prior on appearance probability, they would have changed it in the direction opposite from the actual stimulus probability change experienced locally (since they would prefer the rare stimulus), which is a rather unreasonable thing to do. This suggests that they might have changed their likelihood/noise model.

Why would participants modify the noise model instead of changing their prior? We propose that they choose changing the likelihood because they are highly uncertain in their ability to perceive these noisy stimuli, while they do not have any reason to expect that prior probabilities would change. Our follow-up preliminary findings show that this is in fact, a likely explanation. Training people with volatile stimulus probabilities in a slightly modified version of the original paradigm used in Chapter 4 made people more prone to change their priors on appearance frequencies (Koblinger, Arató, & Fiser, 2018). This result suggests a framework, in which people change parameters of their internal model of the task in an uncertainty-weighted manner similarly to the strategy implemented during optimal cue combination (Ernst & Banks, 2002). If the observer is highly uncertain about her/his ability to perceive the noisy stimuli but had stable balanced training, the likelihood (noise) part of the internal model is updated (as in Chapter 4-5). If she/he is more uncertain about the stimulus appearance probabilities, she/he is more willing to update the prior of the internal model, resulting in an adaptation to changed stimulus probabilities (Koblinger et al., 2018). Therefore, the uncertainty attached to different components of the internal representation of the environmental statistics can explain how these representations are updated under changing circumstances.

Role of uncertainty in past probability effects

Since the model described above suggests that uncertainty plays a crucial role in our findings, it would be very useful to look at measures of confidence/uncertainty and see if they are in fact sensitive to the changes in task conditions in the way forecasted above. An approach that had recent

success in linking subjective uncertainty to sequential perceptual decision-making biases used pupil size as a measure. Pupil size based uncertainty, which is related to a general arousal state of the brain, affects both short-term serial influences (Urai, Braun, & Donner, 2017), and adjustments to long-term stimulus statistics changes (Krishnamurthy, Nassar, Sarode, & Gold, 2017) during perceptual decision making. The main idea is that depending on the level of uncertainty participants integrate past and current trials differently. This approach applied to short-term influences found that higher uncertainty increased the tendency to alternate responses (Urai et al., 2017), while regarding the long-term changes, larger pupil size reliably predicted a smaller decision bias (Krishnamurthy et al., 2017). It would be pertinent to look at similar effects with our paradigm to detect if potential individual differences in bias could be linked to confidence related changes. In line with the above finding of (Krishnamurthy et al., 2017), we expect that, after the period of changing probability resulting in an increased uncertainty, the participants who readjusted their internal model would be more uncertain. This would lead to an increased willingness to adapt to changes in stimulus probabilities, resulting in a smaller bias. Naturally, some participants would exhibit more variability in their behavior, and thus adapt their internal model to the changed statistics, while others stay more rigid, resulting in stronger bias by long-term probabilities (Glaze et al., 2018). It remains to be seen whether such pupil-size-based assessment of uncertainty could link individual variability in adapting to changing probabilities to the predicted bias-variance trade-off across participants (Glaze et al., 2018).

Such an uncertainty-dependent updating could form a link between our two main lines of inquiry as uncertainty could play a pivotal role at different levels in both paradigms. In Study 2, there was a large uncertainty about momentary stimulus identity (as it was jointly influenced by noise and appearance probabilities). In contrast, in Study 1, there was no uncertainty about the stimulus at the currently fixated location, however there was a large uncertainty about the contents of the other areas of the scene. By learning the predictive spatial regularities of pairs, the observer's uncertainty in predicted states of the scenes should decrease. To confirm such a link between subjective

uncertainty and spatial statistical learning in Study 1, pupil size measure can be once again a useful tool. If scenes that violate the learned statistical structure are introduced into the stimulus presentation stream, they should elicit pupil size changes in a learning-dependent manner (Kloosterman et al., 2015).

Sampling of episodic experience

A recent line of research argues that human explorative decisions are influenced by reminders of past contexts: episodic experience is sampled to modulate momentary choices (Bornstein, Khaw, Shohamy, & Daw, 2017; Bornstein & Norman, 2017). The same experimental paradigm has also been applied to perceptual decision making, showing that samples of episodic experience can bias evidence accumulation as well (Bornstein, Aly, et al., 2017). These experiments used reminders of past stimulus contexts, thereby enhancing the retrieval of particular episodic experiences, which in turn influence momentary perceptual decisions based on their stimulus association strength. This integrated view on memory retrieval and perceptual decision making is in line with the proposal of a recent review paper suggesting a parallel between sampling sensory information and sampling from memory (Shadlen & Shohamy, 2016).

Our experiments in Chapters 4-5 did not use explicit reminders of past contexts, but we posit that certain periods of past experience (after the change-points) were deemed more significant. We hypothesize that samples of information collected during this critical period were stored/used with a higher weight than later samples. In essence, it can be argued that our experimental paradigm can be likened to the integrated view above by claiming that the context in our case was heavily determined by the change-points. Since there were no more drastic shifts in stimulus statistics during our experiment, participants implicitly assumed that they remained in the same context afterwards, and therefore, during the formation of their perceptual decisions they relied more heavily on the samples collected soon after the change.

Since the explicitness of the reminders are not a necessary feature for the integrated theory arguing for past episodic influence, this approach could be integrated with our findings. It remains to be seen whether predictions from this proposed approach using the idea of sampling from the critical period could be disentangled from or merged with the predictions of the “changing internal model parameters” account outlined above. If validated with behavioral data, this would be an interesting theoretical advancement arguing that people sample the summary statistics of past trials similarly to episodic content.

Links to reasoning

In Chapter 4, we have shown a remarkable misuse of stimulus probabilities in perceptual decision making. While, to our knowledge, this finding is novel in the domain of perceptual decision-making, it has interesting links to the reasoning and categorization literature. In reasoning studies, the misuse of probabilities has been known for a long-time, the classical finding being the ignorance of base-rate probabilities (Bar-Hillel, 1980; Kahneman & Tversky, 1972). However, it has also been shown that, depending on the context, people can use base-rate information appropriately or can exhibit a counter-intuitive inverse base-rate effect (Johansen, Fouquet, & Shanks, 2007; Medin & Edelson, 1988). This surprising inverse base-rate effect has been first described with a medical classification paradigm, showing that people judge novel ambiguous combinations of symptoms as manifestations of a disease with the lower base-rate (Medin & Edelson, 1988).

Our findings can be considered as an inverse base-rate effect in perceptual decision making, since noisy stimuli are erroneously judged as belonging to the rare category. While this is a remarkable similarity, indeed, the differences are also profound. First, in the medical categorization studies, participants were explicitly trained with unequal base-rates, whereas our effects emerged after a balanced training and unequal rates were introduced only during the unsupervised test. Second, a crucial component of the inverse base-rate effect is the introduction of novel symptom combinations

at test (Kruschke, 1996). In contrast, our paradigm uses the same noisy stimuli during training and test, making it unlikely that high-level reasoning would underlie our findings.

Despite these differences, a finding has been described in a very recent study investigating how unequal base-rates affect categorization judgments that could form a link between our study and the inverse base-rate effect (Levari et al., 2018). Similar to the results of Exp 1 in Chapter 4, these authors found that making a stimulus more frequent can bias people away from choosing it, a phenomenon they named “prevalence-induced concept change”. They found that such probability changes affected the categorization of a wide range of stimuli, from colors to facial expressions and even ethical judgments.

While the described effect was achieved with an extreme probability shift (6% of trials from the rare category), the underlying phenomenon could be of a similar nature to the one we report in our studies. Levari et al. (2018) described their findings as if the “concept” of threat or colors changed due to the manipulation of the probabilities. We think a more principled explanation of their findings is a general influence of probability changes on how people categorize ambiguous visual stimuli, similarly to our results in Chapter 4. Imbalanced base-rates could make people uncertain about their ability to distinguish these ambiguous stimuli (uncertainty in the likelihood/noise model), therefore, they update their noise model instead of updating the prior, resulting in a compensatory choice bias. If there are similar underlying mechanisms between these two studies, those could form an interesting link between the domains of perceptual decision making and explicit categorization. In addition, our method of using temporary sudden shifts to elicit a lasting change in decision criteria could also extend the findings of Levari et al., by showing that their findings are not a sole consequence of a short-term compensatory bias.

While decisions about high-level concepts such as faces had been shown to be amenable to the same sequential biases on a short time-scale as low-level perceptual stimuli (Fischer & Whitney, 2014; Liberman et al., 2014), the link between high-level and low-level influences related to long-term


probabilities have been largely ignored in the literature. Our work paves the road to future research that could focus on whether the link between probability shift eliciting high-level biases could work similarly to those found with low-level biases, and whether more delicate changes similar to those in Chapter 4 and much less drastic than those in Levari et al. (2018) could elicit lasting influences in already balanced periods, as in Chapter 5.

Conclusions


Environmental statistics are continuously built into internal representations. These internal representations, in turn, influence visual decisions continuously at multiple time-scales. Eye-movement-based explorative decisions are influenced by complex spatial regularities, in a close interaction with learning, suggesting an integrated view on memory and visuo-spatial attention. At the same time, depending on the task, learning of complex regularities can take place without influencing eye-movements, confirming the automatic manifestation of statistical learning. Perceptual decisions about the identity of ambiguous visual stimuli are strongly influenced by probabilities of past occurrences. This influence depends on perceived changes in the environmental statistics, thereby pointing to a pivotal role of change dynamics in how past stimulus statistics influence the interpretation of momentary visual input. Taken together, this thesis advances our understanding of how past environmental regularities influence visual decisions under various conditions. We have shown that past statistical influences can scaffold successful interaction with the environment, however in other high uncertainty scenarios, past influences can lead to a wrongly adjusted internal model which can hinder performance. Past statistical biases might not always be beneficial in an experimental set-up, but they are important in real life, where successful interaction with our surroundings is only possible if expectations are well tuned to the regularities we encounter in our daily lives. Good decisions are only possible if they are based on appropriate estimations on how the world works.

