UC San Diego
UC San Diego Electronic Theses and Dissertations

Title
Adaptations of temporal dynamics

Permalink
https://escholarship.org/uc/item/06w2r9j6

Author
Rieth, Cory Alan

Publication Date
2012

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, SAN DIEGO

Adaptations of temporal dynamics: Faces, places, and words

A dissertation submitted in partial satisfaction of the requirements for the degree Doctor of Philosophy

in

Psychology and Cognitive Science

by

Cory Alan Rieth

Committee in charge:

      Professor David E. Huber, Chair
      Professor Garrison W. Cottrell
      Professor Karen R. Dobkins
      Professor Donald MacLeod
      Professor Angela J. Yu

2012

The Dissertation of Cory Alan Rieth is approved, and it is acceptable in quality and form

for publication on microfilm and electronically:

_____

_____

_____

_____

_____
Chair

University of California, San Diego

2012

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGEMENTS

Mendeley. These valuable pieces of free software contributed greatly to this work.

My greatest thanks goes to all of the amazing teachers in my life, both inside and outside of school. I wouldn't be writing this dissertation without all of you. Most importantly thanks to my advisor, Dr. David Huber, who has been the ultimate mentor. His time, support, and assistance have been essential, and he has been generous with all of them. I would not be writing this without him. Thanks also to Dr. Garrison Cottrell, for letting me hang around GURU meetings all these years, and encouraging both hard work and appreciation for the bigger picture. The other members of my committee, Dr. Karen Dobkins, Dr. Donald MacLeod, and Dr. Angela Yu also have my appreciation for the guidance they have provided throughout graduate school in meetings, seminars, and journal clubs. Additionally, thanks to all of my collaborators for their patience in working with someone who habitually takes on more than they should.

The love and support of family and friends was an essential ingredient of this dissertation. You have kept me sane throughout the years, and allowed for a maintained illusion of normalcy among all the work. Naming everyone that deserves it would take far too much space, but I have to express my gratitude to my mother. Without her pushing me to succeed, I wouldn't have come this far. A special thanks to Sarah, who has my utmost appreciation for all her help, patience and belief in me. The last few weeks would have been that much rougher without her around, not to mention all the time before that. Finally, if you are reading this, you deserve recognition as well... thanks!

VITA

2000   Bachelor of Arts, Boston University

2006   Master of Arts, University of California, San Diego

2012   Doctor of Philosophy, University of California, San Diego

PUBLICATIONS

Rieth, C. A., Lee, K., Liu, J., Tian, J., & Huber, D. E. (2011). Faces in the mist: illusory detection of faces and letters. i-Perception 2(5) 458-476.

Liu, J., Li, J., Rieth, C. A., Huber, D. E., Tian, J., & Lee, K,. (2011). A dynamic causal modeling analysis of the effective connectivities underlying top-down letter processing. Neuropsychologia. 49 (5) 1177-1186.

Rieth, C. A., Cai, D. J., McDevitt, E., & Mednick, S. (2010). The role of sleep and practice in motor memory consolidation. Behavioural Brain Research. 214 (2). 470-474.

Li, J., Liu, J., Liang, J., Zhang, H., Zhao, J., Rieth, C. A., Huber, D. E., Shi, G., Tian, J., & Lee, K. (2010). Effective connectivities of cortical regions for top-down face processing: A dynamic causal modeling study. Brain Research, 1340, 40-51.

Huber, D. E., Tomlinson, T. D., Rieth, C. A., & Davelaar, E. J. (2010). Reply to Ba̋uml and Hanslmayr: Adding or subtracting memories? The neural correlates of learned interference versus memory inhibition. Proceedings of the National Academy of Science, 107(2), E4.

Liu, J., Li, J., Liang, J., Tian, J., Zhang, H., Rieth, C. A., Huber, D. E., & Lee, K. (2010). Neural correlates of top-down letter processing. Neuropsychologia, 48 (2), 636-641.

Rieth, C. A., & Huber, D. E. (2010). Priming and habituation for faces: Individual differences and inversion effects. Journal of Experimental Psychology: Human Perception and Performance, 36 (3), 596-618.

Tomlinson, T. D., Huber, D. E., Rieth, C. A., & Davelaar, E. J. (2009). An interference account of cue-independent forgetting in the no-think paradigm. Proceedings of the National Academy of Science, 106, 15588-15593.

Li, J., Liu, J., Liang, J., Zhang, H., Zhao, J., Huber, D. E., Rieth, C. A., Lee, K., Tian, J., & Shi, G. (2009). A distributed neural system for top-down face processing. Neuroscience Letters, 451 (1), 6-10.

Rickard, T. C., Cai, D. J., Rieth, C. A., Jones, J., & Ard, C. (2008). Sleep does not enhance motor sequence learning. Journal of Experimental Psychology: Learning, Memory, and Cognition, 34 (4), 834-842.

Zhang, H., Liu, J., Huber, D. E., Rieth, C. A., Stiles, J., Tian, J., & Lee, K. (2008). Detecting faces in pure noise images: An fMRI study on top-down perception. Neuroreport, 19, 229-233.

Rieth, C. A., & Huber, D. E. (2005). Using a neural network model with synaptic depression to assess the dynamics of feature-based versus configural processing in face identification. Proceedings of the 27th Annual Conference of the Cognitive Science Society. Hillsdale, NJ: Erlbaum Associates.

ABSTRACT OF THE DISSERTATION

Adaptations of temporal dynamics: Faces, places, and words

by

Cory Alan Rieth

Doctor of Philosophy in Psychology and Cognitive Science

University of California, San Diego, 2012

Professor David E. Huber, Chair

The human mind rapidly parses high-dimensional dynamic input into useful
representations. Identifying the source of relevant sensory information in time is critical
to the formation of informative perceptions. One source of information to aid perceptual
processing is the temporal regularity of sensory input. Knowledge of what regularities are
utilized and how they are utilized is critical to understanding how perception functions in
a dynamic world. The objective of this dissertation is to explore the adaptations of
perception to temporal regularities in the environment. Chapters 1 and 2 use short term
priming paradigms to examine the temporal segregation of perceptual information.
Chapter 1 shows that inverted versus upright faces have different time courses of
priming, suggesting that familiarity plays a role in temporal segregation. Furthermore, for
faces, there are individual differences in priming dynamics related to perceptual

identification ability. Chapter 2 examines the influence of occurrence and co-occurrence statistics in the same task. Several incremental reductions of stimulus frequency at different levels of representation all slow priming compared to higher occurring word stimuli. Word associations as a measure of co-occurrence, produce waxing and waning priming effects, with a small effect of the order of the associative relationship. Priming based on direct measures of linguistic co-occurrence produces priming dynamics that are similar to associative priming. Words which are expected to repeat provide a small increase in priming compared to priming by repeating a word. These priming effects are interpreted as the result of habituation. The findings in Chapters 1 and 2 generalize the habituation account of priming effects to faces, demonstrate its dependence on the occurrence rates of stimuli, and show that temporal co-occurrence regularities for particular items only result in small changes to the basic dynamics. Chapter 3 uses a spatial cueing paradigm to test if we adapt to diagnostic spatial-temporal relationships. Participants implicitly learn and generalize diagnostic cue-target relationships with long delays between cues and target. Diagnostic relationships are also learned for short cue-target delays when no shift of expectation is required. As a whole, these experiments find that temporal regularities pertaining to familiarity and general occurrence relationships are strongly reflected in the dynamics of perception, and that directional associations or expectations between individual stimuli deviate only slightly from general dynamics.

INTRODUCTION

Perception extracts relevant relationships from sensory input to inform decision making and actions. An important and often overlooked aspect of perception is that it operates in a dynamic world. Perception must discriminate objects not only in space, but also in time. Through both evolution and development, our minds have adapted to process sensory information effortlessly and reasonably accurately. However, violations of expected temporal relationships produce consistent behavioral consequences. Using inappropriate or insufficient information to make perceptual inferences results in measureable biases. For example, our perceptual systems may not be able to resolve the identity of a quickly flashed (and subsequently masked) word. While the entire word was presented to our retina, we remain unable to identify it definitively, having been both experientially and evolutionarily accustomed to additional time to accomplish this task. This dissertation presents a series of behavioral experiments investigating the nature of adaptations of perception to the temporal structure of the world.

Considering the statistics of the world in psychological studies has a long history (Estes, 1955). In the form of rational analysis (Anderson, 1990) it has been particularly useful for understanding memory. Other frameworks, for example probabilistic topic models (Griffiths, 2004; Griffiths & Steyvers, 2007; Landauer, 1997; Steyvers & Griffiths, 2006; Steyvers, Griffiths, & Dennis, 2006), similarly use environmental statistics to understand cognition (for reviews of the impact and implications of rational analysis see Chater & Oaksford, 1999; Steyvers & Griffiths, 2008). These approaches generally operate on pre-extracted representations, and do not consider the process of learning and activating representations. Similar approaches have also been fruitful for

understanding low-level visual representations. For example, computational investigations have found that Gabor- like components emerge from analysis of natural image statistics (Bell & Sejnowski, 1997; Olshausen & Field, 1996). These components are very similar to the Gabor-like receptive fields of V1 discovered through direct neuroscientific investigations (Hubel & Wiesel, 1968). Thus, like higher memory and cognitive processes, the properties of early visual cortex can be thought of as an adaptation to the statistical properties of its input. The present research concerns what happens in between these two processes. A strict rational approach to this problem is currently an impossible task. It would require an accurate model of all transformations of sensory input. Therefore, instead of measuring the properties of dynamic information in the world and building a rational model characterizing this process, this dissertation studies the role of environmental statistics in the processing of information using a combination of experimentation and computational modeling.

There are two extreme viewpoints concerning the role of environmental statistics in perception. One posits that the mind is not at all adapted to different information and that everything is processed in a general purpose manner. The other is that the mind is acutely adapted to every possible individual stimulus. Clearly neither of these extremes is optimal; one ignores potentially helpful relationships, while the latter makes generalization of knowledge to new situations untenable. The optimal balance lies somewhere between these extremes. Importantly, the structure and algorithms of the mind need not be rational, even after consideration of non-informational constraints (e.g. energy consumption). The purpose of the following studies is to explore the ways

perception is adapted to temporal regularities, at what levels of processing it is adapted, and how flexible these adaptations are.

Chapters 1 and 2 focus on the formation of coherent percepts from dynamic input. Sensory input must be attributed to either static or changing attributes of the world. This attribution can be measured through priming experiments, where a prime precedes a briefly flashed target. In this sequence of events, perceptions resulting from the prime need to be separated from those in response to the target. Varying prime identity and duration offer insights in to this process. Immediate priming experiments have previously showed that word priming shows properties of habituation (Huber, 2008; Huber & O'Reilly, 2003; Huber, Shiffrin, Lyle, & Ruys, 2001). Short primes (17 ms) bias the identification of the target towards attributes of prime item, while actively processed or longer primes (2,000 ms) bias target identification away from attributes of the prime (Huber, 2008; Huber et al., 2001). This interaction with prime duration is reflective of mental processing speed. If the world changes too fast, we attribute novel features to past items, and the world blends together. Chapters 1 and 2 investigate how this habituation process is adapted to temporal regularities. It may be specific to linguistic stimuli, or more domain general. Furthermore, it may or may not be sensitive to different base rates of occurrence or co-occurrence in the world.

Chapter 1 uses face images to test the critical question of whether these priming effects are general properties of perception rather than specific to linguistic stimuli. Furthermore, Chapter 1 explores individual differences in the task which were not found with words, and also the role of information content through comparison of upright with inverted face images. Chapter 2 tests the roles of occurrence and co-occurrence statistics

in perceptual segregation. Even if the segregation process of the mind is attuned to differences in frequency, there are two contrasting rational predictions. One hypothesis is that habituation reflects the completion of information extraction from a stimulus, with discounting occurring once available information is exhausted. According to this hypothesis, stimuli with more information content invoke negative priming only at longer priming durations compared to stimuli conveying little useful information. This predicts that non-words, containing no semantic content, elicit negative priming at shorter prime durations than words. However, an alternative hypothesis predicts that priming dynamics reflect the activation strength or speed of different stimuli. More frequently occurring stimuli should be processed more quickly, and show negative priming for shorter prime durations than unfamiliar stimuli. Furthermore, if general frequency and familiarity characteristics influence the segregation of items, than co-occurrence rates between individual items may also be reflected in priming dynamics. In Chapter 2, associative relationships are used as a measure of co-occurrence to test if the direction of association is reflected in priming dynamics. Chapter 2 also investigates the role of co-occurrence between specific items by directly testing primes which are likely to repeat. The overall priming dynamics may simply reflect expectations against repeated stimuli, resulting in negative priming. If particular items are likely to repeat, negative priming may be lessened or absent for these specific stimuli.

Following the small effects of adaptation to relationships between individual stimuli, Chapter 3 tested if general temporal relationships are learned using a spatial cueing paradigm. Participants were implicitly trained with diagnostic spatial-temporal relationships between the cue and target, and their behavior measured in later test blocks

with no diagnostic relationships. At a conceptual level, typical spatial cueing results have

similar attributes to the priming task in Chapters 1 and 2. They both involve a facilitation

of similarity at short intervals, but a determent of similarity at long intervals. When a cue

is presented just before a target in the same location, participants' reaction time to detect

targets is facilitated relative to when a cue is presented just before a target at a different

location. In contrast, after about 300 ms, reaction time is slower for targets appearing in

the same location as a cue, compared to targets appearing at an alternative location. This

latter effect is commonly referred to as inhibition of return (Posner & Cohen, 1984;

Posner, Rafal, Choate, & Vaughan, 1985). Inhibition of return has in fact been

conceptualized as a habituation of the orienting response (Dukewich, 2009), similar to the

habituation of item representations considered in Chapters 1 and 2. In Chapter 3, several

experiments test to see if reaction time adapts to diagnostic spatial-temporal relationships

between the cue and target. For example, after training when the target constantly appears

at either an opposite location as the cue after a short delay, or at the same location after a

long delay, will inhibition of return disappear? If so, the typical spatial cueing effect

might be an adaptation to relationships between salient items in the world and perceptual

targets. Typically, salient items may predict the occurrence of importance information,

but after determining there is no relevant information at a location, useful information

may be more likely found elsewhere.

Together, these studies use short term priming and spatial cueing investigate how

perception is informed by different types of temporal regularities. Both priming and

spatial cueing show similar patterns of results, with positive effects over short temporal

asynchronies, and negative effects after a delay. Effects of the type and familiarity of

stimuli are investigated, as well as expectation and association relationships between

individual stimuli and general temporal expectations.

References

Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Earlbaum.

Bell, A. J., & Sejnowski, T. J. (1997). The "independent components" of natural scenes are edge filters. *Vision Research*, *37*(23), 3327-3338.

Chater, N., & Oaksford, M. (1999). Ten years of the rational analysis of cognition. *Trends in Cognitive Sciences*, *3*(2), 57-65.

Dukewich, K. R. (2009). Reconceptualizing inhibition of return as habituation of the orienting response. *Psychonomic Bulletin & Review*, *16*(2), 238-251. doi:10.3758/PBR.16.2.238

Estes, W. K. (1955). Statistical theory of spontaneous recovery and regression. *Psychological Review*, *62*(3), 145-154.

Griffiths, T. L. (2004). Finding scientific topics. *Proceedings of the National Academy of Sciences of the United States of America*, *101*(Suppl 1), 5228-5235.

Griffiths, T. L., & Steyvers, M. (2007). Topics in semantic representation. *Psychological review*, *114*(2), 211-244.

Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *The Journal of Physiology*, *195*(1), 215-243.

Huber, D. E. (2008). Immediate priming and cognitive aftereffects. *Journal of Experimental Psychology: General*, *137*(2), 324-347. doi:10.1037/0096-3445.137.2.324

Huber, D. E., & O'Reilly, R. C. (2003). Persistence and accommodation in short-term priming and other perceptual paradigms: temporal segregation through synaptic depression. *Cognitive Science*, *27*(3), 403-430. doi:10.1016/S0364-0213(03)00012-0

Huber, D. E., Shiffrin, R. M., Lyle, K. B., & Ruys, K. I. (2001). Perception and preference in short-term word priming. *Psychological Review*, *108*(1), 149-182. doi:10.1037//0033-295X.108.1.I49

Landauer, T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, *17*, 660-665.

Olshausen, B. A., & Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, *381*(6583), 607-609.

Posner, M. I., & Cohen, Y. (1984). Components of visual orienting. *Attention and performance X: Control of language processes*, *32*, 531-556.

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.

Steyvers, M., & Griffiths, T. L. (2006). Probabilistic topic models. In T. Landauer, D. McNamara, S. Dennis, & W. Kintsch (Eds.), *Handbook of latent semantic analysis*. (pp. 424-440). Lawrence Erlbaum, Mahwah, NJ.

Steyvers, M., & Griffiths, T. L. (2008). Rational analysis as a link between human memory and information retrieval. In N. Chater & M. Oaksford (Eds.), *The probabilistic mind: Prospects for a Bayesian cognitive science* (pp. 329-350). Oxford: Oxford University Press.

Steyvers, M., Griffiths, T. L., & Dennis, Simon. (2006). Probabilistic inference in human semantic memory. *Trends in Cognitive Sciences*, *10*(7), 327-334. doi:10.1016/j.tics.2006.05.005

CHAPTER 1: Priming and habituation for faces: Individual differences and
inversion effects

Searching for a face in a crowd is a seemingly easy task. Yet, to complete this
task, we must identify the subtle differences between different faces while generalizing
face recognition across variations in image size, orientation, hairstyle, lighting, and other
more salient perceptual dimensions. Because a crowd is in constant motion, and because
we generally wish to find a desired face quickly, not much time is afforded to consider
each face. Nonetheless, some minimum duration is needed to integrate sufficient
information for identification. The constant integration of perceptual input over time
implies the possibility that we may incorrectly blend together one face and the next as our
eyes scan across positions, or as the crowd moves. However, we do not suffer greatly
from such source confusion between successive faces. To explain this and other results
involving the effects of successively presented stimuli, we suggest that the perceptual
system includes a discounting mechanism that appropriately reduces the response to a
previously presented face, and this reduction serves to offset the effect of source
confusion.

Huber & O'Reilly (2003) proposed that neural habituation is the basic mechanism
behind temporal discounting, which automatically parses the stream of perceptual events.
This theory was developed to explain priming effects with words, and the model correctly
predicted a variety of lexical effects as well as electrophysiological recordings when
viewing words. Because habituation is a general mechanism, similar effects should exist
with many other tasks and stimuli. To explore the generality of this account, we
examined face identification because it is a type of perceptual expertise that is thought to

9

differ from word reading in a variety of ways. Beyond generalization to a new stimulus class, face identification allowed us to test a key prediction of this account. Because identification is the driving force behind habituation, situations that allow faster identification should also produce faster habituation and thus faster prime discounting. In the reported experiments we tested this prediction in two ways. First, we compared individuals who could identify a target face with short exposure durations to those who needed a longer duration. Second, we compared upright and inverted faces because inverted faces are more difficult to recognize. For both of these tests, we examined immediate face repetition priming to assess temporal discounting.

To motivate Experiment 1.1, which examined individual differences in face habituation as indicated by priming, we next 1) present the a theory of temporal segregation through habituation; 2) review previous work with word priming based on average performance; 3) present a re-analysis of individual differences with word priming, which failed to confirm the predicted individual differences; and 4) review previous word on habituation with faces.

Temporal segregation through habituation

The integration of information over time while avoiding temporal source confusion is a general problem in perception and these effects are likely to exist beyond the domains of face perception and word priming. Our perceptual system effortlessly separates the constant stream of input into discrete objects and events. Yet this temporal segregation is not trivial, requiring neither too little nor too much integration in calculating both where and *when* one object begins and another ends. Temporal integration allows the perceptual system to appropriately combine different features

contained within a single object, such as when an object is partially blocked by a moving occluder. However, too much integration can inappropriately blend perception of a previous object with a subsequent object (e.g. combining Ernie's eyes with Bert's nose). This problem of temporal segregation cannot be solved with fixed discrete timing intervals considering that each stimulus in each situation may necessitate a different time scale for proper temporal segregation. Instead, we suggest that there is a constant interplay between the process of integrating the current percept and the process of discounting percepts that were previously identified.

Mechanisms that reduce inappropriate perceptual blending while maintaining appropriate perceptual binding have been proposed both at the cognitive level, such as in Treisman's feature integration theory (Treisman & Gelade, 1980) or Biedermann's GEON theory (Hummel & Biederman, 1992), as well as at the neural level, such as with Singer's synchrony theory (Singer, 1999) or Grossberg's adaptive resonance theory (Carpenter & Grossberg, 2003). To bridge these levels of description, we have addressed the problem of temporal segregation by proposing interrelated models in terms of optimal Bayesian 'explaining away' (Huber, Shiffrin, Lyle, & Ruys, 2001; Huber, 2008a) as well as with a specific neural account in the form of habituation through synaptic depression (Huber & O'Reilly, 2003). In line with David Marr's levels of explanation (Marr, 1982), we seek to find a mapping between the computational "why" and the implementational "how" of perceptual segregation through discounting.

For word priming, perceptual discounting through neural habituation has been examined both with behavioral measures as well as with neural measures. Behaviorally, repetition word priming was found to produce a gradual transition from positive priming

(more accurate identification of primed targets) to negative priming[1] (less accurate

identification of primed targets) as a function of increasing prime duration (Huber,

2008b, Huber, Shiffrin, Lyle, Quach, 2002; Weidemann, Huber, & Shiffrin, 2005;

Weidemann, Huber & Shiffrin, 2008;). In terms of neural behavior, repetition priming

produced Event-Related Potential (ERP) and evoked Magnetoencephalography (MEG)

responses in early perceptual processing that were modulated by increasing prime

duration (Huber, Tian, Curran, O'Reilly, & Woroch, 2008). Furthermore, the time course

of these neural effects was predicted by the habituation account as previously fit to

behavioral data.

In keeping with the proposal that dynamic discounting is a general mechanism of

perception, habituation through synaptic depression (Abbott, Varela, Sen, & Nelson,

1997; Tsodyks, & Markram, 1997), is found in the majority of pyramidal cells in many

cortical areas (Thomson & West, 1993). Synaptic depression refers to the finding that

post-synaptic depolarization (i.e., the message delivered to the receiving cell) is rapidly

weakened as a function of recent presynaptic action potentials (i.e., the number of

messages recently sent), with this loss of responsiveness lasting anywhere from hundreds

of milliseconds to several seconds. This finding is explained by hypothesizing the

existence of resources that enable effective signaling across the synapse (e.g., presynaptic

neurotransmitter availability, although many other mechanisms produce similar effects),

with these resources becoming depleted as a function of recent activity. Therefore,

synaptic depression is a natural candidate for automatic temporal segregation because it is

the ongoing activation that drives depletion of synaptic resources. Thus, habituation does

not occur until after activation is sufficient both in magnitude and duration (i.e.,
habituation follows on the heels of identification).

This hypothesis predicts a specific relationship between the speed of identification
and the speed of habituation. For instance, individuals who identify a class of object more
quickly (i.e., individuals who have greater perceptual expertise) should produce priming
data indicative of greater habituation. As reported below in a re-analysis of Experiment 1
from Huber (2008b), this prediction was not confirmed for threshold identification of
words. Before presenting this re-analysis, we review results using this particular priming
paradigm with words.

Habituation effects with immediate word priming

Immediate word priming has been used for more than 30 years to measure various
aspects of language (e.g., Meyer & Schvaneveldt, 1971). However, in our work, we use
word priming to investigate the basic mechanisms behind perceptual identification and
lexical retrieval—mechanisms that may well generalize to other stimuli and other tasks.
Priming is traditionally measured with speed of lexical decision, speed of naming, or
accuracy in the form of threshold identification. For threshold identification, the task is
usually to name the briefly flashed word (e.g., Meyer, Schvanevedlt, & Ruddy, 1974),
although forced-choice testing provides a variant of threshold identification that allows
separate measurement of the costs and benefits of priming (e.g., Ratcliff & McKoon,
1997) by comparing conditions that prime the correct answer (the target primed
condition) versus conditions that prime the incorrect answer (the foil primed condition).
Huber et al. (2001) adopted this paradigm for the study of immediate priming effects as a
function of the duration and degree of processing for primes (see Figure 1.1 for the

equivalent paradigm using faces). With this paradigm, short duration primes produced

benefits for the target primed condition and deficits for the foil primed condition (i.e., a

preference to choose the repeated choice alternative, or 'positive priming'). In contrast,

long duration primes produced the opposite pattern for repetition or orthographic

priming, with deficits for the target primed condition and benefits for the foil primed

condition (i.e., a preference against the repeated choice alternative, or 'negative

priming').

Prime
Identical to either
target or foil
17-2000 ms

OR

Target
Ex 1,3 and 4 duration set to threshold
Ex 2 experimentally manipulated
33-100 ms

Pattern Mask
500 ms – target face duration

**Forced Choice
Recognition**
Between target and foil faces
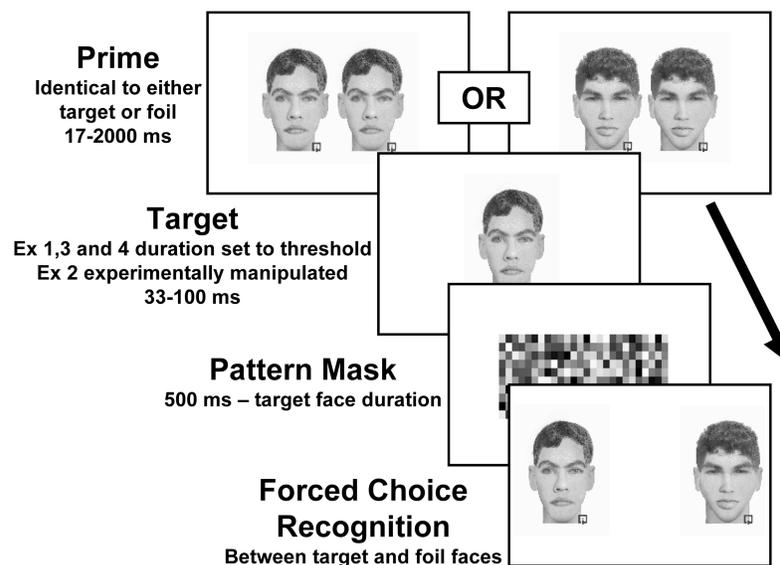
Figure 1.1 Presentation sequence used in Chapter 1. Side by side replicated images of either the foil or the target face appeared as the prime display, followed by a single image of the target face. All experiments manipulated prime duration, using 17, 50, 150, 400, and 2,000 ms (Experiments 1.1, 1.3, and 1.4) or 17, 150, and 2,000 ms (Experiment 1.2). Target duration in Experiments 1.1, 1.3 and 1.4 was set for each participant to a duration that produced 75% accuracy. In Experiment 1.2, 33, 50, and 100 ms target durations were tested for all participants. Target presentation was followed by a checkerboard pattern mask and then two choice faces (the target and foil) that remained on the screen until a response was made. Choices were presented at the top and bottom of the screen in Experiments 1.1, 1.3, and 1.4 or to the sides of the screen as pictured here in Experiment 1.2. All experiments discouraged strategic use of the primes by informing participants that the prime was just as likely to indicate the incorrect answer as the correct answer and by providing trial by trial accuracy feedback.

Huber et al. (2001) explained both the positive and negative priming results with a Bayesian model based upon probabilistic feature activation. In the model, source confusion between the prime and target results in positive priming, but negative priming occurs through the discounting of prime features. The appropriate level of discounting is calculated by lowering the evidence for features known to have been primed. Too little discounting of prime-activated features produces positive priming (i.e., these features lend extra evidence in favor of the primed choice) and too much discounting of features activated by primes produces negative priming (i.e., discounted features result in a relative loss of evidence as compared to no priming). This Bayesian model describes these results in terms of too little or too much discounting, but it does not specify why prime duration results in this change. Therefore, Huber and O'Reilly (2003) developed a model of dynamic discounting through neural habituation in a perceptual cascade of rate-coded neurons (see O'Reilly & Munakata, 2000, for a description of such rate-coded neurons). Huber (2008b) tested this account by parametrically varying prime duration, revealing a gradual non-linear transition from positive to negative priming. Beyond the form of this transition, other predictions of the habituation account were confirmed for the relationship between priming and forward masking (habituation predicted a relationship between the onset of negative priming and the elimination of forward masking) as well as the difference between repetition priming (deficits due to orthographic habituation) and associative priming (benefits due to lexical-semantic top-down expectations).

In the Huber & O'Reilly (2003) model, the transition from positive to negative priming was captured through synaptic depression. For most neurons, if a presynaptic cell

is constantly driven, a receiving postsynaptic cell initially shows a large depolarization, but after a short time (e.g., hundreds of milliseconds), postsynaptic depolarization drops off sizably (Abbott et al., 1997; Tsodyks, & Markram, 1997). Including synaptic depression allowed the model to capture the interaction between priming condition (e.g., target primed versus foil primed) and prime duration. For short duration primes, the prime presentation is too brief to cause much depletion of synaptic resources and so lingering activation from the prime results in enhanced processing of the primed choice word. Longer duration primes cause a build up of synaptic depression for the prime representation, which more than offsets the advantage of lingering activation. In other words, the prime is still active, but because synaptic resources are depleted, the relative advantage of this activation is lost. This depletion results in sluggish responding for the primed representation both in terms of its presentation as a briefly flashed target as well as its presentation as a choice alternative. Not only can this account explain choice accuracy, but it also correctly predicts reaction times; there is a speed up for chosen alternatives (correct response in the target primed condition and error responses in the foil primed condition) following brief primes but a slow down for chosen alternatives following long duration primes (Huber & Cousineau, 2004).

An analysis of individual differences in word priming

Assuming that different individuals differ in their word identification expertise and that the target duration needed for threshold performance is a good indication of these differences, temporal segregation through habituation predicted that individuals with lower threshold durations should reveal faster/stronger habituation as determined by their priming data. As seen in Figure 1.2, which is a median split of the priming data from

Huber (2008b, Experiment 1), this prediction was not confirmed. The median split was based on the target duration needed to achieve 75% accuracy in a condition where the prime was unrelated to the choices (i.e., threshold target duration without priming). The group on the left, labeled the 'fast group', is the 50% of participants with shorter threshold target durations while the group on the right, the 'slow group', was the upper 50%. Even though these groups differed in the duration needed to identify a word, both groups show similar habituation as evidence by the crossover point between the target primed and foil primed conditions. For both the fast and the slow groups of participants, 50 ms of priming produced the maximal positive priming and both groups of participants produced sizable negative priming for the 400 and 2,000 ms prime durations.
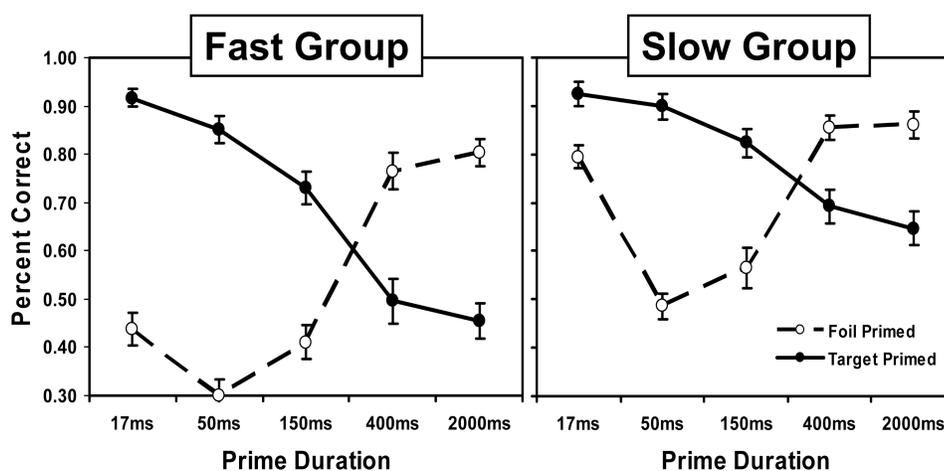
Figure 1.2. Reanalysis of word priming data. Results of Huber (2008 b Experiment 1 after splitting participants into fast and slow groups. Error bars are +/- 1 SE.

This reanalysis failed to support the predicted relationship between the speed of identification and the onset of negative priming. However, it may be that individual differences in the target duration needed for threshold word identification are not an

accurate measure of the speed of perceptual word identification. This may be the case if these target duration differences are due to low level visual processing differences (e.g., visual acuity) or perhaps motivational differences (e.g., staying alert on all trials). Alternatively, the observed individual differences might reflect differences in verbal processing (e.g., the ability to maintain words in verbal working memory) rather than differences in perceptual word identification. Regardless of the reason behind this failure, we sought to reexamine the situation with face priming. By using a face identification task, we manipulated the speed of identification within subjects by comparing upright and inverted faces. In addition, by using images of unfamiliar faces, the task was more clearly perceptual--unlike word identification, individual differences in verbal working memory should not matter because there is no associated verbal label for an unfamiliar face. Because the reported experiments concern the speed of habituation with faces, we next review the literature on face habituation.

Habituation effects with immediate face priming

There is a sizable literature examining habituation[2] aftereffects with faces, although these experiments do not typically use repetition priming or test different durations of initial exposure. One exception is a study by Bennett, Lleras, Oriet & Enns (2007) that examined identification for the emotion of a face. In the basic condition there was a priming benefit when the prime emotion matched the target emotion, but presentation of a neutral expression between the prime and target reversed this effect. Our habituation account might make sense of this result if the lingering activation of the prime, which is the source of positive priming, fades during presentation of the intervening neutral face. Habituation operates on slower time scale, and so an intervening

stimulus might serve to eliminate lingering activation but not eliminate lingering habituation, which is the source of negative priming.

Demonstrating that these face aftereffects are seen with manipulations of prime-target similarity, other studies find that extended viewing of a face can cause shifts in the categorical boundary along a series of morphed faces (Rothstein, Henson, Treves, Driver & Dolan, 2005), criterial shifts in a high dimensional 'face space' (Leopold, O'Toole, & Blanz, 2001), and distortions of face configuration (Rhodes, Jeffery, Watson, Clifford, & Nakayama, 2003) or emotion expression (Fox & Barton, 2007). In light of these results and others, we take the existence of face habituation as a given. However, instead of testing a single duration that produces strong habituation, we seek to examine the dynamic properties of face habituation by using multiple prime durations.