Appendix A


Additional Analysis for Chapters 2-3




Appendix Figure. A.1. Model fit likelihood with different null models. To make sure that choice of null model is not responsible for our main results, we consider an alternative here which uses a random walk instead of the individual transition probability distribution (for Exps 2,1,3). As expected fitting M1 to an individual random null model (Solid line) fits the data better than using a general random null model (Dashed line). Y-axis is the negative log likelihood/trial after fitting M1. Individual null model is based on the individual empirical transition probability distribution as in the main text. Random null model is a random walk with 99% probability to adjacent cells and 1% to other non-adjacent cells. The probability of the different directions of transitions to adjacent cells was based on the empirical data-set (horizontal=.51, vertical=.36, or diagonal=.13, to other cells=.01). The average $-\log(L)$ is increasing over time, since it is the sum of the log likelihood on each trial, and participants performed more transition over time.



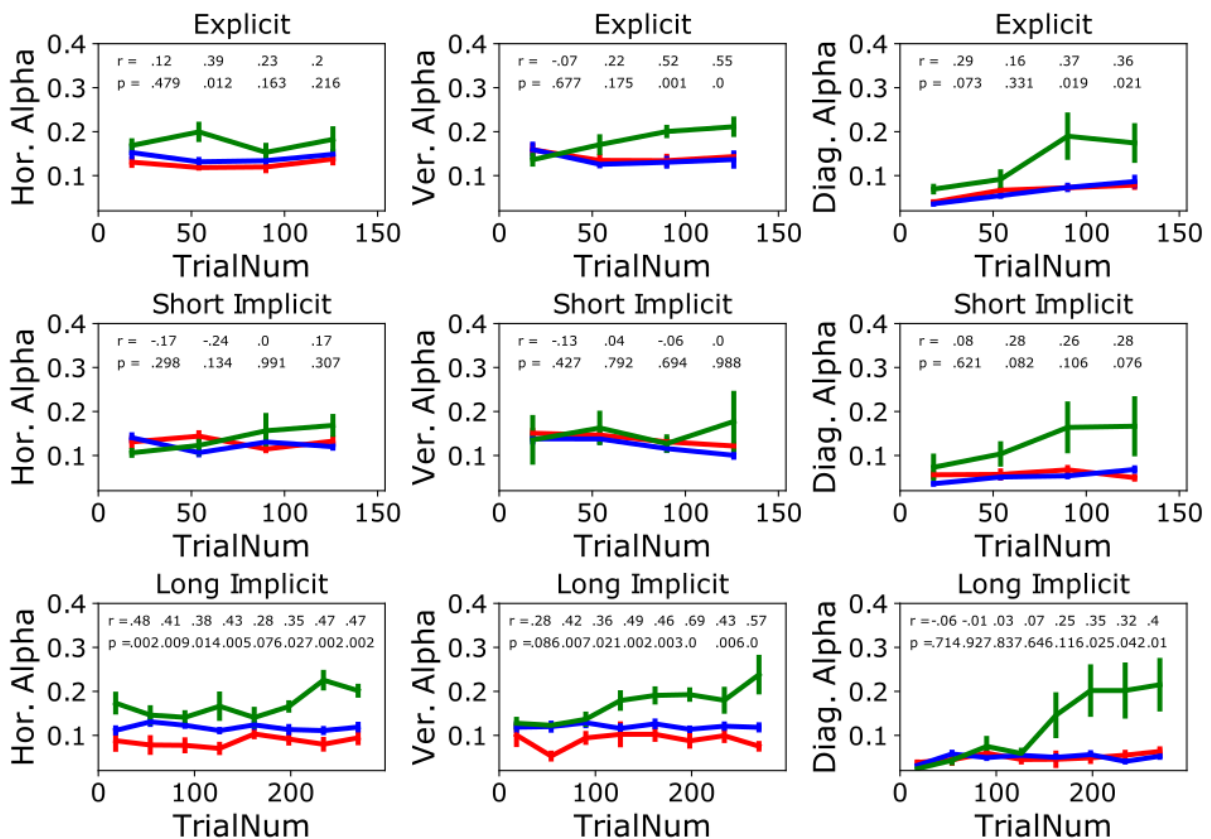
Appendix Figure A.2. Pair influence with alternative null models. The M_1 alpha parameter measuring pair influence was barely influenced by the choice of null model. The individual alpha values are binned by time and test performance as before for Exps 2,1,3. Solid line is the individual null model (as in main text), dashed is the random walk as defined on Appendix Fig A.1. This figure shows that the values of the pair influence parameter alpha are mostly very similar regardless of the choice of null model (despite the considerable difference in the $\log(L)$ of the fit as seen on the previous figure).



Appendix Figure. A.3. Fitting M_1 with larger alpha range. Alpha was constrained on the range $[-1, 2]$ instead of $[0,1]$ as in the main text. The correlation with familiarity test performance (shown on each figure) is significant for Exps 1-3. Furthermore, we can see that pair influence is very close to zero in the working memory experiments. Thus, how the values of alpha are constrained, influences the magnitude of fitted values, but the correlations with learning are very similar to those reported in the main text (or stronger in the case of the Long Implicit Experiment)



Appendix Figure. A.4. Model Selection Results for Exp 1-3, Color shows which model fits the overall exploration data of each participants better, as a function of pair transition rate (x-axis) and familiarity test performance(y-axis). Participants who are better fitted by M_1 tend to be high learners, who make lot of within pair transitions, however, the model is sensitive to pair influences where the overall transition rate does not show a difference from other participants. X axis: pair transition rate (N transitions within pairs/total number of transitions, not used in the main text) For this visualization we use this measure as this shows the closest relationship with our M_1 alpha parameter. Yellow: participants better fitted with M_1 , Blue Better fitted by M_0 . Horizontal lines show thresholds for grouping by familiarity test performance as in the main text. We can conclude that in all three Exps mostly high learners are fitted better by M_1 .



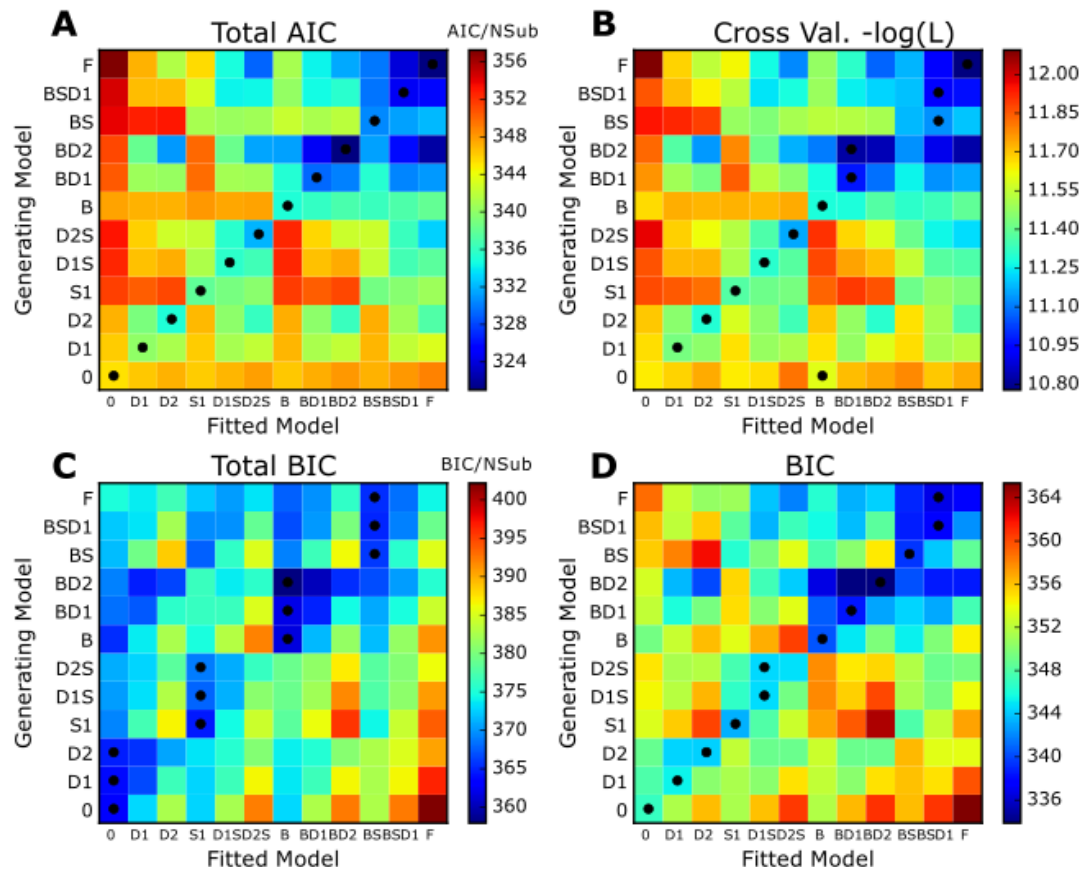
Appendix Figure. A.5. Specific pair influence over-time in Exps 1-3. Temporal direction specific pair influence for Exps 1,2,3 (Top,Middle,Bottom) for Horizontal-Vertical-Diagonal (Left, Middle, Right) statistical influence parameter α_{1-3} from M_2 . Time is shown on the x-axis, pair structure influence on y-axis, participants grouped by performance as before. Pearson correlation values are shown between pair influence parameter(α_{1-3}) and test performance on each direction, for each time-bin. Early emergence of Horizontal pair influence was true for Exp 1 & 3, which was followed by Vertical and Diagonal for both groups. No significant influence was found for any of the time-bins for any of the directions for Exp 2.