The Current Investigation

Experiment 1.1 was a replication of Huber's (2008b) Experiment 1, replacing words with upright views of computer generated faces. In light of the observed individual differences with faces, Experiment 1.2 tested whether target duration manipulations could produce these differences. Experiment 1.3 replicated the individual differences using images of real faces and included a baseline priming condition to rule out a specific strategic interpretation of the results. Besides using faces to test for greater variability in perceptual expertise, another advantage of using faces is that they are viewed in a conical upright orientation more so than words and other objects, and, correspondingly, faces reveal preferentially large inversion effects as compared to other objects (Yin, 1969). With faces, a given individual can be turned from a face expert into a face novice simply by inverting the face. Therefore, Experiment 1.4 replicated Experiment 1.1, but also

introduced the manipulation of face inversion, revealing an apparent interaction between face inversion and individual differences. Finally, the neural habituation account was quantitatively applied to the data of Experiment 1.4, which also allowed us to explore possible accounts of the interaction between individual differences and face inversion.

Experiment 1.1

Experiment 1.1 replicated the immediate repetition word priming found in Experiment 1 of Huber (2008b) except that upright faces were used rather than words. A range of prime durations were tested with both target primed conditions (the benefits of priming), in which the target repeated the prime face, as well as with foil primed conditions (the costs of priming), in which the incorrect choice repeated the prime face. It was predicted that individuals who needed less time to identify the briefly flashed target would also reveal more rapid face habituation. Time to identify a face was measured by the threshold target duration needed to achieve 75% accuracy. Habituation was measured by the prime duration needed to produce negative priming (foil primed performance better than target primed performance).

Method

Participants. Twenty-eight undergraduate psychology students participated in this experiment for course credit.

Procedure. As seen in Figure 1.1, the basic task was to identify the briefly flashed target faces that appeared in the center of the screen, which was immediately replaced with a pattern mask. All faces were novel computer generated male faces and no face appeared on more than one trial during the experiment. Immediately prior to the

target face, a prime face appeared for durations ranging from 17 ms (i.e., subliminal face priming) to 2,000 ms (i.e., excessive face priming). Because these experiments investigated immediate repetition priming, it was not possible to present the prime in the exact same location as the target. In other words, a spatial cue was needed to separate primes from targets. Therefore, two identical versions of the prime face were displayed side by side. In this manner, the prime face also served as a forward mask of the target face. Face identification was tested by means of a forced-choice between the target face and a foil face. Performance was not speeded and accuracy feedback was provided on every trial. Participants were explicitly informed that half the time the prime was identical to the correct answer and half the time the prime was identical to the incorrect answer, and that there was no effective response strategy that used knowledge of the prime alone. Thus, the most effective strategy was to focus on the center and wait for the briefly flashed target face.

The experiment began by establishing the target duration each participant needed to achieve accuracy at the threshold value of 75%. Once this target duration was established, it remained fixed for the remainder of the experimental session. In the analyses of individual differences, a split half grouping of these target durations defined slow and fast face identification groups of participants.

Two variables were manipulated within subject: prime duration with 5 levels (17, 50, 150, 400 or 2,000 ms), and prime type with 2 levels (target primed or foil primed). There were 32 trials for each condition occurring randomly throughout the experiment. Each trial began with a fixation cross presented for a duration such that the combined duration of the fixation cross and the prime face was 2,500 ms. After the target face was

presented, it was immediately masked with a checkerboard pattern mask. The mask was presented for a length of time such that the duration of mask and target combined was 500 ms, thus maintaining a constant duration between onset of the target and onset of choice alternatives. Following presentation of the mask, the target and foil faces of the forced choice decision appeared above and below the center line, separated by approximately 6 degrees of visual angle. Responses were collected by keyboard. The choice faces were presented until a decision was made, after which feedback was presented for 1,500 ms. The position of the target face was randomly counterbalanced across top versus bottom (i.e., in every condition, the correct answer was the top face for half the trials).

During the initial block of trials that determined the appropriate target duration for threshold performance, primes that were different than both the target and the foil were used with a prime duration of 150 ms. In other words, these were 'neither-primed' trials. Unlike Huber's (2008b) Experiment 1, which used a staircase procedure to obtain threshold target duration, this experiment mapped out the psychometric function by testing specific target durations of 33, 50, 67, and 83 ms equally often as randomly mixed across trials. If 75% identification was not achieved with the 83 ms target, a target duration of 100 ms was adopted for the remainder of the experiment. Otherwise, the target duration that produced performance closest to 75% was selected.

At the beginning of the experiment, participants completed 12 practice trials, with the first 6 using a target duration of 167 ms (which was sufficiently easy for all participants) and the second 6 using a target duration of 83 ms to accommodate participants to threshold presentations. Participants then completed a block of 80

threshold determination trials (20 at each target duration, randomized), followed by 4 blocks of 80 priming trials (32 at each prime duration for each prime type, randomized). Before the third priming block, participants completed a ten-minute word search puzzle to give them a rest. The experiment was self paced and lasted approximately 60 minutes on average.

Materials. There were 1,000 face images randomly generated using the FACES computer program that creates faces from composites of features including head shape, eyes, nose, mouth, ears, eye brows, chin shape and facial hair. This was achieved by repeated use of the software's random function followed by elimination of faces that did not appear plausible or sufficiently male. Facial hair and different hair styles were allowed. All features were sampled randomly with replacement by the program, and thus the faces had varying degrees of similarity. All face images were grayscale and 150 by 137 pixels. As viewed each face was approximately 5 degrees visual angle in height and 3.5 degrees in width. The pattern mask presented after the target face was 150 by 140 pixels, consisting of a 5 by 5 grid of blocks. On each trial, a random pattern mask was created by sampling each block of the grid separately from 256 levels of gray.

Equipment. The experiment was conducted on CRT monitors with a 60 HZ vertical refresh rate. Presentation times were synchronized with the display refresh rate. All items were displayed on a gray background and the display resolution was set to 640 by 480 pixels.

Results and Discussion

Collapsing across individuals, a repeated measures ANOVA revealed a main effect of prime type, $F(1,27) = 21.46$, $p < .001$, prime duration, $F(4,108) = 15.33$, $p <$

.001, and a prime type X prime duration interaction, $F(4,108) = 14.51, p < .001$. There was no difference between the target primed and foil primed conditions either at 400 ms, $t(27) = -1.35, p = .18$, or at the longest prime duration of 2,000 ms, $t(27) = 0.81, p = .43$. However for the 50 and 150 ms prime durations, there was sizable positive priming, with target primed performance better than foil primed performance (for both tests, $t(27) < -6.9, p < .001$). This was also true of the shortest prime duration, however the magnitude of difference was smaller, $t(27) = -2.68, p = .013$. Thus, the interaction was a pattern of increasing then decreasing positive priming. This qualitative pattern was the same as previously found with immediate repetition priming of words. However, contrary to words, longer prime durations only served to eliminate positive priming, rather than produce negative priming. Next, individual differences are analyzed based on threshold target duration.

Results from the threshold determination block of trials are presented in the Appendix, which shows results collapsed across all participants, as well as for the median split according to the target duration required to achieve threshold performance. These median split groupings will hereafter be referred to as the fast group (shorter target durations to achieve threshold performance) versus slow group (longer target durations to achieve threshold performance), although this is not meant to imply that there are literally two distinct groups of people[3]. Instead, this is merely a convenient way to parse the data while maintaining sufficient numbers of trials. This median split was created in the following manner for each experiment. Participants were first rank ordered by the target duration needed for threshold performance. Because there are only a few possible target durations due to the refresh rate of the monitors, the median target duration did not

cleanly divide the participants into two equal sized groups. Therefore, individuals with the median target duration were further divided according to their average accuracy during the threshold block of trials. This provided a fine grained measure so that these median target duration participants could be assigned to the fast group or the slow group to produce two equal sized groups. If the number of participants was odd, the median participant was included in the fast group. For Experiment 1.1, each group consisted of 14 participants. The average target duration thresholds for the fast and slow groups were 46.43 ms (SE = 3.11) and 83.33 ms (SE = 3.91), respectively.
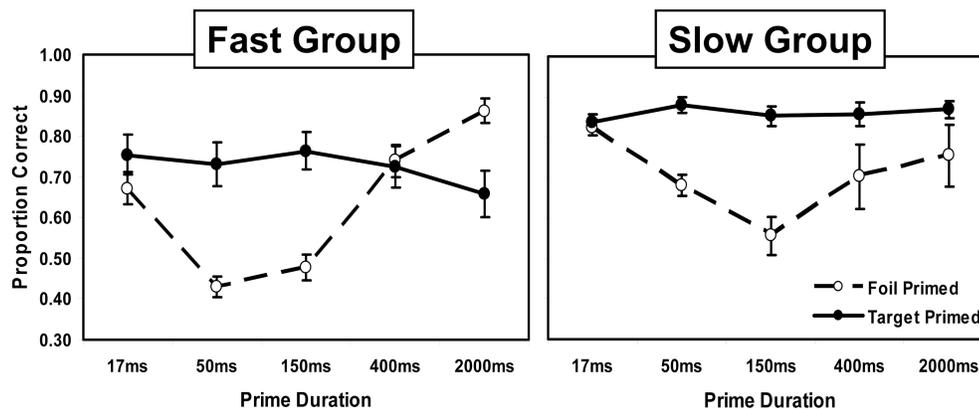
Figure 1.3. Experiment 1.1 results. The results of Experiment 1.1 broken down by fast versus slow median split of participants based on the target duration necessary for threshold performance. Error bars are +/- 1 SE.

Figure 1.3 shows the effects of prime duration and prime type separately for the fast and slow median split groups of participants. Collapsing across conditions, accuracy was .70 (SE = .03) and .77 (SE = .03) for each group respectively. A mixed design ANOVA was conducted by adding the factor of fast/slow group to the analysis of prime type and prime duration. There was a main effect of prime type, $F(1,26) = 21.90$, $p <$ .001, participant group, $F(1,26) = 6.62$, $p < .016$, and prime duration, $F(4,104) = 17.92$, p

< .001. Additionally, there was an interaction between prime duration and group, $F(4,104) = 5.556$, $p < .001$, between prime type and prime duration, $F(4,104) = 17.55$, $p < .001$, and a three-way interaction between prime type, prime duration and group, $F(4,70) = 6.65$, $p < .001$.

As seen in Figure 1.3, the two groups differed on the pattern of the two-way interaction between priming condition and prime duration. Both groups demonstrated an increase followed by a decrease in positive priming as a function of prime duration. However, for the fast group, the prime duration with the maximal positive priming was 50 ms. In contrast, for the slow group, the prime duration with the maximal positive priming did not occur until 150 ms of prime exposure. Furthermore, at the longest prime duration of 2,000 ms, the fast group produced negative priming, with performance in the target primed condition actually worse than the foil primed condition, $t(12) = 4.38$, $p = .001$. In contrast, there was no difference between these conditions for the slow group at the 2,000 ms prime duration, $t(12) = -1.40$, $p = .185$. Thus, only the fast group revealed a full crossover from positive to negative priming.

A main effect of group is visible in Figure 1.3, with the slow group revealing higher accuracy compared to the fast group. Ideally, the use of individually set threshold target durations should have equated performance across the groups. However, the target threshold was determined based on a limited number of forced choice trials and thus reflects some degree of chance. Due to chance, some people did better during the threshold block of trials than they would have over more trials and so these individuals were assigned a target duration that was briefer then their true threshold. Chance also worked in the opposite manner, serving to assign some individuals a target duration that

was too long. When these misassigned participants subsequently performed the priming trials, they did better or worse than expected because they were assigned a target duration that was correspondingly too long or too short. Because the assignment to the slow or fast group is a selection based on extremes, the role of chance worked against the groupings in a systematic manner. Thus, the failure to fully equate the slow and fast groups was expected due to 'regression to the mean', producing better performance for the slow group than for the fast group.

The distinction between the slow and fast group was based on accuracy during the threshold determination block of trials. However, if there is a speed accuracy tradeoff in this paradigm, there may also be a difference in reaction time between the groups. Such a speed difference may be relevant to the transition from positive to negative priming if it takes longer to strategically correct for the influence of the prime. Thus, perhaps the slow group participants responded too quickly, not allowing time to discount the effect of the prime face. We calculated the median reaction time for each participant's correct and incorrect responses both overall and by priming condition, excluding reaction times recorded as shorter than 100 ms or longer than 6,000 ms. No significant differences between the groups nor interactions with the grouping variable were found. The pattern of reaction times over different priming conditions were similar to those found with words, which suggested a race model between choice alternatives (Huber & Cousineau, 2004). For correct trials, reaction time followed a similar pattern to accuracy data, although mirrored, with shorter reaction times to target primed trials, and longer reaction times to foil primed trials. As with the accuracy data, the reaction time difference between the priming conditions increased then decreased with increasing prime duration.

Incorrect reaction times mirrored this pattern (target primed trials were slower than foil

primed trials) but otherwise appeared similar. In summary, explanations based on speed

accuracy tradeoff are ruled out because there were no reaction time differences between

the groups and because correct reaction times were faster for the conditions that produced

higher accuracy.
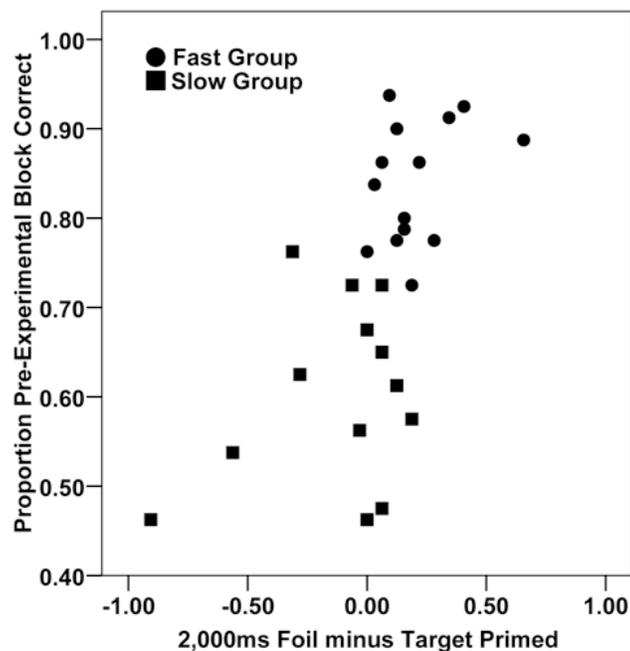
Figure 1.4. Scatter plot of individual differences. This plot compares comparing average accuracy during the threshold determination block of trials versus the accuracy difference between the target primed and foil primed conditions at the 2,000 ms prime duration. Circles and squares indicate individuals placed into the fast and slow groups respectively.

A post hoc median split of participants is merely a convenient method for

exploring the behavior of individuals who are better versus worse at identifying the

briefly flashed target face. The theory that motivated this work makes no particular

predictions for distribution of face processing abilities but it does predict that there

should be a continuous relationship between the speed of face identification and the

speed of habituation. Figure 1.4 provides a scatter plot of individuals comparing their mean accuracy over the threshold determination block of trials (our proxy for rapid face identification) to the difference between foil and target primed accuracy with a 2,000 ms prime duration (our proxy for the strength of face habituation). The correlation between these measures was $r = .598$, $p = .001$, suggesting that the relationship between rapid face identification and face habituation is continuous. As seen in the figure, the two leftmost individuals were possibly outliers. However, elimination of these individuals did not change any of the statistical conclusions drawn from the median split analyses.

Comparison of Figure 1.2 and Figure 1.3 reveals that individual differences are apparent with faces but not words. This comparison also reveals that face habituation (as indicated by the relationship between the priming conditions as a function of prime duration) appears to occur more slowly than word habituation. However, it is misleading to compare prime durations between face and word priming considering that the displays in the two experiments differed in a variety of ways, such as retinal size (the faces were much larger) and the type of masks (the words were masked with line segments rather than checkerboards). Another key difference is that the faces were unknown, which makes it difficult to assign a verbal label to the face, whereas the words were high frequency words.

In summary, face priming produced increasing and then decreasing positive priming as a function of prime duration. Furthermore, there was a relationship between the speed of face identification (target threshold duration) and the speed of the prime duration effects. If the observed individual differences in the transition from positive to negative priming with face priming reflect true differences in face processing expertise,

this confirms the prediction that a higher degree of expertise should correspond to a faster rate of habituation. However, the median split analysis confounds rate of habituation with target duration. This leaves open the possibility that negative priming only occurs for conditions of reduced stimulus energy for the target (i.e., shorter target durations). In other words, a weak target may be more susceptible to the negative effect of a strong prime. The opposite result of this prediction has been found with word priming, in which case reducing target duration actually flipped negative priming to positive priming (Huber et al. 2002; Weidemann et al. 2008), which makes this alternative explanation seem unlikely. Nevertheless, the interaction between target duration and prime duration has not been tested with faces. Therefore, Experiment 1.2 replicated Experiment 1.1 while eliminating the setting of different target durations for different individuals. Instead, Experiment 1.2 fully crossed different prime durations with different target durations for each individual. Thus, the relationship between priming and target duration was examined as a within-subjects manipulation.

Experiment 1.2

Experiment 1.2 manipulated target duration within subjects to see if the interaction between prime type and prime duration is affected by target duration. The between subject difference found in Experiment 1.1 could be explained by the target duration differences between the fast and slow groups. If this explanation is correct, then the within subjects design used in Experiment 1.2 should produce a priming reversal for short target durations but not for long target durations. Experiment 1.2 did not use an initial block of trials to set the target duration, but instead contained 'neither primed'

trials during the experimental blocks of trials to assess individual differences. These trials allowed separation of the participants into fast and slow groups to test for replication of the individual differences observed in Experiment 1.1. Including target duration as a manipulation crossed with prime duration greatly increased the number of conditions, and so only three prime durations and three target durations were tested. In addition, unlike Experiment 1.1, half of the trials reused previously seen faces.

Method

All stimuli, equipment, and procedures were identical to Experiment 1.1 except as noted.

Participants. Fifty-eight undergraduate psychology students participated in this experiment for course credit.

Procedure. Participants were given 48 practice trials, with progressively decreasing target durations, followed by 5 blocks of 96 priming trials with breaks between blocks. The trial procedure was the same as described for Experiment 1.1 except that target duration was manipulated and the choice faces were offset to the left and right instead of top and bottom. Target durations were tested at 33, 50, and 100 ms and were completely crossed with prime durations of 17, 150 and 2,000 ms. In addition to the target and foil primed trials, there were trials that presented a prime face for 150 ms that was unrelated to either choice (i.e., neither-primed trials). Thus, these trials were identical to threshold trials from Experiment 1.1 and were included to determine the threshold of each participant. The 96 trials within each block consisted of 72 target or foil primed trials, with 4 replications at each combination of prime and target duration, and 24 neither primed trials, with 8 replications at each target duration. Because this experiment

involved a greater number of total trials, half of the trials used entirely new faces while the other half of trials used previously seen faces.

Results and Discussion.

A repeated measures analysis determined that there was no main effect of reusing faces, $F(1,56) < 1$, $p = .586$, but there was an interaction between reuse of faces and prime duration, $F(2,114) = 4.201$, $p = .02$. However, reuse of faces did not interact with the prime duration by priming condition interaction, $F(2,56) = 2.273$, $p = .112$, and so subsequent analyses collapsed over trials that used new and reused faces.

To test the relationship between target duration and face priming, we conducted an ANOVA on the factors of prime type, prime duration, and target duration. While there was a 3 way interaction of prime type, prime duration and target duration, $F(4,228) = 5.30$, $p < .001$, it was not in a direction compatible with the proposal that shorter target durations produce a more rapid transition to negative priming. Instead, the nature of this interaction was for the opposite pattern, with a lesser degree of positive priming for longer target durations. Pairwise tests found that the target primed condition was significantly greater than the foil primed condition at all prime durations (for all tests, $t(56) < -3.04$, $p < .004$) except for 100 ms targets and 2,000 ms primes, $t(57) = -1.92$, $p = .06$. Focusing on this result, we performed an ANOVA using the longest prime duration, revealing a target duration X prime type interaction, $F(2,114) = 26.39$, $p < .001$. As seen in Figure 1.5, longer target durations caused the benefit of target priming versus foil priming to *decrease* (i.e., the data moved in the direction of negative priming with increasing target duration). An explanation of Experiment 1.1 based on target duration predicted that priming should have instead moved in the direction of positive priming

with increasing target duration. Thus, the use of different target durations for different individuals in Experiment 1.1 was not the cause of the individual differences in the transition from positive to negative priming.

Results from the 150 ms neither primed trials are presented in the Appendix, both for the collapsed results and for the results broken down by median fast/slow split. Based on these trials, threshold target durations were determined for each participant with the same procedure as Experiment 1.1 (keeping in mind that unlike Experiment 1.1, Experiment 1.2 did not set the target durations to these individual threshold durations). The mean threshold target duration for the fast group was 62 ms (SE = 4.97). All participants in the slow group had threshold target durations of 100 ms (note that only 33, 50, and 100 ms target durations were tested in this experiment and so this estimate of threshold target duration may be somewhat inflated because of the absence of intermediate target durations).
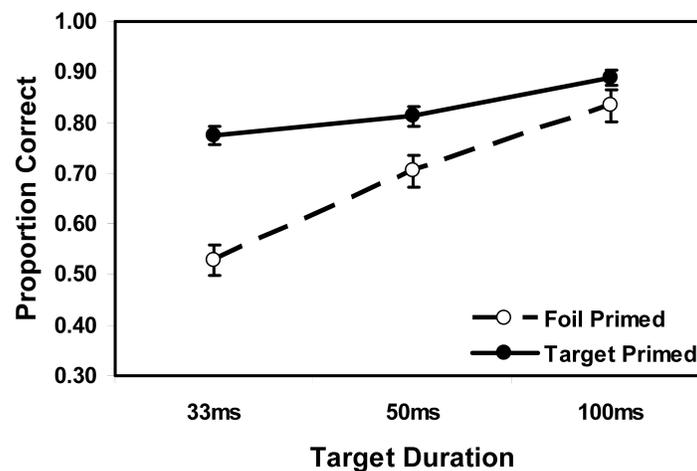
Figure 1.5. Collapsed Experiment 1.2 results. This plot collapses over all participants for the case of strongest prime discounting, which occurred in the 2,000 ms prime duration conditions. Error bars are +/- 1 SE.
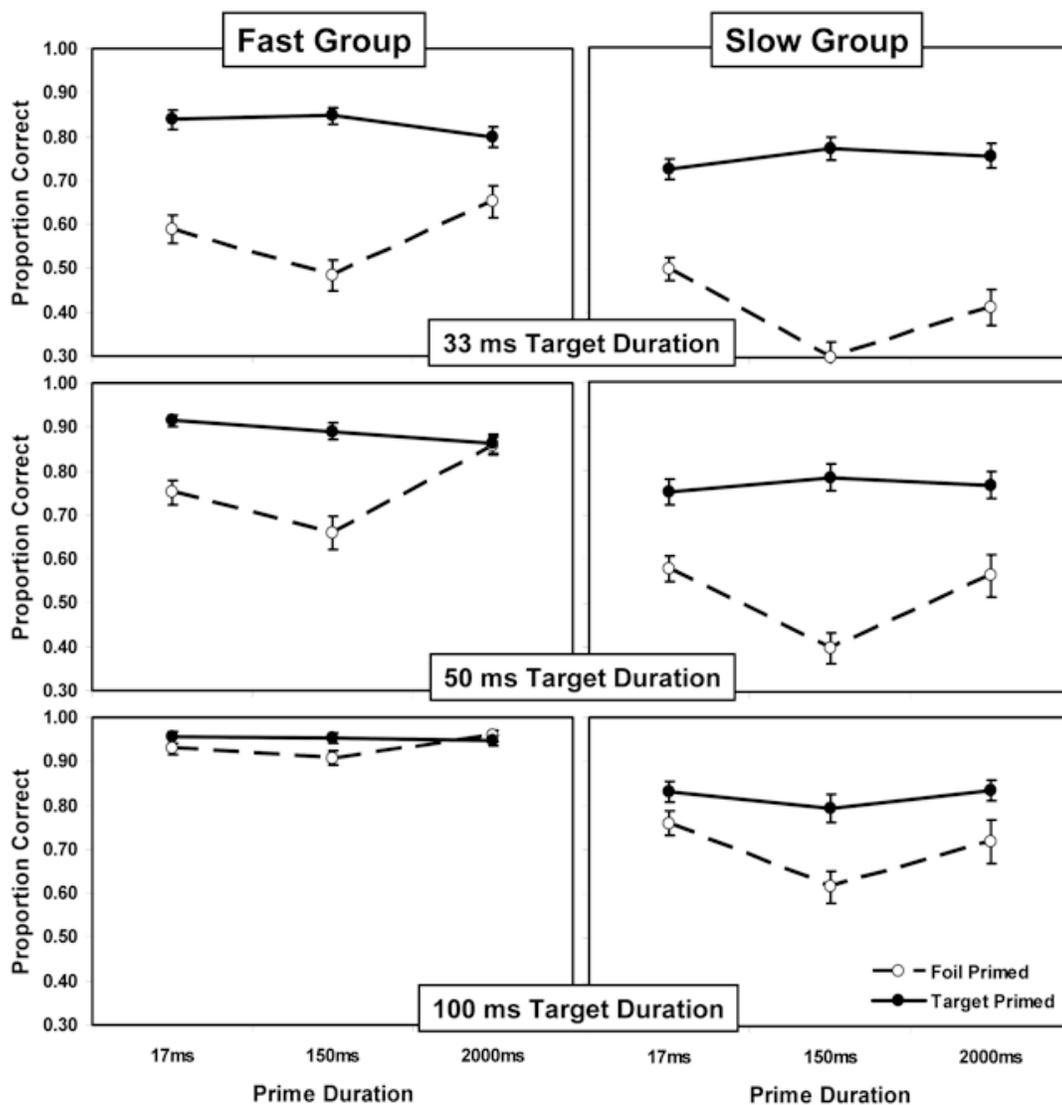
Figure 1.6. Experiment 1.2 results. The results are split by target duration and by fast/slow median split of participants. Error bars are +/- 1 SE.

Figure 1.6 shows the effects of prime type and prime duration on accuracy for the fast and slow groups. To test for individual differences, we conducted a mixed design ANOVA testing the factors of prime type, prime duration, target duration, and median split of the threshold target durations on accuracy. Accuracy was higher with longer

target durations, $F(2,112) = 243.24$, $p < .001$, for the fast median split group, $F(1,56) =$ 61.33, $p < .001$, and for target primed trials, $F(1,56) = 127.77$, $p < .001$. There was a main effect of prime duration, $F(2,112) = 24.33$, $p < .001$, and a prime type X prime duration interaction, $F(2,112) = 18.88$, $p < .001$, similar to Experiment 1.1. There were also interactions of target duration X median split, $F(2,112) = 5.80$, $p = .004$, prime type X median split, $F(1,56) = 11.62$, $p = .001$, prime duration X median split, $F(2,112) = 3.079$, $p = .05$, and target duration X prime type, $F(2,112) = 86.63$, $p < .001$. Importantly, there was an interaction between prime type, prime duration and median split, $F(2,112) = 6.37$, $p = .002$, thus replicating the individual differences of a more rapid elimination of positive priming for the fast participants found in Experiment 1.1, although without the crossover into negative priming. For the slow group, target primed accuracy was higher than foil primed accuracy at all prime durations and target durations (for all tests, $t(28) < -2.38$, $p < .02$). For the fast group, target primed accuracy was higher than foil primed accuracy at all prime and target durations (for all tests, $t(28) < -2.98$, $p < .006$), except when the prime duration was 2,000 ms and the target duration was either 50 ms, $t(28) = -.21$, $p = .836$, or 100 ms, $t(28) = .9$, $p = .38$, or when prime duration was 17 ms and target duration was 100 ms, $t(28) = -1.87$, $p = .07$.

Emphasizing the importance of setting target duration separately for each participant to avoid ceiling or floor effects, some subjects in the 100 ms target duration achieved 100% accuracy. To compensate for this ceiling effect, we reanalyzed the data using d′ (Egan, 1975). Adjustments of +/- 1/40 were made to conditions where a subject achieved 0 or 100% accuracy. No changes in the significance of any main effect or interaction were found.

In summary, Experiment 1.2 demonstrated that use of longer target durations was not the underlying cause of individual differences. While the effect of prime duration was not as striking as in Experiment 1.1, the same general pattern of more rapid elimination of positive priming for participants requiring more brief target flashes was obtained both as measured with proportion correct and d′. Thus, these individual differences in the prime duration effect are not caused by different target durations. Instead, these differences must reflect either differences in how participants perform the task or different face identification abilities.

Experiment 1.3

One difference between Figures 2 and 3 is that there appears to be more symmetry between the target primed and foil primed conditions with word priming; for word priming, prime durations that decreased foil primed performance also increased target primed performance. However, even with word priming, note that the two conditions are not perfect mirror images of each other. These apparent asymmetries between the target primed and foil primed conditions are in truth misleading because neither figure includes the baseline neither-primed condition that presents an unrelated word or face for the appropriate prime duration. Although the above face experiments did not include this condition, the word priming experiment (Huber, 2008b, Experiment 1) did include such a condition, which revealed a u-shaped pattern as a function of prime duration for the baseline neither primed condition. In other words, an intermediate prime duration proved to be the most disruptive to performance even though the prime was unrelated to the choices. In light of this 'forward masking' effect for the baseline condition with words,

there was in fact a fair degree of symmetry for the costs and benefits for the target primed and foil primed conditions.

Due to the differences found between words and faces, it is important to evaluate the priming effects we have found relative to a baseline condition, where the identity of the prime face is neither the target nor the foil choice. Without this condition, we cannot make strong conclusions about the symmetry of priming. Furthermore, considering that the target primed condition appears to be unaffected by priming duration, it is possible that no priming occurred in this condition (rather than increasing then decreasing priming against a baseline that falls then rises). In Experiment 1.3 we include this baseline condition to rule out this possibility. In addition, because there is evidence that computer generated faces may be problematic in the study of face processing (Carlson, & Gronlund, 2007), we used photographs of real faces rather than computer generated faces. These photographs were cropped to include just the interiors of faces, which should reduce any reliance on hairstyle in the identification process. Similar to word priming, we hypothesized that the neither primed baseline condition would lie in between the other priming conditions and thus reveal a u-shaped forward masking effect as a function of increasing prime duration.

Methods

All stimuli, equipment, and procedures were identical to Experiment 1.1 except as noted.

Participants. A total of one hundred and two undergraduate psychology students participated in this experiment for course credit.

Procedure. In addition to target and foil primed conditions there was also a neither-primed condition where the prime identity was different from both the target and

foil choices. Each participant completed 10 priming practice trials. This was followed by 2 blocks of 40 threshold trials testing target durations of 50, 67, 83, and 100 ms. Participants not achieving 75% accuracy at the 100 ms target duration were assigned a target duration of 117 ms. Participants next completed 5 blocks of 60 priming trials. Across all blocks, each participant completed 20 trials of each prime condition. Prime durations of 17, 50, 150, 400 and 2,000 ms were tested. Between blocks, participants took breaks of at least 15 s. Because of the smaller number of available faces, each face was repeated four times in either the threshold trials or priming trials.

Materials. Face stimuli were created by cropping the faces of 238 White faces (Tanaka & Pierce, 2009; Tanaka, (2007), "A race face database") to be 149 pixels in width by 155 pixels in height such that they contained only the interior of each face. This was done to place primes as close together as possible, and also to reduce similarity between the faces (all faces in the database have the same exterior features).

Results and Discussion

The priming results split for the fast and slow groups are presented in Figure 1.7. We conducted a repeated measure ANOVA on the factors of prime type, prime duration, and fast/slow group. There was a significant three way interaction between prime type, prime duration, and participant group $F(8,800) = 9.291$, $p = .002$, as well as a two way prime type and prime duration $F(4,400) = 45.20$, $p < .001$, and main effects of participant group, $F(1,100) = 10.2$ $p = .002$, prime type, $F(24,200) = 208.22$, $p < .001$, and prime duration, $F(4,400) = 43.64$, $p < .001$. Planned comparisons of the difference between the prime type conditions for the 2,000 ms prime duration trials revealed positive priming. Prime (accuracy was lower for foil primed than the target foil primed condition) for the

slow participants, $t(50) = 2.14\ p = .037$, but for fast participants, positive priming was eliminated, $t(50) = 1.138\ 1.13,\ p = .26$.

As expected, when a prime unrelated to either choice was presented, accuracy was between that of the target primed and foil primed conditions. Thus, the priming effects are fairly symmetric about this baseline condition and there is indeed a rise and fall in priming for the target primed condition. Furthermore, as expected, the baseline condition revealed u-shaped forward masking. According to the neural habituation model, this is explained through the rise and fall of activation to the prime face combined with inhibition between any concurrently active faces.
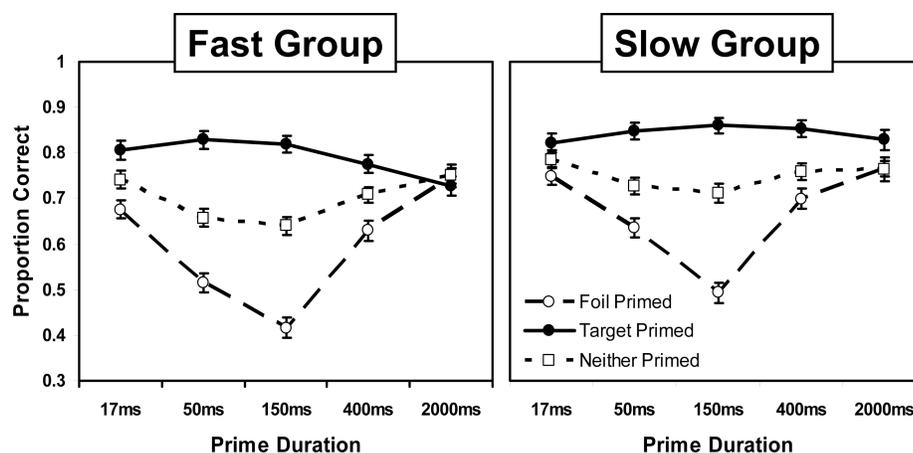
Figure 1.7. Experiment 1.3 results. These results are broken down by fast versus slow median split of participants based on the target duration necessary for threshold performance. Error bars are +/- 1 SE.

Similar to Experiment 1.1, there was an increase and decrease of positive priming with increasing prime duration. Furthermore, the individual differences in these priming effects were again replicated. This generalizes the effect of Experiments 1 and 2 to the use of real face images. Unlike Experiment 1.1, there was no crossover to negative priming for the fast group of participants. However, negative priming is not a qualitative

prediction of the habituation account. According to this account, habituation is designed to offset the unwanted effect of a previous face blending with a subsequent face. Ideally, this offset through habituation should perfectly balance between integration and separation, producing no difference between the target primed and foil primed conditions. With particularly salient primes, there may be excessive habituation, which produces negative priming. For some reason, the images used in Experiment 1.1 produced this excessive habituation whereas these real face images did not. There are a number of possible causes for the lack of negative priming compared to Experiment 1.1, including the change to real faces, repetition of faces, inclusion of exterior face features, and perceived task difficulty. Nevertheless, the predicted differences between groups in terms of the elimination of positive priming as a function of prime duration were confirmed.