Experiment	Pair Exploration Rate	Pair Return Rate	M1 alpha	Shape Exploration Rate
1 Explicit	0.255 +/- 0.05	0.275 +/- 0.1	0.16 +/- 0.09	0.631 +/- 0.03
2 Short Implicit	0.24 +/- 0.04	0.225 +/- 0.04	0.134 +/- 0.07	0.624 +/- 0.02
3 Long Implicit	0.235 +/- 0.05	0.232 +/- 0.08	0.162 +/- 0.07	0.623 +/- 0.03
4 Implicit WM	0.255 +/- 0.03	0.273 +/- 0.06	0.103 +/- 0.03	0.636 +/- 0.01
5 Imp. WMshuf.	0.248 +/- 0.03	0.25 +/- 0.05	0.099 +/- 0.03	0.63 +/- 0.02

Appendix Table A.1. Descriptives of the statistical influence measures for Exps 1-5 on the overall data-set. These four measures were used to describe statistical influences in Chapters 2-3. Mean +/- SD

Appendix B

Additional Analysis for Chapters 4-5



Appendix Figure B.1. Validation of approach with simulation. Answers of 12 candidate models to the stimulus sequence presented to all participants of Exps 1-8 were evaluated, where the models included all combinations of Present Stimulus (null: 0), Bias (B), Previous Decision (D1), Previous two Decisions (D2), Previous Stimulus (S), and the full model with all 4 parameters (F). Results from a randomly selected round of simulation (12 generating * 12 fitted model to the stimulus sequence of all participants). Color code: best model=dark blue, worst model=dark red. For each generating model on the Y axis, a black dot shows the best Fitted model. Dots along the diagonal indicate when the true underlying model was found. AIC was the most accurate in finding the true mode (more info in Supplementary text), results of individually calculated and corrected AIC showed identical conclusions with the Total AIC, therefore we only show the latter. A) AIC calculated from total likelihood (divided by number of participants) B) Mean of Cross-Validated Negative Log Likelihood calculated individually C) BIC from total likelihood (divided by number of participants) D) Mean of individually calculated BIC.

	Bias	Current Stimulus	Decisions -1	Decision -2
Mean Weights	-0.05	1.104	0.189	0.202
CI	[-0.14, 0.04]	[1.04, 1.17]	[0.13, 0.24]	[0.17, 0.24]
Absolute Weights	0.427	1.104	0.312	0.2429
CI	[0.37, 0.49]	[1.04, 1.17]	[0.27, 0.35]	[0.22, 0.27]

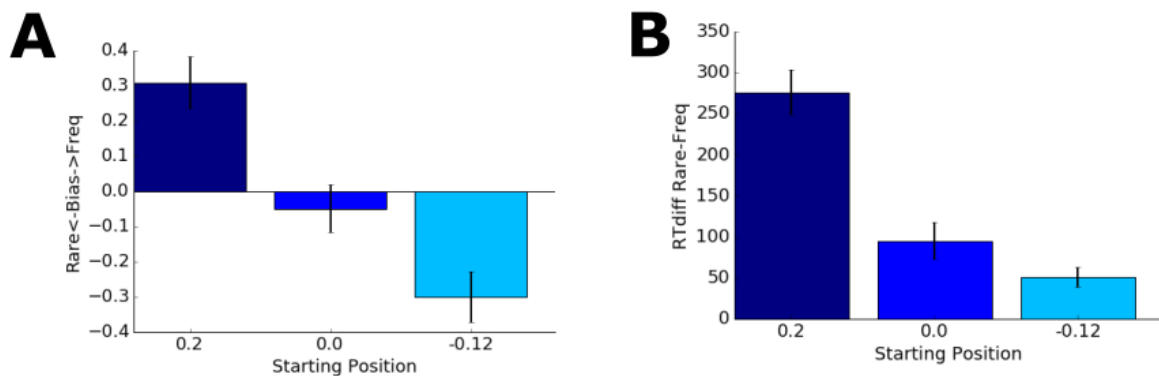
Appendix Table B.1. Logistic Model fit parameters on the overall data-set. These are the mean and absolute weights of the best fitting model with four parameters as selected on **Fig 4.1C**. Decision -1 &-2 refer to the previous decision and the decision two trials ago. Confidence intervals were obtained by bootstrapping 5000 times. On the overall dataset average bias is close to 0, looking at the absolute weights shows it has the second strongest influence on decisions. The predictor weights were standardized according to the method described in the main text to make the values comparable regardless of the scale of encoding the predictors

Appendix Text B.1.

Validation of Noise Model

To confirm that characteristics of the randomly selected shape stimuli and the Gaussian additive noise used in the experiment did not contribute to our results, we investigated how noise could affect the encoding of stimuli in our experiments. We used the N-dimensional pixel space for this analysis, where N denotes the number of pixels in the stimulus image and hence, each of the two noise-free target shape images as well as any noisy version of those images defined individual single points in the space. For each participant and each trial, we calculated the distance in this pixel space between the two noise-free targets and the noisy stimulus presented in the trial by projecting the noisy stimulus position onto the vector connecting the two noise-free targets in the pixel space. While the addition of random noise to the shape images in each trial introduced variability in the location where the projected line bisected the vector between the two targets, this location did not correlate with the participant's response bias. As a second control, instead of using a simple linear metric to quantify the noise levels based on the percentage perturbed pixels, we encoded the noise level of the stimuli based on vector distance obtained from the projection described above and replicated our analysis with the logistic regression model (Fig 4.1) with this encoding. The two

encoding methods yielded very similar results with the original linear encoding of noise performing slightly better in explaining the participants responses (obtaining a higher negative log likelihood). Based on these results, systematic biases in the noise model or interaction between target shapes and noisy stimuli could not explain our main results.



Appendix Figure B.2. Result patterns of Exps 3,2,1 replicated by drift diffusion model simulation (note the reverse order). **A)** Bias obtained from simulated data (Color code from dark to light blue Exp3, Exp2, Exp1) **B)** Results of RT differences in the simulated data. An asymmetric drift rate (faster drift for frequent stimulus) and a biased starting position could reproduce the main pattern of the findings in Exps 1-3, most notably the opposing effects on RT and decision bias in Exp 1. The same asymmetric scaling of the drift rate was used for each bar of the above figures, while the value of starting positions of the DDM were selected as noted on the x axis. The simulated RT results do not include non-decision time that would compress the differences between the three RT bars making them more similar to human results. (see detailed description in Supplementary Text)

Appendix Text B.2.

Drift Diffusion Model

In order to capture the RT and decision bias asymmetry in our data we independently manipulated two parameters in the standard DDM model (Ratcliff & Rouder, 1998). We simulated data with trial-by-trial DDM to the stimulus sequence of each participant from Exps 1-3, and fitted the above used logistic regression model to the simulated data. To fit the decision bias we manipulated the starting value of the diffusion process. In order to capture the opposing bias in RT-s, we scaled the accumulation rate of the diffusion process asymmetrically for the two stimuli. A simulation with the

starting value biased toward the rare element but with faster accumulation rate toward the frequent element could reproduce the main trend of our findings in Experiments 1-3 (Suppl Fig. 5).

Psychometric PSE	Experiment P Structure	df	Mean	CI	<i>t</i> -test <i>t</i>	<i>t</i> -test <i>p</i>	Cohen's <i>d</i>	Bayes Factor
Exp 1	50-65	19	0.5581	0.2616 0.8854	3.4228	.0029	0.7654	14.79
Exp 2	65-65	19	- 0.3419	-0.7936 0.1039	1.4668	.1588	-0.3280	.59
Exp 3	75-65	19	- 1.5722	-2.5611 -0.7718	3.3142	.0036	-0.7411	11.99
Exp 4	75-50	19	- 1.4927	-2.0547 -0.8735	4.8758	.0001	-1.0903	266.46
Exp 5	50-65Gr	19	- 0.6184	-1.3496 -0.0377	1.7578	.0949	-0.3931	.85
Exp 6	50-65-50Gr	19	0.6428	0.2229 1.1156	2.7514	.0127	0.6152	4.2
Exp 7	50-75-50Gr	20	0.4184	0.0360 0.7884	2.1013	.0485	0.4585	1.4

Appendix Table B.2. Point of subjective equivalence (PSE) of psychometric curves for experiments 1-7 with descriptive stats and statistical tests whether the average observer was biased in each experiment. CI is obtained via bootstrapping. *d* is effect size (Cohen's *d*). Bayes Factor. Here higher numbers mean a stronger preference for the rare stimulus in each experiment.

Logistic Model Bias	Mean	CI	<i>t</i> -test <i>t</i>	<i>t</i> -test <i>p</i>	Cohen's <i>d</i>	Bayes Factor
Exp 1	-0.336	-0.4815, -0.1907	4.431	.0003	0.9908,	109.3
Exp 2	-0.058	-0.2468, 0.1123	0.6227	.5409	0.1392	0.28
Exp 3	0.338	0.1077, 0.5816	2.7457	.0129	0.614	4.16
Exp 4	0.418	0.1367, 0.7208	2.7406	.0130	0.6128	4.12
Exp 5	-0.001	-0.1539, 0.1669	0.0073	.9943	0.0016	.23
Exp 6	-0.35	-0.514 , -0.196	4.1714	.0005	0.9328	64.96
Exp 7	-0.347	-0.5394, -0.1514	3.4377	.0026	0.7502	15.22

Appendix Table B.3. Logistic model bias for Exps 1-7 with descriptive stats and statistical tests whether the overall bias is different from zero. Values below zero represent a bias toward the rare, above zero toward the frequent stimulus (unlike in the previous table with PSE)

Logistic Model Bias over-time	Initial	1 st Half	2 nd Half
Exp 1 Mean CI	0.049 [-0.26 0.35]	-0.303 [-0.46 -0.14]	-0.326 [-0.48 -0.17]
Exp 2 Mean CI	0.381 [0.11 0.65]	0.04 [-0.15 0.24]	-0.129 [-0.33 0.07]
Exp 3 Mean CI	0.57 [0.31 0.8]	0.409 [0.17 0.63]	0.309 [0.06 0.57]
Exp 4 Mean CI	0.36 [-0. 0.7]	0.387 [0.13 0.66]	0.423 [0.14 0.72]
Exp 5 Mean CI	0.1 [-0.18 0.36]	0.092 [-0.12 0.34]	0.001 [-0.14 0.16]
Exp 6 Mean CI	0.213 [0.01 0.42]	-0.269 [-0.47 -0.07]	-0.373 [-0.52 -0.23]
Exp 7 Mean CI	0.172 [-0.08 0.42]	-0.373 [-0.55 -0.2]	-0.277 [-0.49 -0.07]

Appendix Table B.4. Decision Bias over-time for Exps 1-7. Initial is calculated based on the first 30 trials of the test block, to see the bias during initial test exposure. 1st and 2nd half are the two halves of the entire test sessions for Exps 1-4, and the two halves of the 300 trial period after the gradual change for Exps 5-7. Confidence intervals were obtained by bootstrapping. Values below zero represent bias toward the rare, above zero toward the frequent stimulus.

RTdiff	Mean (ms)	CI (bootstrap)	t-test t	t-test p	Effect size (d)	Wilcoxon Z	Wilcoxon p	BF
Exp 1	53.198	20.2194 83.2898	3.2096	.0046	0.7177	3.0239	.0025	9.82
Exp 2	29.975	-3.8205 60.4835	1.8	.0878	0.4025	1.9786	.0479	0.9
Exp 3	74.318	35.8719 117.955	3.4559	.0026	0.7728	3.0239	.0025	15.77
Exp 4	33.985	0.0889 67.7941	1.937	.0678	0.4331	1.6053	.1084	1.1
Exp 5	56.335	32.1521 82.6002	4.2201	.0005	0.9436	3.4346	.0006	71.61
Exp 6	3.882	-28.7359 19.7899	0.3058	.7631	0.0684	0.2987	.7652	0.24
Exp 7	13.893	-33.9326 5.8561	1.3357	.1966	0.2915	1.3034	.1924	0.49

Appendix Table B.5. Descriptives of reaction time difference (rare-frequent), higher values mean bigger advantage of frequent responses. CI is obtained by bootstrapping. In testing whether the data is different from zero, we report both t-test and non-parametric statistics, since there were some strong outliers in the RT data as can be seen in Fig 4.4D. Bayes factor is calculated from t-value.

References

- Acerbi, L., Vijayakumar, S., & Wolpert, D. M. (2014). On the origins of suboptimality in human probabilistic inference. *PLoS Computational Biology*, 10(6), e1003661. <https://doi.org/10.1371/journal.pcbi.1003661>
- Ackermann, J. F., & Landy, M. S. (2014). Suboptimal decision criteria are predicted by subjectively weighted probabilities and rewards. *Attention, Perception & Psychophysics*, 638–658. <https://doi.org/10.3758/s13414-014-0779-z>
- Ahmad, S., Jolla, L., & Yu, A. J. (2013). Active Sensing as Bayes-Optimal Sequential Decision-Making. *ArXiv*, arXiv:1305.
- Akaishi, R., Umeda, K., Nagase, A., & Sakai, K. (2014). Autonomous mechanism of internal choice estimate underlies decision inertia. *Neuron*, 81(1), 195–206. <https://doi.org/10.1016/j.neuron.2013.10.018>
- Ayton, P., & Fischer, I. (2004). The hot hand fallacy and the gambler's fallacy: two faces of subjective randomness? *Memory & Cognition*, 32(8), 1369–1378. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/15900930>
- Bar-Hillel, M. (1980). The base-rate fallacy in probability judgments. *Acta Psychologica*, 44(3), 211–233. [https://doi.org/10.1016/0001-6918\(80\)90046-3](https://doi.org/10.1016/0001-6918(80)90046-3)
- 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. <https://doi.org/10.1016/j.cognition.2013.07.003>
- Barlow, H. B., & Hill, R. M. (1963). Evidence for a physiological explanation of the waterfall phenomenon and figural after-effects. *Nature*, 200(4913), 1345.
- Behrens, T. E. J., Woolrich, M. W., Walton, M. E., & Rushworth, M. F. S. (2007). Learning the value of information in an uncertain world. *Nature Neuroscience*, 10(9), 1214–1221. <https://doi.org/10.1038/nn1954>
- Bell, A. H., Summerfield, C., Morin, E. L., Malecek, N. J., & Ungerleider, L. G. (2016). Encoding of Stimulus Probability in Macaque Inferior Temporal Cortex. *Current Biology*, 26(17), 2280–2290. <https://doi.org/10.1016/j.cub.2016.07.007>
- Bell, A. H., Summerfield, C., Morin, E. L., Malecek, N. J., & Ungerleider, L. G. (2017). Reply to Vinken and Vogels. *Current Biology*, 27(22), R1212–R1213. <https://doi.org/10.1016/j.cub.2017.09.021>
- Bertels, J., Boursain, E., Destrebecqz, A., & Gaillard, V. (2014). Visual statistical learning in children and young adults: How implicit? *Frontiers in Psychology*, 5(OCT), 1–11. <https://doi.org/10.3389/fpsyg.2014.01541>
- Bertels, J., Franco, A., & Destrebecqz, A. (2012). How implicit is visual statistical learning? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 38(5), 1425.
- Bishara, A. J., & Hittner, J. B. (2012). Testing the significance of a correlation with nonnormal data: comparison of Pearson, Spearman, transformation, and resampling approaches. *Psychological Methods*, 17(3), 399.

- Bohil, C. J., & Wismer, A. J. (2014). Implicit learning mediates base rate acquisition in perceptual categorization. *Psychonomic Bulletin & Review*, 22(2), 586–593. <https://doi.org/10.3758/s13423-014-0694-2>
- Borji, A., & Itti, L. (2013). State-of-the-art in visual attention modeling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(1), 185–207. <https://doi.org/10.1109/TPAMI.2012.89>
- Bornstein, A. M., Aly, M., Feng, S., Turk-Browne, N. B., Norman, K. A., & Cohen, J. D. (2017). Perceptual decisions result from the continuous accumulation of memory and sensory evidence. *BioRxiv*, 186817. <https://doi.org/10.1101/186817>
- Bornstein, A. M., Khaw, M. W., Shohamy, D., & Daw, N. D. (2017). What's past is present: Reminders of past choices bias decisions for reward in humans. *Nature Communications*, 8(June 2017), 16–21. <https://doi.org/10.1101/033910>
- Bornstein, A. M., & Norman, K. A. (2017). Reinstated episodic context guides sampling-based decisions for reward. *Nature Neuroscience*, 20(7), 997–1003. <https://doi.org/10.1038/nn.4573>
- Brady, T. F., Konkle, T., & Alvarez, G. a. (2009). Compression in visual working memory: using statistical regularities to form more efficient memory representations. *Journal of Experimental Psychology. General*, 138(4), 487–502. <https://doi.org/10.1037/a0016797>
- Brady, T. F., Störmer, V. S., & Alvarez, G. A. (2016). Working memory is not fixed-capacity: More active storage capacity for real-world objects than for simple stimuli. *Proceedings of the National Academy of Sciences*, 113(27), 7459–7464. <https://doi.org/10.1073/pnas.1520027113>
- Braun, A., Urai, A. E., & Donner, T. H. (2018). Adaptive History Biases Result from Confidence-weighted Accumulation of Past Choices. *The Journal of Neuroscience*, 38(10), 2189–17. <https://doi.org/10.1523/JNEUROSCI.2189-17.2017>
- Bronfman, Z. Z., Brezis, N., Moran, R., Tsetsos, K., Donner, T., & Usher, M. (2015). Decisions reduce sensitivity to subsequent information. *Proceedings of the Royal Society B: Biological Sciences*, 282(1810), 1–10. <https://doi.org/10.1098/rspb.2015.0228>
- Burns, B. D., & Corpus, B. (2004). Randomness and inductions from streaks: “Gambler’s fallacy” versus “hot hand”. *Psychonomic Bulletin & Review*, 11(1), 179–184.
- Burns, K. J., & Demaree, H. A. (2009). A chance to learn: On matching probabilities to optimize utilities. *Information Sciences*, 179(11), 1599–1607.
- Cho, R. Y., Nystrom, L. E., Brown, E. T., Jones, A. D., Braver, T. S., Holmes, P. J., & Cohen, J. D. (2002). Mechanisms underlying dependencies of performance on stimulus history in a two-alternative forced-choice task. *Cognitive, Affective & Behavioral Neuroscience*, 2(4), 283–299. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/12641174>
- Chopin, A., & Mamassian, P. (2012). Predictive properties of visual adaptation. *Current Biology*, 22(7), 622–626. <https://doi.org/10.1016/j.cub.2012.02.021>
- Chopin, A., & Mamassian, P. (2013). Response: Genuine long-term positive aftereffects. *Current Biology : CB*, 23(10), R439-40. <https://doi.org/10.1016/j.cub.2013.03.025>
- Chun, M. M., & Jiang, Y. (1998). Contextual Cueing: Implicit Learning and Memory of Visual Context Guides Spatial Attention. *Cognitive Psychology*, 36(1), 28–71.
- Chun, M. M., & Jiang, Y. (1999). Top-down attentional guidance based on implicit learning of visual covariation. *Psychological Science*, 10(4), 360–365. Retrieved from <http://journals.sagepub.com/doi/pdf/10.1111/1467-9280.00168>