Experiment 1.4

Experiments 1-3 demonstrated the predicted relationship between the speed of face identification and the speed of face habituation by comparing individuals. However, because this involved a subject-variable, a comparison across individuals may involve more than just differences in the speed of face identification. Therefore, Experiment 1.4 sought to manipulate face expertise in a within-subjects design. This was done by comparing upright versus inverted faces.

While the basic mechanisms that support face processing are hotly debated (Gauthier & Logothetis, 2000; Gauthier & Nelson 2001; Gauthier, Tarr, Anderson, Skudlarski, & Gore, 1999; Kanwisher, 2000; McKone, Kanwisher, Duchaine, 2006) it is uncontested that we are skilled at identifying upright faces. Investigations into the nature

of face expertise suggest that face identification is reliant on information contained in more than just feature identity, such as the precise configurations of features within a face or holistic information over large regions of a face (Maurer, Le Grand & Mondlock, 2002; Rakover, 2002). Configuration refers to the unique spatial arrangement of features (e.g., eyes, nose, mouth) on each face. Similarly, holistic face processing occurs without decomposition into separate face parts (Farah, Wilson, Drain & Tanaka, 1998; although for evidence that holistic face processing can result from decisional factors see Richler, Gauthier, Wenger & Palmeri, 2008; Wenger & Ingvalson, 2002, 2003). It has been suggested that face inversion disrupts holistic or configuration information (Farah, Tanaka & Drain, 1995; Freire, Lee & Symons, 2000; Leder & Bruce 2000; Rossion & Gauthier, 2002), leading to the face inversion effect, which is a disproportionate deficit for inversion of faces as compared to inversion of other objects (Yin, 1969). In the present work, the aspect of face perception that is disrupted by inversion is not under investigation, and so we will refer to this as face configuration for brevity, keeping in mind that identical arguments could be made for holistic information. Regardless of the mechanism, we seek to exploit face inversion to manipulate the strength of response for the information that underlies face identification.

Experiment 1.4 manipulated face expertise by comparing priming of upright versus inverted faces. The design was identical to Experiment 1.1, except that half the trials used upright presentations of prime, target, and choice faces, and the other half of trials used inverted faces for all presentations. Because both upright and inverted faces were drawn from the same pool of face stimuli, this allows manipulation of expertise while equating many low-level visual aspects of the images. It was predicted that because

recognition ability for inverted faces is diminished, there should be less habituation for inverted faces. According to this prediction, there should be a slower transition from positive to negative priming for immediate repetitions of inverted faces.

Method. All stimuli, equipment, and procedures were identical to Experiment 1.1 except as noted.

Participants. Forty undergraduate psychology students participated in this experiment for course credit.

Procedure. Half of the trials for each condition, including practice and threshold blocks, used inverted face stimuli for primes, targets, and choice faces. This resulted in 10 trials per condition in threshold trials, and 16 trials per condition in priming trials. The calculation of each subject's threshold duration was taken over both upright and inverted threshold trials and the same target duration was used for both trial types during the experimental blocks of trials.

Materials. Inverted faces were 180 degree rotations from the upright faces. Each face was only presented once, either upright or inverted.

Results and Discussion

Figure 1.8 shows the interaction between prime type and prime duration separated by fast and slow groups and also by face orientation. These results also appear in the Appendix, along with the threshold results, and results collapsed across groups. It is clear that for upright faces, both the overall prime duration effect and the median split differences were replicated. First, we consider only the upright face conditions, which replicate the results of Experiment 1.1, and then we consider inverted faces and include face orientation as a factor in the ANOVA.
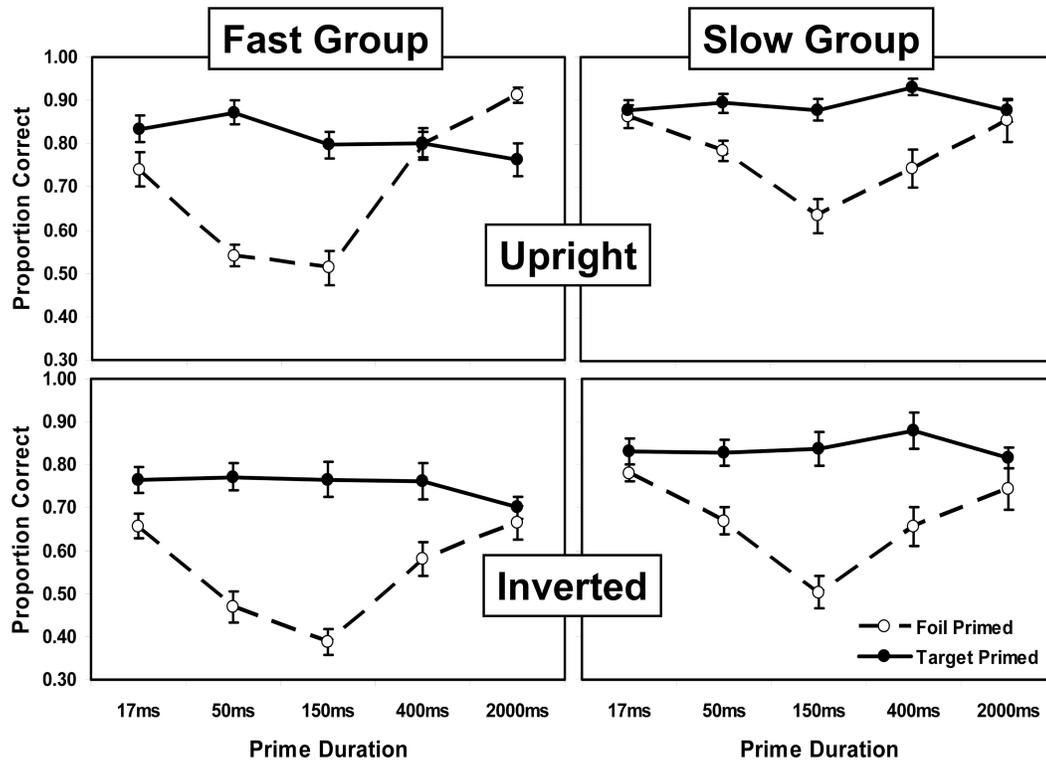
Figure 1.8. Experiment 1.4 results. The results are separated by upright and inverted conditions and by fast/slow median split of participants. Error bars are +/- 1 SE.

For upright faces, as in Experiment 1.1, there were effects on accuracy of prime type, $F(1,38) = 80.89$, $p < .001$, prime duration, $F(4,152) = 31.063$, $p < .001$, and a prime type X prime duration interaction, $F(4,152) = 32.36$, $p < .001$. Replicating the median split distinction, there was a three way interaction between prime type, prime duration and participant group, $F(4,152) = 8.16$, $p < .001$. There was also a main effect of group, $F(1,38) = 12.87$, $p = .001$, where overall accuracy was higher in the slow group, which as elaborated previously is explained by regression to the mean.

For inverted faces, the time course of the prime effect looks similar for both groups. Accuracy was lower for inverted faces compared to upright faces, $F(1,38) = 160.04$, $p < .001$, as expected. The effect of prime type was larger with inversion, $F(1,38)$

= 30.10, $p < .001$, presumably due to the lack of any negative priming for inverted faces, even for the fast group. The three way interaction of inversion, prime type, and prime duration was marginally significant, $F(4,152) = 2.15$, $p = .08$. The four-way interaction between orientation, prime type, prime duration and participant group, $F(4,152) = 2.23$, $p = .07$, was marginally significant, and a planned comparison of prime type, orientation, and participant group for only the 2,000 ms prime duration highlighted the manner in which face inversion affected the groups differently, $F(1,38) = 4.16$, $p = .05$. For the fast group, 2,000 ms of foil priming resulted in higher accuracy than target priming (i.e., a pattern of negative priming) when viewing upright faces, $t(19) = 5.02$, $p < .001$, but at this prime duration, there was no effect of prime type when viewing inverted faces, $t(19) = -.85$, $p = .41$. In contrast to this inversion priming effect for the fast group, for the slow group there was neither a priming effect for upright faces, $t(19) = -.498$, $p = .62$, nor for inverted faces, $t(19) = -1.306$, $p = .21$ in the 2,000 ms prime duration conditions.

In summary, the individual differences with upright faces were again replicated, but these differences apparently vanished with inverted faces. For inverted faces, the rate of prime habituation, as revealed by the transition from positive to negative priming, was essentially the same for both participant groups, and similar to the upright slow group. For the fast group, the comparison between upright and inverted faces demonstrates that a within subjects manipulation of expertise can produce the predicted change in the speed of habituation. However, the absence of a habituation face inversion effect for the slow group was unexpected. To explore alternative interpretations of this apparent interaction between individual differences and face orientation, the habituation model was applied to these data, as reported next.

Experiment 1.5: Simulation Study

As predicted from a neural habituation explanation of face priming, both individual differences and face orientation modulated the transition from positive to negative priming—in both cases, better performance (i.e., upright faces or individuals with low perceptual thresholds) produced more sizable negative priming following long duration primes. These results were predicted under the assumptions that activation drives habituation and that a higher degree of performance corresponds to stronger activation. However, the precise manner in which individuals differ and upright versus inverted faces differ is unclear, and there may be several alternative implementations of this theory that are qualitatively compatible with these results. In this simulation study, we explored one such model to assess whether it was able to quantitatively handle these results, and whether it could shed light on the apparent interaction between face orientation and individual differences. Additionally, by using the same model that explained similar results with word priming (Huber, 2008 b), this simulation study might explain the different time course for word versus face priming. The dynamic activation functions of the model and the three layer structure are identical to that presented by Huber and O'Reilly (2003), and are presented here briefly for completeness.

Model Structure

As applied to words, Huber (2008b) assumed that the top layer of this three layer model represented lexical-semantic information, the mid level represented orthographic information, and the bottom visual line segments, thus accounting for different dynamics for masks, orthographic priming and associative priming. In the current application, we

do not have results that constrain the exact nature of the information at each level of face processing beyond inversion effects (i.e., we didn't manipulate face similarity), and so we more generically refer to the bottom layer as visual processing, the mid layer as processing of face parts (which are equally represented when inverted) and the top layer as processing of face wholes (which are disrupted when inverted). We are not making any strong claims about what these face parts or whole are (e.g. we are not claiming that these face parts are necessarily face features such as the eyes, nose and mouth), just that the higher level of perception is disrupted by inversion. Thus, the goal of this simulation study is to investigate the dynamics of face processing rather the representation of faces.
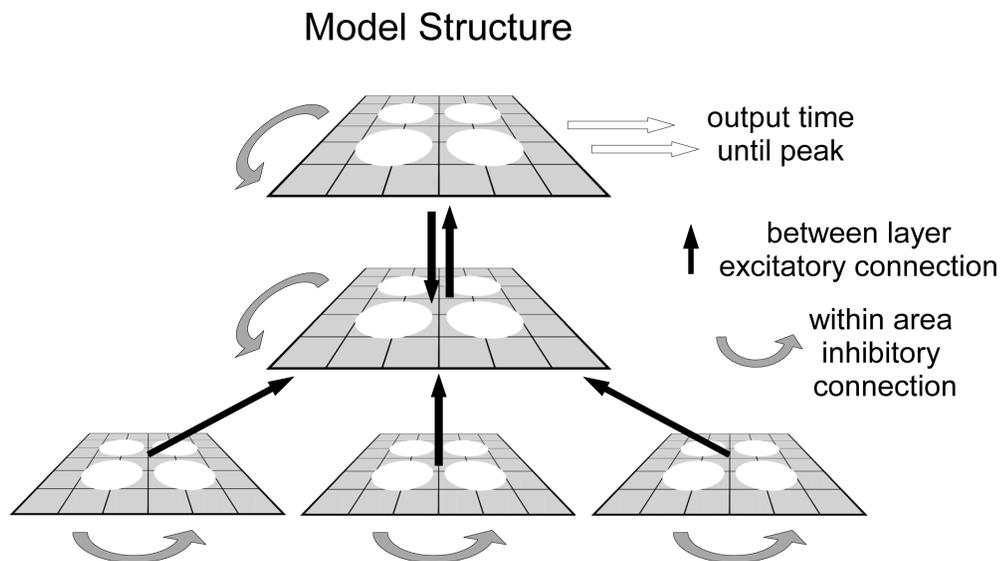
Model Structure

Figure 1.9. Model Structure. The model applied to the reported results had a 3 layer structure. The bottom layer is sensitive to spatial position, and so the same face presented in different locations (e.g., left, center, or right) activates different nodes. Each possible item in a trial (target, foil, and mask) is represented at this layer. At higher layers, connections from layer 1 converge such that the same face presented in different locations contacts the same nodes. A simple one to one mapping is used in connecting layers 2 and 3 and the use of two separate layers is designed to capture differences in the dynamics of activation, rather than making specific representational assumptions. Activation of each node is achieved through a real value (i.e., rate-coded rather than spiking). Within each area, nodes inhibit each other. The time to achieve maximum response in layer 3 is used to compute accuracy.

The model has a three layer structure as shown in Figure 1.9. Each node in the model can be viewed as describing the average firing rate of a large group of similarly connected neurons. In the bottom layer of the model there exists a simulated node for each possible presented visual object (including pattern masks) in each possible spatial location. The projections of these nodes to the next layer converge so that each node in the second layer codes for the face parts unique to each face, regardless of where the face is presented. Pattern masks, being a simple visual pattern, are only represented at the bottom layer. The model used a localist representation such that each node projects only to higher level nodes that code for the same face; a distributed representation was not necessary because we did not manipulate similarity. Thus, layers two and three were simply implemented through a single node for each face. Connection weights were set at fixed parameter values, and no learning was used in the simulations. The weights between the first and second layers of the model were fixed at one for connections between the same face at each level and zero otherwise. Connections between the second and third layers were likewise set to zero between different faces, but were set to one of two possible free parameter values for same face connections, one for upright faces and the other for inverted faces. In this manner, the strength of response for upright versus inverted faces was manipulated. Additionally, by allowing these two values to differ between the fast and slow groups, strength of response for individual differences was manipulated.

Neural Dynamics

The dynamics of each node use the same assumptions as found in O'Reilly and Munakata's (2000) LEABRA framework, but with the addition of synaptic depression.

Lateral inhibition was used within each layer of the model, simulating the effects of

inhibitory interneurons, which dampen excessive excitatory activity. The membrane

potential (i.e., activation) of each node is updated by Equation 1.1. The equations

provided here assume updates are performed every millisecond. More model details can

be found in Huber & O'Reilly (2003). Membrane potential (v_i^n) for node i in layer n is

increased by excitatory input weighted by the connection strength between sending (j)

and receiving (i) node ($w_{ij}o_i^{n-1}$) from the previous layer, from input, or from feedback

connections (F), if any, from nodes in layer $n+1$, indexed by k. Membrane potential is

decreased by within layer inhibition (I) and constant leak currents (L).

$$\frac{\Delta v_i^n}{S_n} = (1 - v_i^n)\left(\sum_{\forall j} w_{ij}o_j^{n-1} + \sum_{\forall k} F^n o_k^{n+1}\right) - (v_i^n)\left(I\sum_{\forall j} o_j^n + L\right) \quad (1.1)$$

$$o = \begin{cases} (v - \theta)a & v > \theta \\ 0 & v \le \theta \end{cases} \quad (1.2)$$

$$\frac{\Delta a}{S_n} = R(1 - a) - Do \quad (1.3)$$

Unlike traditional artificial neural networks, which only have an activation value

that varies with time, the nodes of this model contain two time varying parameters that

multiply to produce synaptic output as determined by presynaptic activity as well as by

the level of synaptic resources. In this manner, the model includes habituation due to

recent activity in the specific form of synaptic depression (Abbott et al. 1997; Tsodyks, &

Markram, 1997). The output (o) of each node is 0.0 until the membrane potential

increases past a threshold value (θ), after which it is scaled by a dynamically varying

factor (a) representing available resources for that particular connection (see Equation 1.2). In other words, the effect of a node is the product of two terms, with the first representing the on-average firing rate (v - θ), and the second representing the effect of each action potential in light of the available non-depleted resources (a). One example of synaptic resources is available presynaptic neurotransmitter—if a neuron has recently been highly active and has depleted its neurotransmitter supply, there will be little effect of ongoing action potentials. Equation 1.3 specifies the dynamic update for synaptic resources, and contains a constant for the recovery rate (R), and a constant for the rate of depression of activity (D). The dynamic update Equations 1.1 and 1.3 additionally contain a parameter, S_n, which regulates the speed of information integration for layer n, thus specifying the speed of processing at each layer of the model. The model was run in time steps of one millisecond. Input to a bottom layer node was set to 0.0 in general, but changed to 1.0 for the time steps when a visual object was presented that was the preferred input for that node.

Decision Rule

Because the behavioral task was threshold identification based on partial information, the decision rule in the model was to choose the test face identified more rapidly. As previously assumed in modeling threshold identification of words, we assumed that this perceptual fluency was used to assess the residual activation from the briefly flashed target face. This fluency measure was calculated as the "time to peak" response found separately for each choice face. In other words, if the target choice face reached its peak activation first, it was chosen, resulting in an accurate trial, but if the foil choice face reached peak activation first, an error occurred. Thus, the choice face

identified more quickly was chosen. This measure is analogous to a horserace model of forced choice, which was demonstrated to capture correct and error reaction time distributions with the word priming version of this task (Huber & Cousineau, 2004). Although we do not include it in the current model, the calculation of perceptual fluency could be implemented in simulated neurons through the difference between a fast and a slow integrator, with the latter inhibiting the activation of the former (Huber & O'Reilly, 2003).