- Chun, M. M., & Turk-Browne, N. B. (2007). Interactions between attention and memory. *Current Opinion in Neurobiology*, 17(2), 177–184. <https://doi.org/10.1016/j.conb.2007.03.005>
- Cicchini, G. M., Mikellidou, K., & Burr, D. (2017). Serial dependencies act directly on perception. *Journal of Vision*, 17(14), 6. <https://doi.org/10.1167/17.14.6>
- Conway, C. M., & Christiansen, M. H. (2005). Modality-constrained statistical learning of tactile, visual, and auditory sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 31(1), 24.
- Craig, C. C., & Fisher, R. A. (1936). The Design of Experiments. *The American Mathematical Monthly*, 43(3), 180. <https://doi.org/10.2307/2300364>
- Dienes, Z. (2014). Using Bayes to get the most out of non-significant results. *Frontiers in Psychology*, 5(July), 1–17. <https://doi.org/10.3389/fpsyg.2014.00781>
- Drugowitsch, J., Wyart, V., Devauchelle, A.-D., & Koechlin, E. (2016). Computational Precision of Mental Inference as Critical Source of Human Choice Suboptimality. *Neuron*, 92(6), 1–14. <https://doi.org/10.1016/j.neuron.2016.11.005>
- Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review*, 96(3), 433.
- Ellis, R. (2009). Implicit and explicit learning, knowledge and instruction. *Implicit and Explicit Knowledge in Second Language Learning, Testing and Teaching*, 42, 3–25.
- Ernst, M. O., & Banks, M. S. (2002). Humans integrate visual and haptic information in a statistically optimal fashion. *Nature*, 415(6870), 429–433. <https://doi.org/10.1038/415429a>
- Firestone, C., & Scholl, B. J. (2014). “Top-down” effects where none should be found: the El Greco fallacy in perception research. *Psychological Science*, 25(1), 38–46. <https://doi.org/10.1177/0956797613485092>
- Fischer, J., & Whitney, D. (2014). Serial dependence in visual perception. *Nature Neuroscience*, 17(5), 738–743. <https://doi.org/10.1038/nn.3689>
- Fiser, J. (2009). Perceptual learning and representational learning in humans and animals. *Learning & Behavior*, 37(2), 141–153.
- Fiser, J., & Aslin, R. N. (2001). Unsupervised statistical learning of higher-order spatial structures from visual scenes. *Psychological Science*, 12(6), 499–504. <https://doi.org/10.1111/1467-9280.00392>
- 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.
- 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. <https://doi.org/10.1016/j.tics.2010.01.003>
- Foley, N. C., Kelly, S. P., Mhatre, H., Lopes, M., & Gottlieb, J. (2017). Parietal neurons encode expected gains in instrumental information. *Proceedings of the National Academy of Sciences*, 114(16), E3315–E3323.
- Friston, K., Adams, R. a, Perrinet, L., & Breakspear, M. (2012). Perceptions as hypotheses: saccades as experiments. *Frontiers in Psychology*, 3(May), 151. <https://doi.org/10.3389/fpsyg.2012.00151>
- Fritsche, M., Mostert, P., & de Lange, F. P. De. (2017). Opposite Effects of Recent History on Perception and Decision Report Opposite Effects of Recent History on Perception and Decision.

- Current Biology*, 1–6. <https://doi.org/10.1016/j.cub.2017.01.006>
- Fründ, I., Wichmann, F. A., & Macke, J. H. (2014). Quantifying the effect of intertrial dependence on perceptual decisions. *Journal of Vision*, 14(7), 1–16. <https://doi.org/10.1167/14.7.9>
- Gaissmaier, W., & Schooler, L. J. (2008). The smart potential behind probability matching. *Cognition*, 109(3), 416–422. <https://doi.org/10.1016/j.cognition.2008.09.007>
- Gallistel, C. R., Krishan, M., Liu, Y., Miller, R., & Latham, P. E. (2014). The perception of probability. *Psychological Review*, 121(1), 96–123. <https://doi.org/10.1037/a0035232>
- Gibson, B. R., Rogers, T. T., & Zhu, X. (2013). Human Semi-Supervised Learning. *Topics in Cognitive Science*, 5(1), 132–172.
- Glaze, C. M., Filipowicz, A. L. S., Kable, J. W., Balasubramanian, V., & Gold, J. I. (2018). A bias-variance trade-off governs individual differences in on-line learning in an unpredictable environment. *Nature Human Behaviour*, 2(3), 213–224. <https://doi.org/10.1038/s41562-018-0297-4>
- Glaze, C. M., Kable, J. W., & Gold, J. I. (2015). Normative evidence accumulation in unpredictable environments, (August). <https://doi.org/10.7554/eLife.08825>
- Gold, J. I., & Stocker, A. A. (2017). Visual Decision-Making in an Uncertain and Dynamic World. *Annual Review of Vision Science*, 3(1), 227–250. <https://doi.org/10.1146/annurev-vision-111815-114511>
- Goodnow, J. (1955). Determinants of Choice-Distribution in Two-Choice Situations. *University of Illinois Press*, 68(1), 106–116.
- Gottlieb, J. (2012). Attention, learning, and the value of information. *Neuron*, 76(2), 281–295. <https://doi.org/10.1016/j.neuron.2012.09.034>
- Gottlieb, J., Hayhoe, M., Hikosaka, O., & Rangel, A. (2014). Attention, Reward, and Information Seeking. *Journal of Neuroscience*, 34(46), 15497–15504. <https://doi.org/10.1523/JNEUROSCI.3270-14.2014>
- Gottlieb, J., Oudeyer, P.-Y., Lopes, M., & Baranes, A. (2013). Information-seeking, curiosity, and attention: computational and neural mechanisms. *Trends in Cognitive Sciences*, 17(11), 585–593. <https://doi.org/10.1016/j.tics.2013.09.001>
- Graf, P., & Schacter, D. L. (1985). Implicit and explicit memory for new associations in normal and amnesic subjects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11(3), 501.
- Green, D. M., & Swets, J. A. (1966). Signal detection theory and psychophysics. *Wiley; New York*.
- Gregory, R. L. (1970). The Intelligent Eye. Retrieved from <https://eric.ed.gov/?id=ED064889>
- Grotheer, M., & Kovacs, G. (2014). Repetition Probability Effects Depend on Prior Experiences. *Journal of Neuroscience*, 34(19), 6640–6646. <https://doi.org/10.1523/JNEUROSCI.5326-13.2014>
- Hanks, T. D., Mazurek, M. E., Kiani, R., Hopp, E., & Shadlen, M. N. (2011). Elapsed decision time affects the weighting of prior probability in a perceptual decision task. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, 31(17), 6339–6352. <https://doi.org/10.1523/JNEUROSCI.5613-10.2011>
- Hannula, D. E. (2010). Worth a glance: using eye movements to investigate the cognitive neuroscience of memory. *Frontiers in Human Neuroscience*, 4(October), 1–16. <https://doi.org/10.3389/fnhum.2010.00166>

- Hayhoe, M., & Ballard, D. (2005). Eye movements in natural behavior. *Trends in Cognitive Sciences*, 9(4), 188–194. <https://doi.org/10.1016/j.tics.2005.02.009>
- Hollingworth, A., Williams, C. C., & Henderson, J. M. (2001). To see and remember: Visually specific information is retained in memory from previously attended objects in natural scenes. *Psychonomic Bulletin and Review*, 8(4), 761–768. <https://doi.org/10.3758/BF03196215>
- Hoppe, D., Helfmann, S., & Rothkopf, C. A. (2018). Humans quickly learn to blink strategically in response to environmental task demands. *Proceedings of the National Academy of Sciences*, 201714220.
- Hoppe, D., & Rothkopf, C. A. (2016). Learning rational temporal eye movement strategies. *Proceedings of the National Academy of Sciences*, 113(29), 8332–8337. <https://doi.org/10.1073/pnas.1601305113>
- Hoppe, D., & Rothkopf, C. A. (2017). Optimal planning of eye movements. *BioRxiv*, 240010.
- Horowitz, T. S., & Wolfe, J. M. (1998). Visual search has no memory. *Nature*, 394(6693), 575–577. <https://doi.org/10.1038/29068>
- Howard Jr, J. H., & Howard, D. V. (1997). Age differences in implicit learning of higher order dependencies in serial patterns. *Psychology and Aging*, 12(4), 634.
- Huelsenbeck, J. P., & Crandall, K. A. (1997). Phylogeny estimation and hypothesis testing using maximum likelihood. *Annual Review of Ecology and Systematics*, 28(1), 437–466.
- Itti, L., & Baldi, P. (2009). Bayesian surprise attracts human attention. *Vision Research*, 49(10), 1295–1306. <https://doi.org/10.1016/j.visres.2008.09.007>
- Jepma, M., Verdonschot, R. G., van Steenbergen, H., Rombouts, S. A. R. B., & Nieuwenhuis, S. (2012). Neural mechanisms underlying the induction and relief of perceptual curiosity. *Frontiers in Behavioral Neuroscience*, 6(February), 1–9. <https://doi.org/10.3389/fnbeh.2012.00005>
- Jiang, J., Summerfield, C., & Egner, T. (2013). Attention Sharpens the Distinction between Expected and Unexpected Percepts in the Visual Brain. *Journal of Neuroscience*, 33(47), 18438–18447. <https://doi.org/10.1523/JNEUROSCI.3308-13.2013>
- Jiang, Y. V, Won, B.-Y., & Swallow, K. M. (2014). First saccadic eye movement reveals persistent attentional guidance by implicit learning. *Journal of Experimental Psychology. Human Perception and Performance*, 40(3), 1161–1173. <https://doi.org/10.1037/a0035961>
- Johansen, M. K., Fouquet, N., & Shanks, D. R. (2007). Paradoxical effects of base rates and representation in category learning. *Memory & Cognition*, 35(6), 1365–1379.
- Jones, J. L., & Kaschak, M. P. (2012). Global statistical learning in a visual search task. *Journal of Experimental Psychology: Human Perception and Performance*, 38(1), 152.
- Jones, M., & Dzharfarov, E. N. (2014). Unfalsifiability and mutual translatability of major modeling schemes for choice reaction time. *Psychological Review*, 121(1), 1.
- Jungé, J. A., Scholl, B. J., & Chun, M. M. (2007). How is spatial context learning integrated over signal versus noise? A primacy effect in contextual cueing. *Visual Cognition*, 15(1), 1–11. <https://doi.org/10.1080/13506280600859706>
- Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 3(3), 430–454.
- Kaliukhovich, D. A., & Vogels, R. (2010). Stimulus repetition probability does not affect repetition suppression in macaque inferior temporal cortex. *Cerebral Cortex*, 21(7), 1547–1558.

- Kaneko, Y., & Sakai, K. (2015). Dissociation in decision bias mechanism between probabilistic information and previous decision. *Frontiers in Human Neuroscience*, 9(May), 1–11. <https://doi.org/10.3389/fnhum.2015.00261>
- Kang, M. J., Hsu, M., Krajbich, I. M., Loewenstein, G., McClure, S. M., Wang, J. T., & Camerer, C. F. (2009). The wick in the candle of learning. *Psychological Science*, 20(8), 963–974. <https://doi.org/10.1111/j.1467-9280.2009.02402.x>
- Karuza, E. A., Farmer, T. A., Alex, B., Smith, F. X., & Jaeger, T. F. (2014). On-line Measures of Prediction in a Self - Paced Statistical Learning Task. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 1(36), 725–730.
- Kersten, D., Mamassian, P., & Yuille, A. (2004). Object perception as Bayesian inference. *Annu. Rev. Psychol.*, 55, 271–304.
- Kidd, C., & Hayden, B. Y. (2015). The Psychology and Neuroscience of Curiosity. *Neuron*, 88(3), 449–460. <https://doi.org/10.1016/j.neuron.2015.09.010>
- Kidd, C., Palmeri, H., & Aslin, R. N. (2012). Rational snacking : Young children ' s decision-making on the marshmallow task is moderated by beliefs about environmental reliability.
- Kidd, C., Piantadosi, S. T., & Aslin, R. N. (2012). The Goldilocks Effect: Human Infants Allocate Attention to Visual Sequences That Are Neither Too Simple Nor Too Complex. *PLoS ONE*, 7(5), e36399. <https://doi.org/10.1371/journal.pone.0036399>
- Kidd, C., Piantadosi, S. T., & Aslin, R. N. (2014). The Goldilocks effect in infant auditory attention. *Child Development*, 85(5), 1795–1804.
- 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. <https://doi.org/10.1016/j.neulet.2009.06.030>
- Kloosterman, N. A., Meindertsma, T., van Loon, A. M., Lamme, V. A. F., Bonneh, Y. S., & Donner, T. H. (2015). Pupil size tracks perceptual content and surprise. *European Journal of Neuroscience*, 41(8), 1068–1078. <https://doi.org/10.1111/ejn.12859>
- Koblinger, Á., Arató, J., & Fiser, J. (2018). Complex adaptive internal model subserves perceptual sequential decision making. In *COSYNE abstract* (p. III-46, 216).
- Kok, P., Bains, L. J., van Mourik, T., Norris, D. G., & de Lange, F. P. (2016). Selective Activation of the Deep Layers of the Human Primary Visual Cortex by Top-Down Feedback. *Current Biology*, 26(3), 371–376. <https://doi.org/10.1016/j.cub.2015.12.038>
- Kok, P., Jehee, J. F. M., & de Lange, F. P. (2012). Less Is More: Expectation Sharpens Representations in the Primary Visual Cortex. *Neuron*, 75(2), 265–270. <https://doi.org/10.1016/j.neuron.2012.04.034>
- Körding, K. P., & Wolpert, D. M. (2004). Bayesian integration in sensorimotor learning. *Nature*, 427(6971), 244–247. <https://doi.org/10.1038/nature02169>
- Koski, J., Olson, I. R., & Newcombe, N. S. (2013). Tracking the eyes to see what children remember. *Memory*, 21(3), 396–407. <https://doi.org/10.1080/09658211.2012.735241>
- Kovács, G., Iffland, L., Vidnyánszky, Z., & Greenlee, M. W. (2012). Stimulus repetition probability effects on repetition suppression are position invariant for faces. *NeuroImage*, 60(4), 2128–2135. <https://doi.org/10.1016/j.neuroimage.2012.02.038>
- Kovacs, G., Kaiser, D., Kaliukhovich, D. A., Vidnyanszky, Z., & Vogels, R. (2013). Repetition Probability

- Does Not Affect fMRI Repetition Suppression for Objects. *Journal of Neuroscience*, 33(23), 9805–9812. <https://doi.org/10.1523/JNEUROSCI.3423-12.2013>
- Krishnamurthy, K., Nassar, M. R., Sarode, S., & Gold, J. I. (2017). Arousal-related adjustments of perceptual biases optimize perception in dynamic environments. *Nature Human Behaviour*, 1(6). <https://doi.org/10.1038/s41562-017-0107>
- Kruschke, J. K. (1996). Base rates in category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22(1), 3.
- Kunar, M. A., Flusberg, S., Horowitz, T. S., & Wolfe, J. M. (2007). Does Contextual Cuing Guide the Deployment of Attention? *Journal of Experimental Psychology: Human Perception and Performance*, 33(4), 816–828. <https://doi.org/10.1037/0096-1523.33.4.816>
- Laquitaine, S., & Gardner, J. L. (2018). A Switching Observer for Human Perceptual Estimation. *Neuron*, 97(2), 462–474.e6. <https://doi.org/10.1016/j.neuron.2017.12.011>
- Levari, D. E., Gilbert, D. T., Wilson, T. D., Sievers, B., Amodio, D. M., & Wheatley, T. (2018). Prevalence-induced concept change in human judgment. *Science*, 360(6396), 1465–1467. <https://doi.org/10.1126/SCIENCE.AAP8731>
- Lieberman, A., Fischer, J., & Whitney, D. (2014). Serial Dependence in the Perception of Faces. *Current Biology : CB*, 24(21), 2569–2574. <https://doi.org/10.1016/j.cub.2014.09.025>
- Loftus, G. R., & Mackworth, N. H. (1978). Cognitive determinants of fixation location during picture viewing. *Journal of Experimental Psychology: Human Perception and Performance*, 4(4), 565.
- Maljkovic, V., & Nakayama, K. (1994). Priming of pop-out: I. Role of features. *Memory & Cognition*, 22(6), 657–672.
- Maljkovic, V., & Nakayama, K. (1996). Priming of pop-out: II. The role of position. *Perception & Psychophysics*, 58(7), 977–991.
- Maloney, L. T., Dal Martello, M. F., Sahm, C., & Spillmann, L. (2005). Past trials influence perception of ambiguous motion quartets through pattern completion. *Proceedings of the National Academy of Sciences of the United States of America*, 102(8), 3164–3169. <https://doi.org/10.1073/pnas.0407157102>
- Markant, D. B., & Gureckis, T. M. (2014). Is it better to select or to receive? Learning via active and passive hypothesis testing. *Journal of Experimental Psychology: General*, 143(1), 94–122. <https://doi.org/10.1037/a0032108>
- Maus, G. W., Chaney, W., Liberman, A., & Whitney, D. (2013). The challenge of measuring long-term positive aftereffects. *Current Biology : CB*, 23(10), R438-9. <https://doi.org/10.1016/j.cub.2013.03.024>
- McDonald, S. A., & Shillcock, R. C. (2003). Eye movements reveal the on-line computation of lexical probabilities during reading. *Psychological Science*, 14(6), 648–652. https://doi.org/10.1046/j.0956-7976.2003.psci_1480.x
- Medin, D. L., & Edelson, S. M. (1988). Problem structure and the use of base-rate information from experience. *Journal of Experimental Psychology: General*, 117(1), 68.
- Menard, S. (2004). Six Approaches to Calculating Standardized Logistic Regression Coefficients. *The American Statistician*, 58(3), 218–223.
- Morvan, C., & Maloney, L. T. (2012). Human visual search does not maximize the post-saccadic probability of identifying targets. *PLoS Computational Biology*, 8(2), e1002342.