Simulations were not run with any specific source of noise, and were thus a single deterministic run for each condition. In order to capture trial by trial variability, the simulated time to peak for the target and foil nodes were assumed to be normally distributed, and the deterministic simulation values provided the average time to peak for the target and foil in that particular condition. Assuming that variance is proportional to the mean, a noise constant (m) was used to calculate the variance of a condition based on the average time to peak. This assumption is appropriate if each millisecond of processing provides an independent normally distributed error term with mean zero and variance m. Accuracy then becomes a convolution of two normal distributions. More simply, this can be reformulated as a single normal distribution for the difference between the peak time for the target and the peak time for the foil, with the average difference found by subtracting the average peak time of target and foil, and the variance of the difference found by adding the variances of peak times for target and foil. Thus, probability correct is the inverse z-transformation (i.e., cumulative standard normal) based on the value $(F-T) / [m(T^2+F^2)]^{1/2}$ where T is the time to peak for the target choice face and F is the time to peak for the foil choice face in the condition of interest. This is a

slight, but not critical change from the method used in Huber and O'Reilly (2003). The

logistic function used by Huber and O'Reilly works equally well, although the current

use of separate distributions for target and foil is somewhat better motivated on

theoretical grounds. Because the parameter m only operates on the generated peak times,

it cannot influence the dynamics of the model, and only serves to transform model results

into the accuracy scale (see Appendix Figure A1.1 for example simulations

demonstrating the role of the free parameter m). Thus, allowing different values of m in

modeling the two groups of participants cannot explain the different habituation rates.

Modeling the Current Experiments

The fundamental test of the theory posed by these experiments was to examine

the relationship between the speed of face identification and the speed of habituation,

with the speed of face identification manipulated both by individual differences and by

face inversion. However, there are in fact two methods for manipulating the speed of face

identification to capture different degrees of face expertise. These correspond to changing

the connection weights between layers (w) or changing the processing speed within a

layer (S). Increasing either the connection strength or processing speed produces faster

activation for a briefly presented face (i.e., a shorter target duration is needed for

threshold performance) as well as a faster accrual of habituation, thus producing a more

rapid change from positive to negative priming. Therefore, we cannot distinguish

between these two explanations. However, if the speed of processing for a layer of

representation is related to an anatomical constraint, then it not obvious why this

parameter would vary between individuals. In keeping with a long tradition of

connectionist modeling (Rumelhart, & McClelland, 1986a; Rumelhart, & McClelland,

1986b; O'Reilly, & Munakata, 2000), we hypothesized that connection weights are modulated by experience. Therefore, we parameterized connection weights in different conditions to capture individual differences in face processing and to capture the difference between upright versus inverted faces.

To capture the differences between the dynamics of words and faces, we allowed several parameters to vary, although many were kept the same. Leak ($L = 0.15$), depletion ($D = 0.0324$), recovery ($R = 0.022$), and firing threshold ($q = 0.15$), are considered generic properties of all neurons and these were fixed to the same values for all simulated nodes as reported by Huber and O'Reilly (2003). Likewise, the speed of processing ($S_1 = 0.054$) and inhibition ($I_1 = 0.30$) for the visual input layer of the model is presumably the same type of visual response as for words (e.g., primary visual cortex). Finally, the same degree of feedback ($F = 0.25$) from layer three onto layer two was used to capture attractor dynamics. The three parameters that were truly free, and allowed to take on different values as compared to the previously published parameters used for word priming, were the speed of processing for layer two (S_2), and layer three (S_3), and the inhibition for layers two and three (I_{23}).

The within subject effect of inversion and the between subject difference between groups were modeled with the connection weight between the second and third layer (four values for w_{23}). Additionally, each group was allowed to have a different variance multiplier (two values for m), corresponding to different levels of variability for different individuals. As discussed above, this parameter does not affect the model dynamics, just the manner in which the model dynamics map into the accuracy scale. The model was fit

with chi-squared error as calculated by log likelihoods (e.g., Batchelder & Riefer, 1990),

using 9 free parameters to capture the 40 conditions, with 360 data points per condition.

Results and Discussion

Figure 1.10 shows the fit of the model to the data from Experiment 1.4. The

model clearly fits the qualitative pattern of the data, and is quantitatively accurate for

most conditions. The median chi-square per conditions was 1.01, and the model was not

statistically different from the data in 31 of the 40 conditions (according to a chi-squared

test with 31/40 degrees of freedom, considering the ratio of 9/40 of parameters per

condition). The best-fit parameters were as follows: $S_2 = 0.015$ (speed of integration for

the second layer), $S_3 = 0.022$ (speed of integration for the third layer), $I_{23} = 0.52$

(inhibition in the second and third layer), w_{23}-*fast-upright* = 2.05, w_{23}-*slow-upright*

=0.83, w_{23}-*fast-inverted* =0.73, w_{23}-*slow-inverted* = 0.50, m_{fast} = 0.21, and m_{slow} = 0.10.

A number of interesting observations can be made beyond demonstrating the

sufficiency of the model in capturing these data. Indeed, in descriptive modeling, such as

an application of signal detection theory (Egan, 1975), the goal is not just a good fit to

data, but additionally to transform data into theoretically meaningful parameter values.

Looking at the best-fit parameters, it is notable that the processing speed of the third layer

is slightly faster (larger numerically) then that of the second layer (the word priming

values were $S_2 = 0.046$, $S_3 = 0.015$). The overall differences in speed when comparing

faces and words does not necessarily indicate anything important about these classes of

visual objects considering that the types of displays (e.g., retinal size) and masks were

vastly different. However, that the third layer runs faster than the second layer for faces

but not for words suggests that higher level face information is processed quickly relative
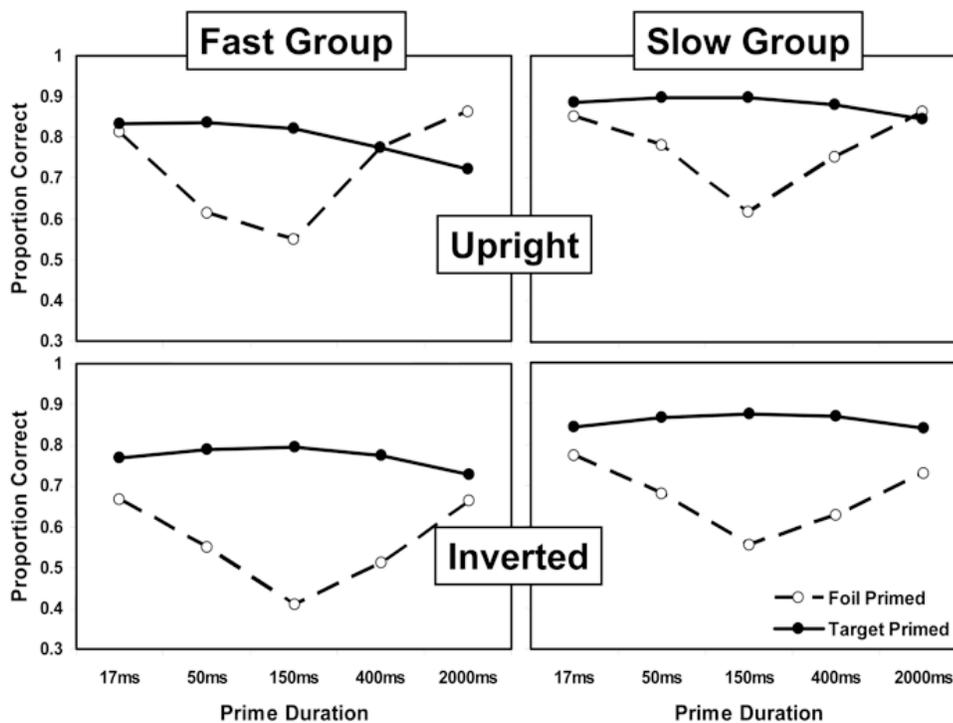
Figure 1.10. Experiment 1.4 model results from best-fitting parameters. The connection weights between layers 2 and 3 were allowed to vary for each of the four conditions. The performance gain parameter, m, was allowed to vary between the fast and slow groups. All other parameters were the same in all conditions and many were set to default constants suggested by Huber & O'Reilly (2003).

to lower level face information. This is sensible if face wholes are based on spatial configuration without needing to wait for full identification of the particular face features. In contrast, word identification critically hinges upon the letters that comprise the word. It is also interesting to note that inhibition for face processing is high relative to words (all I values were 0.3 for words). Anecdotally this makes sense. We commonly need to maintain several words at once in order to extract meaning over a sentence, but it might be confusing to simultaneously identify and maintain several faces (i.e., we usually only attend to one face at a time).

Despite the apparent behavioral trends, both the fast and slow groups reveal a sizable effect of inversion for the connection weights. In other words, application of the

model places a cautionary note on the apparent conclusion from the behavioral data that only the fast group produced an inversion effect. The observed pattern of data might result from greater sensitivity in the behavioral measure for the fast group; the very small inversion effects for the slow group may in fact represent sizable changes to the underlying representation. Sensibly, the connection strengths for the fast group were overall greater than the corresponding values for the slow group, corresponding to the fast group's ability to more rapidly identify the target face and the fast group's faster transition from positive to negative priming with upright faces. Interestingly, connection strength differences between groups were actually larger in magnitude than the inversion effects within groups, suggesting that these individual differences in face processing may be something of importance and are perhaps more sizable than face inversion.

General Discussion

Summary of results

Four behavioral studies examined immediate face repetition priming and confirmed predictions of the claim that habituation temporally segregates streams of perceptual events. According to this theory, temporal integration between recently presented stimuli (i.e., primes) and currently presented stimuli (i.e., targets) produces positive priming, as revealed by accuracy benefits when priming the target but deficits when priming the incorrect choice (i.e., the foil). However, this positive priming is offset and possibly reversed by habituation to primes presented for longer durations. This theory predicted that the speed of identification for a perceptual object should relate to the speed of the transition from positive to negative priming, and this prediction was

confirmed both through individual differences in face identification and by manipulating face orientation.

Experiment 1.1 used a paradigm that previously established the positive to negative priming transition with words, and found similar priming effects with computer generated face images. However, unlike word priming, there were reliable individual differences found in all four experiments. Participants who could more rapidly identify the target face (as determined by threshold target duration) produced a faster transition from positive to negative priming. Experiment 1.2 demonstrated that this was not due to the use of different target durations for the fast and slow median split groups of participants. Experiment 1.3 included a condition where neither the target nor the foil face was primed and tested photographs of real faces cropped to the interior of the face. This verified equal costs and benefits of priming relative to a neutral baseline. Experiment 1.4 found that upright faces produced a faster transition to negative priming as compared to inverted faces. This inversion effect was particularly prominent for the fast group of participants and seemingly absent for the slow group. Experiment 1.5 was a simulation study that applied a dynamic neural network model with habituation to the individual difference and inversion effect data from Experiment 1.4. This provided an accurate account of the data, demonstrating that these qualitative predictions were also quantitatively in accord with neural habituation. The best-fitting connection strength parameters revealed sizable effects of face inversion for both the fast and slow groups, which suggests that the apparent lack of inversion effect with the slow group was a floor effect for the ability of habituation to produce negative priming.

Relationship to other paradigms and measures

This habituation theory may explain a wide variety of paradigms that produce

transitory deficits or changes in neural activation after prolonged exposure to a stimulus.

For instance, similar phenomena occur with studies of perception, lexical processing,

attention, evaluation, and episodic recognition (see Huber, 2008b for a review).

Behavioral, ERP, and MEG experiments have confirmed predictions from the claim that

habituation between perceptual and lexical-semantic processing underlies immediate

repetition priming deficits with words (Huber, Tian, Curran, O'Reilly, & Woroch, 2008).

The idea that habituation occurs due to the loss of association between levels of

representation has also been applied to semantic satiation, which is the feeling that a

word looses meaning when repeated many times. Experiments ruled out explanations

based on lexical or semantic fatigue, suggesting that it is the inability to access meaning

from the repeated word form that explains the phenomenon (Tian & Huber, 2010).

Demonstrating the generality of this theory, something similar to our results was found in

the phenomenon of 'repetition blindness', in which an observer fails to report the second

occurrence of a word in a sentence presented in rapid serial order (Kanwisher, 1987).

This has been explained as a failure to separate each word as a separate word token (i.e.,

the first occurrence versus the second occurrence) despite adequate identification of the

word type (i.e., evidence that a particular word was seen recently, regardless of where

and when). Habituation may be the underlying cause of this failure to produce

identifiably separate responses to repeated words. Similar to our face inversion results,

Coltheart and Langdon (2003) produced the usual repetition blindness effect with known

words, but in the same experiment, repetitions of non-words instead produced a benefit

(i.e., more likely to report a second occurrence of a non-word). These results are explained by habituation because non-words are unfamiliar and habituate more slowly than words; because the first occurrence of an unfamiliar non-word produces little or no habituation, it only serves to provide lingering activation that helps identify the second occurrence (i.e., positive priming).

This theory also may also provide insight into the nature of repetition suppression, which is the attenuation of the neural response to a repeated stimulus as measured with single cell recording, fMRI, or EEG/MEG (Grill-Spector, Henson & Martin, 2006; Ranganath & Rainer, 2003). An account of these effects due to habituation depends on the time scale over which the repetitions occur. Behavioral facilitation is typically seen for a target that is an immediate repetition of a brief or subliminal prime when there is no delay between prime and target. In contrast, longer prime durations, multiple prime presentations, or the inclusion of a very short delay (seconds or less) between prime and target often produce a behavioral and neural deficit, such as with repetition blindness. The model used in Experiment 1.5 naturally produces behavioral facilitation with subliminal no-delay priming due to temporal integration between prime and target. The model also produces reduced neural responses following longer duration primes (e.g., when habituation more than offsets the benefits of temporal integration) or when there is a delay between prime and target. With a short delay between prime and target, residual prime activation fades (i.e., there is no temporal integration and thus no facilitation). However, recovery from habituation operates on a slower time scale, and so the neural activation to the target is weak due to lingering habituation. On a longer time scale, with seconds or longer between prime and target, behavioral benefits are almost always found.

These longer time scale benefits are presumably due to learning, which is not included in the current model. Habituation is a temporary effect, and with longer delays, the neural response recovers, thus unmasking the behavioral benefits of learning. Although the neural response is temporarily placed in a habituated state through repetitions, the brain is presumably learning the repeated stimulus, which supports a behavioral benefit after a longer delay. Furthermore, a well learned visual object (e.g., an upright face rather than an inverted face) is expected to produce a faster transition to habituation, and so a measure with poor temporal resolution, such as fMRI, is expected to produce less activation across an averaged time window. The prediction for well learned visual objects as measured with high temporal resolution techniques (e.g., scalp EEG) are more complicated, as discussed next.

The explanation of the reported face inversion effects assumed that upright faces have an initially stronger response than inverted faces, which then falls more rapidly to a habituated state. However, this assumption seems at odds with the finding that inverted faces produce larger N170 ERPs than upright faces (Rossion et al, 1999). However, this finding is only problematic if more negative N170 responses are assumed to indicate more neural activation for the neurons responsible for face identification. However, source localization of ERP responses is complicated not only by anatomical orientation (e.g., there is often a larger P170 as measured at frontal regions concurrent with the reduced N170) but also because multiple cortical sources may be simultaneously active (e.g., non-face identification processes may contribute to the N170). Therefore, an ERP difference between conditions cannot be uniquely attributed to a particular source, and, furthermore, the direction of the difference does not necessarily indicate a particular

direction of change in the underlying neural activity. Because there is no unique solution to the so-called 'inverse problem' of electrophysiology (Mosher, Baillet, & Leahy, 1999), all that can be concluded is that the summed neural activity for upright faces is different than that of inverted faces at ~170 ms after presentation. Thus, in the absence of source modeling, the ERP N170 face inversion results are equally compatible with the claim that the neurons most directly responsible for face identification are initially more active for upright faces as compared to inverted faces.

Implications of the reported individual differences

Perhaps the most well studied individual difference with faces is the deficit of prosopagnosia, which is characterized by a selective inability to identify faces as compared to other visual objects (Duchanie, Yovel, Butterworth, Nakayama, 2006). Demonstrating the opposite end of the face expertise continuum, recent investigations have uncovered individuals who are extremely skilled at face identification, to the point of needing to hide this ability in social situations where it might seem inappropriate to easily identify someone seen just fleetingly many years beforehand (Russell, Duchaine, & Nakayama, 2009). In combination with our results with normal individuals, these clinical deficits and enhancements suggest there may be larger variations in face perception ability than has been previously appreciated. For instance, application of the habituation model to our Experiment 1.4 data revealed a larger effect for the median split of individuals as compared to the effect of face inversion. However, it has been reported that use of computer generated faces may tend to reduce inversion effects as compared to photographs of real faces (Carlson, & Gronlund, 2007). Providing some closure on this issue, we recently replicated the inversion and individual difference results of Experiment

1.4 using face photographs similar to those used in Experiment 1.3, again finding that individual difference effects were larger than inversion effects. For reasons of space we did not currently report these results.