<https://doi.org/10.1371/journal.pcbi.1002342>

- Mulder, M. J., Wagenmakers, E.-J., Ratcliff, R., Boekel, W., & Forstmann, B. U. (2012). Bias in the Brain: A Diffusion Model Analysis of Prior Probability and Potential Payoff. *Journal of Neuroscience*, 32(7), 2335–2343. <https://doi.org/10.1523/JNEUROSCI.4156-11.2012>
- Najemnik, J., & Geisler, W. S. (2005). Optimal eye movement strategies in visual search. *Nature*, 434(7031), 387–391. <https://doi.org/10.1038/nature03390>
- Nassar, M. R., Helmers, J. C., & Frank, M. J. (2018). Chunking as a rational strategy for lossy data compression in visual working memory. *Psychological Review*, 125(4), 486–511. <https://doi.org/10.1037/rev0000101>
- Nassar, M. R., Wilson, R. C., Heasley, B., & Gold, J. I. (2010). An Approximately Bayesian Delta-Rule Model Explains the Dynamics of Belief Updating in a Changing Environment. *Journal of Neuroscience*, 30(37), 12366–12378. <https://doi.org/10.1523/JNEUROSCI.0822-10.2010>
- Nelson, J. D., Divjak, B., Gudmundsdottir, G., Martignon, L. F., & Meder, B. (2014). Children's sequential information search is sensitive to environmental probabilities. *Cognition*, 130(1), 74–80. <https://doi.org/10.1016/j.cognition.2013.09.007>
- Nemeth, D., Janacek, K., & Fiser, J. (2013). Age-dependent and coordinated shift in performance between implicit and explicit skill learning. *Frontiers in Computational Neuroscience*, 7(October), 147. <https://doi.org/10.3389/fncom.2013.00147>
- Nissen, M. J., & Bullemer, P. (1987). Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology*, 19(1), 1–32.
- Norton, E. H., Fleming, S. M., Daw, N. D., & Landy, M. S. (2017). Suboptimal Criterion Learning in Static and Dynamic Environments. *PLoS Computational Biology*, 13(1), 1–28. <https://doi.org/10.1371/journal.pcbi.1005304>
- Orban, 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. <https://doi.org/10.1073/pnas.0708424105>
- Owen, A. M., McMillan, K. M., Laird, A. R., & Bullmore, E. (2005). N-back working memory paradigm: A meta-analysis of normative functional neuroimaging studies. *Human Brain Mapping*, 25(1), 46–59.
- Paeye, C., Schütz, A. C., & Gegenfurtner, K. R. (2016). Visual reinforcement shapes eye movements in visual search. *Journal of Vision*, 16(10), 15. <https://doi.org/10.1167/16.10.15>
- Parkhurst, D., Law, K., & Niebur, E. (2002). Modeling the role of salience in the allocation of overt visual attention. *Vision Research*, 42(1), 107–123. Retrieved from <http://www.ncbi.nlm.nih.gov/pubmed/11804636>
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., ... Dubourg, V. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12(Oct), 2825–2830.
- Peterson, M., Kramer, A., Wang, R. F., Irwin, D. E., & McCarley, J. S. (2001). Visual Search Has Memory. *Psychological Science*, 12(4), 287–292. <https://doi.org/10.1089/ast.2015.1437>
- Peterson, M. S., & Kramer, A. F. (2001). Attentional guidance of the eyes by contextual information and abrupt onsets. *Perception and Psychophysics*, 63(7), 1239–1249. <https://doi.org/10.3758/BF03194537>

- Posner, M. I., Rafal, R. D., Choate, L. S., & Vaughan, J. (1985). Inhibition of return: Neural basis and function. *Cognitive Neuropsychology*, 2(3), 211–228.
- Rahnev, D., & Denison, R. N. (2018). Suboptimality in Perceptual Decision Making. *Behavioral and Brain Sciences*, 1–107. <https://doi.org/10.1017/S0140525X18000936>
- Ratcliff, R., & Rouder, J. N. (1998). Modeling Response Times for Two-Choice Decisions. *Psychological Science*, 9(5), 347–356. <https://doi.org/10.1111/1467-9280.00067>
- Rayner, K., & Raney, G. E. (1996). Eye movement control in reading and visual search: Effects of word frequency. *Psychonomic Bulletin & Review*, 3(2), 245–248. <https://doi.org/10.3758/BF03212426>
- Reber, A. S., Walkenfeld, F. F., & Hernstadt, R. (1991). Implicit and explicit learning: Individual differences and IQ. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17(5), 888.
- Renninger, L. W., & Coughlan, J. (2007). Where to look next ? Eye movements reduce local uncertainty, 7, 1–17. <https://doi.org/10.1167/7.3.6.Introduction>
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. *Classical Conditioning II Current Research and Theory*, 21(6), 64–99. <https://doi.org/10.1101/gr.110528.110>
- Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin & Review*, 16(2), 225–237. <https://doi.org/10.3758/PBR.16.2.225>
- Ryan, J. D., & Villate, C. (2009). Building visual representations: The binding of relative spatial relations across time. *Visual Cognition*, 17(1–2), 254–272. <https://doi.org/10.1080/13506280802336362>
- 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., Johnson, E. K., Aslin, R. N., & Newport, E. L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition*, 70(1), 27–52. [https://doi.org/10.1016/S0010-0277\(98\)00075-4](https://doi.org/10.1016/S0010-0277(98)00075-4)
- Sagi, D. (2011). Perceptual learning in vision research. *Vision Research*, 51(13), 1552–1566.
- Sanborn, A. N., & Chater, N. (2016). Bayesian Brains without Probabilities. *Trends in Cognitive Sciences*, 20(12), 883–893. <https://doi.org/10.1016/j.tics.2016.10.003>
- Santolin, C., & Saffran, J. R. (2018). Constraints on Statistical Learning Across Species. *Trends in Cognitive Sciences*, 22(1), 52–63. <https://doi.org/10.1016/j.tics.2017.10.003>
- Schrater, P., & Acuna, D. E. (2010). Structure Learning in Human Sequential. *PLoS Computational Biology*, 6(12). <https://doi.org/10.1371/Citation>
- Schutz, A. C., Trommershauser, J., & Gegenfurtner, K. R. (2012). Dynamic integration of information about salience and value for saccadic eye movements. *Proceedings of the National Academy of Sciences*, 109(19), 7547–7552. <https://doi.org/10.1073/pnas.1115638109>
- Settles, B. (2010). Active Learning Literature Survey.
- Shadlen, M. N. N., & Shohamy, D. (2016). Decision Making and Sequential Sampling from Memory. *Neuron*, 90(5), 927–939. <https://doi.org/10.1016/j.neuron.2016.04.036>
- Siegelman, N., Bogaerts, L., Christiansen, M. H., & Frost, R. (2017). Towards a theory of individual

- differences in statistical learning. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 372(1711), 20160059. <https://doi.org/10.1098/rstb.2016.0059>
- Smith, C. N., & Squire, L. R. (2008). Experience-Dependent Eye Movements Reflect Hippocampus-Dependent (Aware) Memory. *Journal of Neuroscience*, 28(48), 12825–12833. <https://doi.org/10.1523/JNEUROSCI.4542-08.2008>
- Spalek, T. M., & Hammad, S. (2005). The left-to-right bias in inhibition of return is due to the direction of reading. *Psychological Science*, 16(1), 15–18.
- Spelke, E. S. (1985). Preferential-looking methods as tools for the study of cognition in infancy. *Measurement of Audition and Vision in the First Year of Postnatal Life: A Methodological Overview*, 323–363. <https://doi.org/10.1336/0893911305>
- St John-Saaltink, E., Kok, P., Lau, H. C., & de Lange, F. P. (2016). Serial Dependence in Perceptual Decisions Is Reflected in Activity Patterns in Primary Visual Cortex. *The Journal of Neuroscience : The Official Journal of the Society for Neuroscience*, 36(23), 6186–6192. <https://doi.org/10.1523/JNEUROSCI.4390-15.2016>
- Stanislaw, H., & Todorov, N. (1999). Calculation of signal detection theory measures. *Behavior Research Methods, Instruments, & Computers*, 31(1), 137–149. <https://doi.org/10.3758/BF03207704>
- Steyvers, M., Lee, M. D., & Wagenmakers, E. J. (2009). A Bayesian analysis of human decision-making on bandit problems. *Journal of Mathematical Psychology*, 53(3), 168–179. <https://doi.org/10.1016/j.jmp.2008.11.002>
- Summerfield, C., Behrens, T. E., & Koechlin, E. (2011). Perceptual classification in a rapidly changing environment. *Neuron*, 71(4), 725–736. <https://doi.org/10.1016/j.neuron.2011.06.022>
- Summerfield, C., & Koechlin, E. (2010). Economic Value Biases Uncertain Perceptual Choices in the Parietal and Prefrontal Cortices. *Frontiers in Human Neuroscience*, 4(November), 1–12. <https://doi.org/10.3389/fnhum.2010.00208>
- Summerfield, C., Trittschuh, E. H., Monti, J. M., Mesulam, M. M., & Egner, T. (2008). Neural repetition suppression reflects fulfilled perceptual expectations. *Nature Neuroscience*, 11(9), 1004–1006. <https://doi.org/10.1038/nn.2163>
- Summerfield, J. J., Lepsien, J., Gitelman, D. R., Mesulam, M. M., & Nobre, A. C. (2006). Orienting attention based on long-term memory experience. *Neuron*, 49(6), 905–916.
- Tatler, B. W. (2007). The central fixation bias in scene viewing : Selecting an optimal viewing position independently of motor biases and image feature distributions, 7, 1–17. <https://doi.org/10.1167/7.14.4.Introduction>
- Taubert, J., Alais, D., & Burr, D. (2016). Different coding strategies for the perception of stable and changeable facial attributes. *Scientific Reports*, 6, 32239. <https://doi.org/10.1038/srep32239>
- Téglás, E., & Bonatti, L. L. (2016). Infants anticipate probabilistic but not deterministic outcomes. *Cognition*, 157, 227–236. <https://doi.org/10.1016/j.cognition.2016.09.003>
- Thompson, P., & Burr, D. (2009). Visual aftereffects. *Current Biology*, 19(1), 11–14.
- Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 267–288.
- Tipper, S. P. (1985). The negative priming effect: Inhibitory priming by ignored objects. *The Quarterly Journal of Experimental Psychology Section A*, 37(4), 571–590.