The reliability of these individual differences demonstrates that this immediate priming paradigm may be useful in clinical diagnoses and treatment of perceptual processing disorders. For instance, this forced choice priming paradigm has been used with words, revealing that the extent of language disorder in schizophrenics relates to the magnitude of semantic priming benefits (Quelen, Grainger, and Raymondet, 2005). This was found for the both-primed condition (not currently used), which, according to Huber (2008b), reflects the strength of top-down facilitation as compared to bottom-up support. The current application of the priming paradigm with multiple prime durations demonstrates the predicted link between speed of identification for a stimulus and the rate at which rapid presentations can be handled with minimal source confusion. However, the reported individual differences are correlative in nature, and a training study is needed. Providing an example of such a training study, it has been found that training on sound identification increases reading speed (Merzenich, Jenkins, Johnston, Schreiner, Miller, & Tallal, 1996). Finally, our paradigm for identifying and understanding individual differences may be useful in the study of autism-spectrum disorder (ASD). Although much research on ASD has focused on a lack of social motivation, perceptual processing in general, and face processing in particular, is also disrupted (Deruelle, Rondan, Gepner & Tardif, 2004; Dawson, Webb & McPartland, 2005; McCleery, Allman, Carver & Dobkins, 2007). Variants of our priming paradigm might, for instance,

identify whether ASD face processing deficits are in part due to difficulty segregating different views of a face.

Conclusions

Four experiments demonstrated that immediate face priming produces positive priming for short prime durations but negative or no priming for longer target durations. This effect was previously shown with immediate word priming and this generalization to faces was expected by a theory of temporal segregation through habituation. According to this theory, the temporal integration between short duration prime faces and target faces produces source confusion (see also, Anaki, Boyd, & Moscovitch, 2007). However, habituation following long duration prime faces more than offsets this effect, producing negative priming. More importantly, the reported results tested the prediction that more rapid face identification should produce more rapid habituation and, therefore, more quickly produce negative priming as a function of increasing prime duration. This prediction was confirmed both with individual differences relating face detection threshold to priming and also by comparing priming for upright versus inverted faces. The observed differences in face perception were large and found in all experiments. Beyond verifying key predictions of this habituation account, these results suggest a new technique for examining individual differences in face processing.

Chapter 1, in full and pending permission from the copyright holder, APA, is adapted from published material Copyright © 2010 by the American Psychological Association. Adapted with permission. The official citation that should be used in referencing this material is Rieth, C. A., & Huber, D. E. (2010). Priming and habituation for faces: Individual differences and inversion effects. Journal of Experimental

Footnotes

[1] We use the phrase 'negative priming' to refer to a situation in which performance is worse for primed trials as compared to unprimed trials, and this should not be confused with the specific paradigm of *negative priming* (Tipper, 1985), in which priming deficits are found as a function of a prime's status as a to-be-ignored distractor.

[2] This literature often refers to these effects as face adaptation rather than habituation. However, because there may be a tendency to confuse the term adaptation with long-term learning and memory, and because adaptation does not necessarily imply a negative effect, we instead adopt the term habituation.

[3] Performing a regression analysis on continuously varying measures of individual differences is statistically preferred to median split analyses (Irwin & McClelland, 2003). However, the current situation only included 5 possible values for the measure of individual differences. In any case, reliability of these individual differences is additionally assessed through replication across all 4 experiments, and our use of a median split is used to illustrate the qualitative difference between ends of the face identification spectrum.
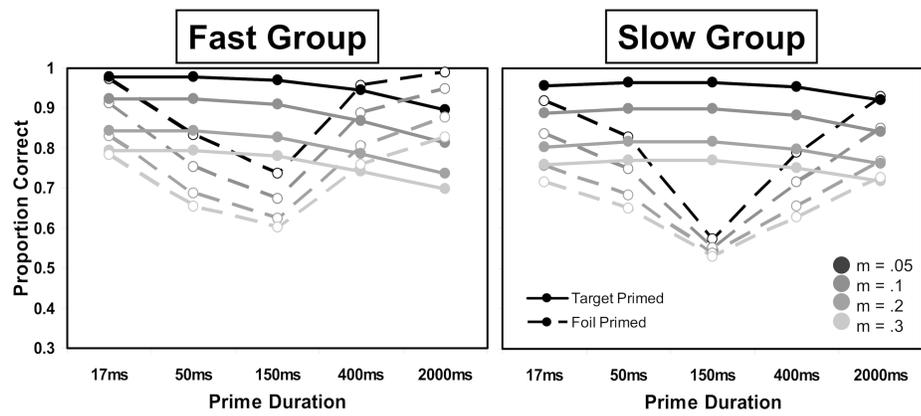
Figure A1.1. Varying model variance multiplier. Variations in parameter *m* while holding other parameters to the best-fit values influence overall accuracy of the model, but do not influence priming dynamics.

Appendix

To facilitate future modeling research, we provide mean accuracy and standard errors (in parentheses) for all conditions in all experiments collapsed across participants, as well as broken down by fast/slow group.

Table A1.1.

Experiment 1.1, Neither Primed Results

Target Duration	Fast Group	Slow Group	All
33ms	.66 (.03)	.48 (.03)	.57 (.03)
50ms	.86 (.02)	.53 (.03)	.69 (.04)
67ms	.90 (.02)	.66 (.05)	.78 (.03)
83ms	.94 (.02)	.78 (.04)	.86 (.03)
N	14	14	28

Table A1.2.

Experiment 1.1 Priming Results

Prime Duration	Fast Group	Slow Group	All
		Foil Primed	
17ms	.67 (.04)	.82 (.02)	.75 (.03)
50ms	.43 (.03)	.68 (.03)	.56 (.03)
150ms	.48 (.03)	.56 (.05)	.52 (.03)
400ms	.74 (.04)	.70 (.08)	.72 (.04)
2,000ms	.86 (.03)	.75 (.08)	.81 (.04)
		Target Primed	
17ms	.75 (.05)	.84 (.02)	.80 (.03)
50ms	.73 (.05)	.88 (.02)	.80 (.03)
150ms	.76 (.05)	.85 (.02)	.81 (.03)
400ms	.73 (.05)	.86 (.03)	.79 (.03)
2,000ms	.66 (.06)	.87 (.02)	.76 (.04)
N	14	14	28

Table A1.3.

Experiment 1.2, 33 ms Target Duration Results

Prime

Duration	Fast Group	Slow Group	All
		Foil Primed	
17ms	.59 (.03)	.50 (.02)	.54 (.02)
150ms	.48 (.04)	.30 (.03)	.39 (.03)
2,000ms	.65 (.04)	.41 (.04)	.53 (.03)
		Target Primed	
17ms	.84 (.07)	.73 (.02)	.78 (.02)
150ms	.85 (.10)	.77 (.03)	.81 (.02)
2,000ms	.80 (.12)	.75 (.03)	.78 (.02)
		Neither Primed	
150ms	.67 (.02)	.54 (.01)	.60 (.01)
N	28	28	56

Table A1.4.

Experiment 1.2, 50 ms Target Duration Results

Prime

Duration	Fast Group	Slow Group	All
		Foil Primed	
17ms	.75 (.03)	.58 (.03)	.66 (.02)
150ms	.66 (.04)	.40 (.03)	.52 (.03)
2,000ms	.86 (.02)	.56 (.05)	.70 (.03)
		Target Primed	
17ms	.91 (.01)	.75 (.03)	.83 (.02)
150ms	.89 (.02)	.78 (.03)	.84 (.02)
2,000ms	.86 (.02)	.77 (.03)	.81 (.02)
		Neither Primed	
150ms	.80 (.02)	.58 (.02)	.69 (.02)
N	28	28	56

Table A1.5.

Experiment 1.2, 100 ms Target Duration Results

Prime Duration	Fast Group	Slow Group	All
		Foil Primed	
17ms	.93 (.01)	.76 (.03)	.84 (.02)
150ms	.91 (.02)	.62 (.04)	.76 (.03)
2,000ms	.96 (.01)	.72 (.05)	.83 (.03)
		Target Primed	
17ms	.96 (.01)	.83 (.02)	.89 (.02)
150ms	.95 (.01)	.80 (.03)	.87 (.02)
2,000ms	.95 (.01)	.84 (.02)	.89 (.01)
		Neither Primed	
150ms	.94 (.01)	.74 (.03)	.84 (.02)
N	28	28	56

Table A1.6.

Experiment 1.3, Neither Primed Results

Target Duration	Fast Group	Slow Group	All
33ms	.67 (.02)	.54 (.02)	.61 (.01)
50ms	.81 (.01)	.57 (.01)	.69 (.02)
67ms	.84 (.02)	.71 (.02)	.78 (.01)
83ms	.90 (.01)	.81 (.02)	.85 (.01)
N	51	51	102

Table A1.7.

<u>Experiment 1.3 Priming Results</u>

Prime Duration	Fast Group	Slow Group	All
		Foil Primed	
17ms	.68 (.02)	.75 (.02)	.75 (.01)
50ms	.52 (.02)	.64 (.02)	.64 (.02)
150ms	.42 (.02)	.49 (.02)	.49 (.02)
400ms	.63 (.02)	.70 (.02)	.70 (.02)
2,000ms	.75 (.03)	.76 (.03)	.76 (.02)
		Target Primed	
17ms	.81 (.02)	.82 (.02)	.82 (.01)
50ms	.83 (.02)	.85 (.02)	.85 (.01)
150ms	.82 (.02)	.86 (.02)	.86 (.01)
400ms	.78 (.02)	.85 (.02)	.85 (.01)
2,000ms	.73 (.02)	.83 (.02)	.83 (.02)
		Neither Primed	
17ms	.74 (.02)	.79 (.02)	.76 (.01)
50ms	.66 (.02)	.73 (.02)	.69 (.01)
150ms	.64 (.02)	.71 (.02)	.68 (.01)
400ms	.71 (.02)	.76 (.02)	.73 (.01)
2,000ms	.75 (.02)	.76 (.02)	.76 (.01)
N	51	51	102

Table A1.8.

Experiment 1.4, Neither Primed Results

Target

Duration	Fast Group	Slow Group	All
	Upright Faces		
33ms	.54 (.03)	.52 (.03)	.53 (.02)
50ms	.74 (.04)	.61 (.04)	.67 (.03)
67ms	.88 (.03)	.59 (.04)	.73 (.03)
83ms	.90 (.03)	.71 (.03)	.80 (.03)
	Inverted Faces		
33ms	.56 (.04)	.48 (.04)	.52 (.03)
50ms	.61 (.04)	.49 (.03)	.55 (.03)
67ms	.73 (.04)	.51 (.04)	.62 (.03)
83ms	.83 (.04)	.69 (.04)	.76 (.03)
N	20	20	40

Table A1.9.

Experiment 1.4, Upright Face Results

Prime

Duration	Fast Group	Slow Group	All
		Foil Primed	
17ms	.74 (.04)	.86 (.03)	.80 (.02)
50ms	.54 (.02)	.78 (.02)	.66 (.03)
150ms	.51 (.04)	.63 (.04)	.57 (.03)
400ms	.80 (.03)	.74 (.04)	.77 (.03)
2,000ms	.91 (.02)	.85 (.05)	.88 (.03)
		Target Primed	
17ms	.83 (.03)	.88 (.02)	.86 (.02)
50ms	.87 (.03)	.89 (.02)	.88 (.02)
150ms	.80 (.03)	.88 (.02)	.84 (.02)
400ms	.80 (.04)	.93 (.02)	.87 (.02)
2,000ms	.76 (.04)	.88 (.03)	.82 (.02)

Table A1.10.

Experiment 1.4, Inverted Face Results

Prime

Duration	Fast Group	Slow Group	All
		Foil Primed	
17ms	.66 (.03)	.78 (.02)	.72 (.02)
50ms	.47 (.04)	.67 (.03)	.57 (.03)
150ms	.39 (.03)	.50 (.04)	.45 (.03)
400ms	.58 (.04)	.66 (.05)	.62 (.03)
2,000ms	.67 (.04)	.74 (.05)	.70 (.03)
		Target Primed	
17ms	.77 (.03)	.83 (.03)	.80 (.02)
50ms	.77 (.03)	.83 (.03)	.80 (.02)
150ms	.77 (.04)	.84 (.02)	.80 (.02)
400ms	.76 (.04)	.88 (.03)	.82 (.03)
2,000ms	.70 (.03)	.82 (.03)	.76 (.02)
N	20	20	40

References

Abbott, L. F., Varela, J. A., Sen, K., & Nelson, S. B. (1997). Synaptic depression and cortical gain control. *Science, 275* (5297), 220-224.

Anaki D., Boyd J., & Moscovitch M. (2007). Temporal integration in face perception: evidence of configural processing of temporally separated face parts. *Journal of Experimental Psychology: Human Perception and Performance, 33(1),* 1-19.

Batchelder, W. H., & Riefer, D. M. (1990). Multinomial processing models of source monitoring. *Psychological Review*, *97(4)*, 548-564.

Bennett, J. D., Lleras, A., Oriet, C., & Enns, J. T. (2007). A negative compatibility effect in priming of emotional faces. *Psychonomic Bulletin & Review*, *14(5),* 908-912.

Carlson, C. A., & Gronlund, S. D. (2007, November). *Computer-generated faces are not processed like real faces.* Presented at Psychonomics, Long Beach, CA.

Carpenter, G. A., & Grossberg, S. (2003), Adaptive resonance theory, In M.A. Arbib (Ed.), *The Handbook of Brain Theory and Neural Networks, Second Edition* (pp. 87-90). Cambridge, MA: MIT Press.

Coltheart, V., & Langdon, R. (2003). Repetition blindness for words yet repetition advantage for nonwords. *Journal of Experimental Psychology: Learning, Memory and Cognition, 29(2)*, 171-185.

Dawson, G., Webb, S. J., & McPartland, J. (2005). Understanding the nature of face processing impairment in autism: insights from behavioral and electrophysiological studies. *Developmental Neuropsychology. 27(3),* 403-424.

Deruelle, C., Rondan, C., Gepner, B., & Tardif, C. (2004). Spatial frequency and face processing in children with autism and asperger syndrome. *Journal of Autism and Developmental Disorders, 34(2),* 199-209.

Duchanie, B. C., Yovel, G., Butterworth, E.J., & Nakayama, K. (2006). Prosopagnosia as an impairment to face-specific mechanisms: Elimination of the alternative hypothesis in a developmental case. *Cognitive Neurophychology, 23(5),* 714-747.

Egan, J. P. (1975). *Signal Detection Theory and ROC Analysis*. New York: Academic Press.

Farah, M. J., Tanaka, J. W., & Drain, H. M. (1995). What causes the face inversion effect? *Journal of Experimental Psychology: Human Perception and Performance, 21(3),* 628-634.

Farah, M. J., Wilson, K. D., Drain, H. M., & Tanaka, J. R. (1998). What is 'special' about face perception? *Psychological Review*, *105(3)*, 482-498.

Fox, C. J., & Barton, J. J. S. (2007). What is adapted in face adaptation? The neural representations of expression in the human visual system. *Brain Research, 1127(1),* 80-89.

Freire, A. Lee, K., & Symons, L. A. (2000). The face-inversion effect as a deficit in the encoding of configural information: Direct evidence. *Perception, 29(2),* 159-170.

Gauthier, I., & Logothetis, N. K. (2000). Is face recognition not so unique after all? *Cognitive Neuropsychology, 17(1-3),* 125-142.

Gauthier, I., & Nelson, C. A. (2001). The development of face expertise. *Current Opinion in Neurobiolog*y, *11(2),* 219-24.

Gauthier, I., Tarr, M. J., Anderson, A. W., Skudlarski, P., & Gore, J. C. (1999). Activation of the middle fusiform 'face area' increases with expertise in recognizing novel objects. *Nature Neuroscience, 2,* 568-573.

Grill-Spector, K., Henson, R., & Martin, A. (2006). Repetition and the brain: neural models of stimulus-specific effects. *Trends in Cognitive Sciences*, *10(1),* 14–23.

Huber, D. E. (2008a). Causality in time: Explaining away the future and the past. In M. Oaksford and N. Chater (Eds.). *The probabilistic mind: Prospects for rational models of cognition.* Oxford University Press. 351-376.

Huber, D. E. (2008b). Immediate priming and cognitive aftereffects. *Journal of Experimental Psychology: General. 137,* 324-347.

Huber, D. E., & Cousineau, D. (2004). A race model of perceptual forced choice reaction time. R. Alterman & D. Kirsh (Eds.), P*roceedings of the 25th Annual Conference of the Cognitive Science Society* (pp. 687-692). Hillsdale, NJ: Erlbaum.

Huber, D. E., & O'Reilly, R. C. (2003). Persistence and accommodation in short-term priming and other perceptual paradigms: Temporal segregation through synaptic depression. *Cognitive Science, 27(3),* 403-430.

Huber, D. E., Shiffrin, R. M., Lyle, K. B., & Quach, R. (2002). Mechanisms of source confusion and discounting in short term priming 2: Effects of prime similarity and target duration. *Journal of Experimental Psychology: Learning Memory and Cognition, 28,* 1120-1136.

Huber, D. E., Shiffrin, R. M., Lyle, K. B., & Ruys, K. I. (2001). Perception and preference in short-term word priming. *Psychological Review, 108(1)*, 149-182.

Huber, D. E., Tian, X., Curran, T., O'Reilly, C., & Woroch, B. (2008). The dynamics of integration and separation: ERP, MEG, and neural network studies of immediate repetition effects. *Journal of Experimental Psychology: Human Perception and Performance. 34(6),* 1389-1416.

Hummel, J. E., & Biederman, I. (1992). Dynamic binding in a neural network for shape recognition. *Psychological Review, 99(3)*, 480-517.

Irwin, J. R., & McClelland, G. H. (2003). Negative consequences of dichotomizing continuous predictor variables. *Journal of Marketing Research (JMR)*, *40(3),* 366-371.

Kanwisher, N. G. (1987). Repetition blindness: Type recognition without token individuation. *Cognition*, *27(2),* 117-143.