- Toro, J. M., Sinnett, S., & Soto-Faraco, S. (2005). Speech segmentation by statistical learning depends on attention. *Cognition*, 97(2), B25–B34.
- Treisman, A. (1992). Perceiving and re-perceiving objects. *American Psychologist*, 47(7), 862.
- Trueswell, J. C., Medina, T. N., Hafri, A., & Gleitman, L. R. (2013). Propose but verify: Fast mapping meets cross-situational word learning. *Cognitive Psychology*, 66(1), 126–156. <https://doi.org/10.1016/j.cogpsych.2012.10.001>
- Tulving, E., & Schacter, D. L. (1990). Priming and Human Memory Systems. *Science*, 247(4940), 302–306. <https://doi.org/10.1126/science.2296719>
- 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. <https://doi.org/10.1037/0096-3445.134.4.552>
- 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. <https://doi.org/10.1523/JNEUROSCI.0858-10.2010>
- Tversky, A., & Kahneman, D. (1971). Belief in the law of small numbers. *Psychological Bulletin*, 76(2), 105.
- Tversky, A., & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science*, 185(4157), 1124–1131. <https://doi.org/10.1017/CBO9780511809477.002>
- Urai, A. E., Braun, A., & Donner, T. H. (2017). Uncertainty and Alters Serial Choice Bias. *Nature Communications*, 8, 1–11. <https://doi.org/10.1038/ncomms14637>
- Urai, A. E., Gee, J. W. De, & Donner, T. H. (2018). Choice history biases subsequent evidence accumulation. *BioRxiv*, 1–24. <https://doi.org/https://doi.org/10.1101/251595>
- Urgolites, Z. J., Smith, C. N., & Squire, L. R. (2018). Eye movements support the link between conscious memory and medial temporal lobe function. *Proceedings of the National Academy of Sciences*, 115(29), 7599–7604. <https://doi.org/10.1073/pnas.1803791115>
- van der Heijden, a H. (1989). Probability matching in visual selective attention. *Canadian Journal of Psychology*, 43(1), 45–52. <https://doi.org/10.1037/h0084208>
- van Kerkoerle, T., Self, M. W., & Roelfsema, P. R. (2017). Layer-specificity in the effects of attention and working memory on activity in primary visual cortex. *Nature Communications*, 8, 13804. <https://doi.org/10.1038/ncomms13804>
- Vinken, K., & Vogels, R. (2017). Adaptation can explain evidence for encoding of probabilistic information in macaque inferior temporal cortex. *Current Biology*, 27(22), R1210–R1212. <https://doi.org/10.1016/j.cub.2017.09.018>
- Vinken, K., Vogels, R., & Beeck, H. Op De. (2017). Recent Visual Experience Shapes Visual Processing in Rats through Stimulus-Specific Adaptation and Response Enhancement. *Current Biology*, 27(6), 1–6. <https://doi.org/10.1016/j.cub.2017.02.024>
- Võ, M. L. H., & Wolfe, J. M. (2012). When does repeated search in scenes involve memory? Looking at versus looking for objects in scenes. *Journal of Experimental Psychology: Human Perception and Performance*, 38(1), 23–41. <https://doi.org/10.1037/a0024147>
- Võ, M. L. H., & Wolfe, J. M. (2013). The interplay of episodic and semantic memory in guiding repeated search in scenes. *Cognition*, 126(2), 198–212. <https://doi.org/10.1016/j.cognition.2012.09.017>

- Vul, E., Goodman, N., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and Done? Optimal Decisions From Very Few Samples. *Cognitive Science*. <https://doi.org/10.1111/cogs.12101>
- Wagenmakers, E. J., Love, J., Marsman, M., Jamil, T., Ly, A., Verhagen, J., ... Morey, R. D. (2018). Bayesian inference for psychology. Part II: Example applications with JASP. *Psychonomic Bulletin and Review*, 25(1), 58–76. <https://doi.org/10.3758/s13423-017-1323-7>
- Walker-Smith, G. J., Gale, A. G., & Findlay, J. M. (2013). Eye movement strategies involved in face perception. *Perception*, 42(11), 1120–1133.
- Walther, C., Schweinberger, S. R., Kaiser, D., & Kovács, G. (2013). Neural correlates of priming and adaptation in familiar face perception. *Cortex*, 49(7), 1963–1977. <https://doi.org/10.1016/j.cortex.2012.08.012>
- Walthew, C., & Gilchrist, I. D. (2006). Target location probability effects in visual search: An effect of sequential dependencies. *Journal of Experimental Psychology: Human Perception and Performance*, 32(5), 1294–1301. <https://doi.org/10.1037/0096-1523.32.5.1294>
- Wei, X.-X., & Stocker, A. A. (2017). Lawful relation between perceptual bias and discriminability. *Proceedings of the National Academy of Sciences*, 114(38), 201619153. <https://doi.org/10.1073/pnas.1619153114>
- Weiss, Y., Simoncelli, E. P., & Adelson, E. H. (2002). Motion illusions as optimal percepts. *Nature Neuroscience*, 5(6), 598.
- Wilder, M. H., Jones, M., Ahmed, A. a, Curran, T., & Mozer, M. C. (2013). The persistent impact of incidental experience. *Psychonomic Bulletin & Review*, 20(6), 1221–1231. <https://doi.org/10.3758/s13423-013-0406-3>
- Willingham, D. B., & Goedert-Eschmann, K. (1999). The relation between implicit and explicit learning: Evidence for parallel development. *Psychological Science*, 10(6), 531–534.
- Winterbottom, M. (2012). Active Learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 6(1), 1–2. https://doi.org/10.1300/J172v06n01_05
- Wolfe, J. M. (1984). Short test flashes produce large tilt aftereffects. *Vision Research*, 24(12), 1959–1964. [https://doi.org/10.1016/0042-6989\(84\)90030-0](https://doi.org/10.1016/0042-6989(84)90030-0)
- Wolfe, J. M. (1998). What can 1 million trials tell us about visual search? *Psychological Science*, 9(1), 33–39. <https://doi.org/10.1111/1467-9280.00006>
- Wozny, D. R., Beierholm, U. R., & Shams, L. (2010). Probability matching as a computational strategy used in perception. *PLoS Computational Biology*, 6(8). <https://doi.org/10.1371/journal.pcbi.1000871>
- Wu, S.-W., Delgado, M. R., & Maloney, L. T. (2009). Economic decision-making compared with an equivalent motor task. *Proceedings of the National Academy of Sciences of the United States of America*, 106(15), 6088–6093. <https://doi.org/10.1073/pnas.0900102106>
- Xiao, L.-Q., Zhang, J.-Y., Wang, R., Klein, S. A., Levi, D. M., & Yu, C. (2008). Complete transfer of perceptual learning across retinal locations enabled by double training. *Current Biology*, 18(24), 1922–1926.
- Yamamoto, N., & Philbeck, J. W. (2013). Peripheral vision benefits spatial learning by guiding eye movements. *Memory & Cognition*, 41(1), 109–121. <https://doi.org/10.3758/s13421-012-0240-2>
- Yang, S. C. H., Lengyel, M., & Wolpert, D. M. (2016). Active sensing in the categorization of visual patterns. *ELife*, 5(FEBRUARY2016), 1–22. <https://doi.org/10.7554/eLife.12215>

- Yang, S. C. H., Wolpert, D. M., & Lengyel, M. (2016). Theoretical perspectives on active sensing. *Current Opinion in Behavioral Sciences*, 11, 100–108. <https://doi.org/10.1016/j.cobeha.2016.06.009>
- Yonelinas, A. (2001). Components of episodic memory: the contribution of recollection and familiarity. ... *Transactions of the Royal Society of* Retrieved from <http://rstb.royalsocietypublishing.org/content/356/1413/1363.short>
- Yu, A. J., & Cohen, J. D. (2008). Sequential effects : Superstition or rational behavior ? *NIPS*, 1–8.
- Yu, C., & Smith, L. B. (2011). What you learn is what you see: Using eye movements to study infant cross-situational word learning. *Developmental Science*, 14(2), 165–180. <https://doi.org/10.1111/j.1467-7687.2010.00958.x>
- Yu, R. Q., & Zhao, J. (2015). The persistence of the attentional bias to regularities in a changing environment. *Attention, Perception, and Psychophysics*, 77(7), 2217–2228. <https://doi.org/10.3758/s13414-015-0930-5>
- Yuille, A., & Kersten, D. (2006). Vision as Bayesian inference: analysis by synthesis? *Trends in Cognitive Sciences*, 10(7), 301–308.
- Zhao, C., Seriès, P., Hancock, P. J. B., & Bednar, J. A. (2011). Similar neural adaptation mechanisms underlying face gender and tilt aftereffects. *Vision Research*, 51(18), 2021–2030.
- Zhao, J., Al-Aidroos, N., & Turk-Browne, N. B. (2013). Attention is spontaneously biased toward regularities. *Psychological Science*, 24(5), 667–677. <https://doi.org/10.1177/0956797612460407>