Kanwisher, N. (2000). Domain specificity in face perception. *Nature Neuroscience, 3(8),* 759-763.

Leder, H., & Bruce, V. (2000). When inverted faces are recognized: The role of configural information in face recognition. *The Quarterly Journal of Experimental Psychology Section A, 53(2)*, 513-536.

Leopold, D. A., O'Toole A. J., Vetter T., & Blanz V. (2001). Prototype-referenced shape encoding revealed by high-level aftereffects. *Nature Neuroscience 4(1),* 89-94.

Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. New York: Freeman.

Maurer, D., Le Grand, R., & Mondlock, C. J. (2002). The many faces of configural processing. *Trends in Cognitive Sciences, 6(6),* 255-260.

McCleery, J. P., Allman, E., Carver, L. J., & Dobkins, K. R. (2007). Abnormal magnocellular pathway visual processing in infants at risk for Autism. *Biological Psychiatry*, *62(9),* 1007-1014.

McKone, E., Kanwisher, N., & Duchaine, B. C. (2006). Can generic expertise explain special processing for faces? *Trends in Cognitive Sciences, 11(1),* 8-15.

Merzenich, M. M., Jenkins, W. M., Johnston, P., Schreiner, C., Miller, S. L., & Tallal, P. (1996) Temporal processing deficits of language-learning impaired children ameliorated by training. *Science, 271(5245)*, 77-81.

Meyer, D. E., & Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: evidence of a dependence between retrieval operations. *Journal of Experimental Psychology, 90(2),* 227-234.

Meyer, D. E., Schvaneveldt, R. W., & Ruddy, M. G. (1975). In P. M. A. Rabbitt (Ed.), *Attention and performance V* (pp. 98-118). London: Academic Press.

Mosher, J. C., Baillet, S., & Leahy, R. M. (1999). EEG source localization and imaging using multiple signal classification approaches. *Journal of Clinical Neurophysiology, 16(3),* 225-238.

O'Reilly, R. C., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience.* Cambridge, Massachusetts: The MIT Press.

Quelen, F., Grainger, J., & Raymondet, P. (2005). An investigation of semantic priming in schizophrenia using a new priming paradigm. *Schizophrenia Research, 80(2-3),* 173-183.

Rakover, S. S. (2002). Featural vs. configurational information in faces: A conceptual and empirical analysis. *British Journal of Psychology*, *93(1),* 1-30.

Ranganath, C., & Rainer, G. (2003). Neural mechanisms for detecting and remembering novel events. *Nature Reviews Neuroscience, 4(3),* 193-202.

Ratcliff, R., & McKoon, G. (1997). A counter model for implicit priming in perceptual word identification. *Psychological Review, 104(2),* 319-343.

Rhodes, G., Jeffery, L., Watson, T. L, Clifford, C. W. G., & Nakayama, K. (2003). Fitting the mind to the world: face adaptation and attractiveness aftereffects. *Psychological Science, 14(6),* 558-566.

Richler, J.J., Tanaka, J.W., Brown, D.D. & Gauthier, I. (2008). Why does selective attention to parts fail in face processing? *Journal of Experimental Psychology: Learning, Memory and Cognition, 34,* 1356-1368.

Rossion, B., Delvenne, J. F., Debatisse, D., Goffaux, V., Bruyer, R., Crommelinck, M., et al. (1999). Spatio-temporal localization of the face inversion effect: An event-related potentials study. *Biological Psychology*, *50(3),* 173-189.

Rossion, B., & Gauthier, I. (2002). How does the brain process upright and inverted faces? *Behavioral and Cognitive Neuroscience Reviews*, *1(1),* 63-75.

Rotshtein, P., Henson, R. N. A., Treves, A., Driver, J., & Dolan, R. J. (2005). Morphing Marilyn into Maggie dissociates physical and identity face representations in the brain. *Nature Neuroscience, 8(1),* 107-113.

Rumelhart, D. E., & McClelland, J. L. (1986a). *Parallel distributed processing: explorations in the microstructure of cognition, vol. 1: foundations*. MIT Press Cambridge, MA, USA.

Rumelhart, D. E., & McClelland, J. L. (1986b). *Parallel distributed processing: explorations in the microstructure of cognition, vol. 2: psychological and biological models*. . MIT Press Cambridge, MA, USA.

Russell, R., Duchaine, B., Nakayama, K. (2009). Super-recognizers: People with extraordinary face recognition ability. *Psychonomic Bulletin & Review. 16(2),* 252-257.

Singer, W. (1999). Neuronal synchrony: A versatile code review for the definition of relations? *Neuron, 24,* 49-65.

Tanaka, J.W. A race face database: Standardized set of over 200 Caucasian, African American, Asian and Hispanic faces. Department of Psychology, University of Victoria, Victoria, BC, Canada.

Tanaka, J. W., & Pierce, L. J. (2009). The neural plasticity of other-race face recognition. *Cognitive, Affective, & Behavioral Neuroscience, 9,* 122-131.

Thomson, A., M., & West, D., C. (1993). Fluctuations in pyramid-pyramid excitatory postsynaptic potentials modified by presynaptic firing pattern and postsynaptic membrane potential using paired intracellular recordings in rat neocortex, *Neuroscience, 54(2),* 329-346.

Tian, X. & Huber, D. E. (2010). Testing an associative account of semantic satiation. *Cognitive Psychology. 60.* 267-290.

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.

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology, 12(1),* 97-136.

Tsodyks, M. V., & Markram, H. (1997). The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proceedings of the National Academy of Sciences, 94(2),* 719-723.

Weidemann, C. T., Huber, D. E., & Shiffrin, R. M. (2005). Confusion and compensation in visual perception: Effects of spatiotemporal proximity and selective attention. *Journal of Experimental Psychology: Human Perception and Performance, 31(1),* 40-61.

Weidemann, C. T., Huber, D. E., & Shiffrin, R. M. (2008). Prime diagnosticity in short-term repetition priming: Is primed evidence discounted, even when it reliably indicates the correct answer? *Journal of Experimental Psychology: Learning Memory and Cognition, 34 (2).* 257-281.

Wenger, M. J., & Ingvalson, E. M. (2002). A decisional component of holistic encoding. *Journal of Experimental Psychology Learning Memory and Cognition, 28(5),* 872-892.

Wenger, M. J., & Ingvalson, E. M. (2003). Preserving informational separability and violating decisional separability in facial perception and recognition. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, *29*(6), 1106-1118.

Yin, R. K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology*, *81,* 141-145.

CHAPTER 2: The temporal dynamics of occurrence and co-occurrence in written word identification

As we experience each moment, we effortlessly form coherent perceptions from dynamic sensory input. Our mind interprets a highly multidimensional sensory stream quickly and seemingly accurately, segregating perceptions not only in space, but also in time. Wildly different sensory signals are caused by the same source. For example, a street sign will elicit varying sensory signals as you approach it. However, these signals are all attributable to the same source. The assignment of sensory activation to either novel causes in the world or to previous events is critical to parsing the sensory stream. Doing so efficiently necessitates using statistical properties of the world to guide perception. If you are not quite able to read a street sign as you pass it, prior knowledge about the occurrence and co-occurrence of words on street signs helps you to identify their contents. For example, if the second word on a street sign started with 'ro' and was about four letters long, it is likely to be 'road'. This paper explores the influence of occurrence and co-occurrence statistics on the dynamics of word identification using immediate priming to elicit temporal confusion.

Immediate word priming paradigms (Huber, 2008; Huber, Shiffrin, Lyle, & Quach, 2002; Huber, Shiffrin, Lyle, & Ruys, 2001; Norris & Kinoshita, 2008) provide an avenue to study how word occurrence and co-occurrence statistics influence temporal segregation. In these paradigms, rapid presentation of primes and targets create confusion about the identity of a target. Manipulating the prime's salience through prime duration can result in either assimilation or segregation of the prime features from the target features (Huber & O'Reilly, 2003; Huber, Shiffrin, Quach, & Lyle, 2002). These results

spawned a theory of habituation as a mechanism for the identification of rapidly presented stimuli (Huber, 2008; Huber & O'Reilly, 2003). Two complementary models explain the process of identification under temporal uncertainty at different levels of description (Marr, 1982), the probabilistic Responding Optimally with Unknown Sources of Evidence (ROUSE) model (Huber et al., 2001), and the artificial neural network model, nROUSE (Huber & O'Reilly, 2003). The nROUSE model posits that perceptual representations habituate through activity-dependent depletion of resources. A key component of this model is that representations saturate after extended activation, allowing novel representations to drive perception when supported by input. If input changes quickly, representations will not habituate, and perceptual inputs from different sources will, to an extent, blend. However, if input changes more slowly, habituation may in some cases overreact, resulting in a bias to report change, when in fact no change has occurred. In essence, habituation encodes implicit memory about the temporal dynamics of the world which are combined with perceptual input to aid temporal segregation. The current research is concerned with the degree and nature of this interaction relative to the occurrence and co-occurrence rates of different words.

In this paper, we investigate the influence of occurrence and co-occurrence on word identification. The consideration of frequency in word identification provides obvious benefits in uncertain situations. In general, if two alternative interpretations are equally plausible based upon input, the alternative occurring more often is the better choice. However, it is unknown if the processing dynamics of frequent items are distinct from those of less frequent items. Taking advantage of co-occurrence relationships between words is similarly advantageous. As one representation is activated, expectation

of co-occurring items, particularly items expected to occur after the activated item, can

aid perception. Here we present three experiments testing the influence of occurrence and

co-occurrence on the dynamics of perceptual activation, and extend the nROUSE model

to account for these effects. First we introduce the details of the behavioral paradigm

used to measure these effects along with previous behavioral results, followed by an

explanation of the nROUSE model.
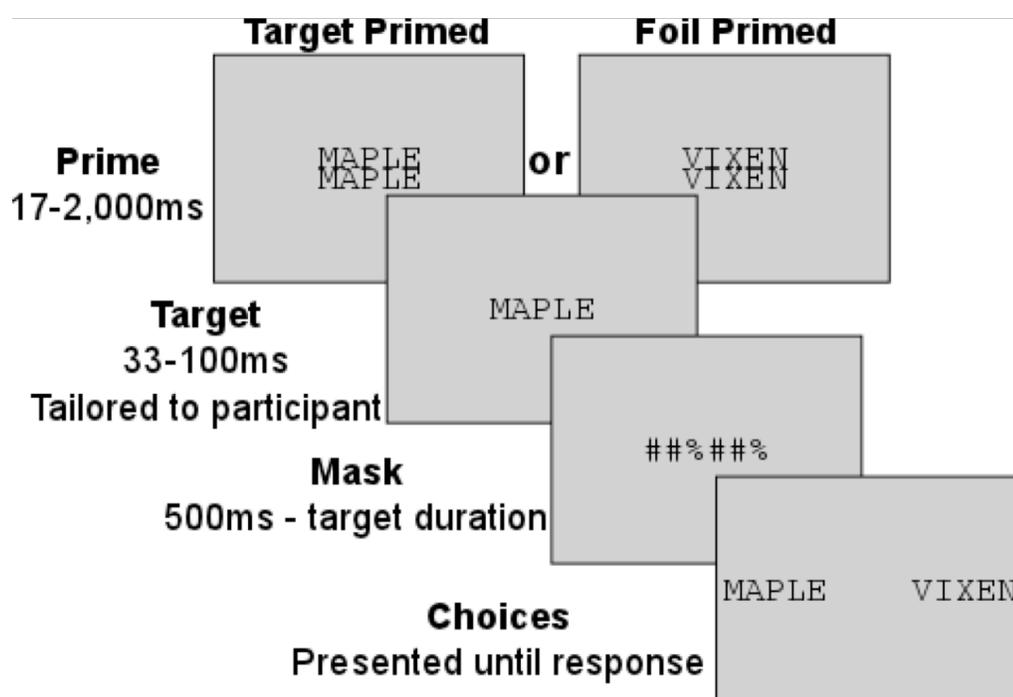
Figure 2.1. Basic trial structure in Chapter 2 experiments. Trials in Experiments 2.1 through 2.3 followed this same general trial procedure with minor variations. After an initial fixation, a prime appeared for a manipulated duration, from 17 to 2,000 ms. The relationship between the prime and the primed item varied between experiment. In Experiment 2.1, the prime was identical to the primed item. In Experiment 2.2, the prime was an associate of the primed item, either in the forward or backwards direction, depending on the condition. In Experiment 2.3, the prime was either be a repetition of the primed item or a word likely to precede the primed item. After the prime, the target was displayed for between 33 and 100 ms. The target duration was tailored to each participant, so that they could achieve at least 75% accuracy with an unrelated 150 ms prime. The target was immediately followed by a mask, presented for a duration such that the target and mask combined lasted 500 ms. Finally, two choices were presented, left and right of where the target appeared, one identical to the target, and another an incorrect foil choice. Participants then made their responses using a button box.

Measuring assimilation and segregation with immediate repetition priming

The present ideas surrounding habituation's role in perception originated from studies of short-term priming. Huber et al. (2001) used a short-term forced choice priming procedure to discriminate priming's influence on perception (affecting only the perception of the target) from preference (influence on later decisions). Short-term priming paradigms allow the perceptual stream to be carefully manipulated, and its influences on subjective perception measured. Unlike related paradigms, such as long-term lexical decision priming (Meyer & Schvaneveldt, 1971; Meyer, Schvaneveldt, & Ruddy, 1974), priming effects are measured primarily in terms of forced-choice identification accuracy, which is less sensitive to response strategies and allows the measurement of bias. Although forced choice accuracy is typically the primary dependent variable, reaction time can also be recorded and analyzed (Huber & Cousineau, 2003). Similarly, other investigations use named identifcation in place of forced-choice identificaiton (Pecher, Zeelenberg, & Raaijmakers, 2002). Figure 2.1 presents an overview of the general trial procedure. After an initial fixation cross, a prime word is presented for a variable amount of time. Typically, the duration of this prime varies from very short and barely perceptible to a few seconds long. In initial studies, active processing tasks involving the prime, such as animacy judgments, were used in place of long primes. Long primes were later found to have similar effects (Huber, Shiffrin, Quach, et al., 2002). Participants are correctly informed that the prime is non-diagnostic, but may be identical to one of the choices (repetition priming), share orthographic features with the choices (orthographic priming), or be an associate of one of the choices (associative priming). Immediately following the prime, a target word briefly appears for

a duration so that a given participant's baseline performance (without priming effects) is around 75%. After the target presentation, a pattern mask appears, followed by two choice words, the correct target word, and an incorrect foil word. Performance is measured as the proportion of times the participant chooses the target.

Using this procedure, Huber et al. (2001) found opposing priming effects for different prime exposures. With short prime durations, repetition priming produced a benefit for the repeated word, termed positive priming. However, long durations or actively processed primes produced a repetition deficit, or negative priming. This was explained by the ROUSE model as consequences of under- or over-discounting the influences of activation due to the prime (Huber et al., 2001; Huber, Shiffrin, Quach, et al., 2002). To the degree that the prime word matches the target word, its features are attributed to the target and the prime increases accuracy. However, for stronger pre-activations of the target by, for example, a long duration prime identical to the target, feature activation of the prime is discounted and accuracy decreases. With long enough priming, features similar to the prime are over-discounted, and the priming effect reverses into a deficit. Activation of the foil choice by the prime has the opposite effect. Weak activations of the foil representation hurt accuracy because primed features, which are similar to the foil, are attributed to the target presentation. However, stronger activations by a foil prime actually help accuracy through excessive discounting of foil related features. The ROUSE model explains both positive short duration priming effects and negative long duration priming effects as the result of attributing activation to sources with varying degrees of prime discounting.

The nROUSE model

The nROUSE model (Huber & O'Reilly, 2003) uses an artificial neural network with activity-dependent synaptic depression to extend the ROUSE model. It has been used successfully to account for a variety of word priming data, including immediate repetition and associative priming (Huber & O'Reilly, 2003), as well as face identity priming (Rieth & Huber, 2010). Here we present a conceptual overview of the general model. Further details are presented in Appendix A. There are three critical components of the model as applied to word priming, 1) a perceptual cascade structure, 2) activity-dependent synaptic depression, and 3) selection of response through a race process.

Figure 2.2 presents an overview of the three-layer model structure. In word priming, the levels are conceptualized as coding visual, orthographic, and semantic information. Perceptual processing is often characterized as a processing hierarchy, from simple visual features to more complex representations. Previous perceptual identification studies have shown that perception occurs in a cascade, such that information at higher levels of processing (e.g. semantic) begins to accrue before the influence of lower level processing (e.g. orthography) is completed (Duñabeitia, Carreiras, & Perea, 2008; Pecher, Zeelenberg, & Wagenmakers, 2005; Rumelhart & McClelland, 1982; Weldon, 1993). In the model, this is directly reflected by the activity of feed-forward connections from lower to higher levels. There are additionally feedback connections from the third level to the second level, and inhibitory connections within levels coding for the same spatial area. As applied to immediate repetition priming, there is a node representing each word stimulus at each layer, and a node representing the mask at the first and second layer (the mask has no semantic content). Applied to associative

priming, shared representations are added at the semantic level coding for semantic content captured by the association between two items (Huber & O'Reilly, 2003). Shared nodes receive input from, and provide feedback to, both associates. The model nodes are considered to be proxies for rich representations of different visual stimuli. Connections between nodes in different layers representing the same stimuli are built into the model. The first layer is assumed to code for both location and content. Therefore, in the first visual layer there are representations of the target and foil for each presentation location and inhibition only occurs within a particular visual location. Input causes the corresponding node to activate and propagating output according to its connections.

The output activation of each node is a function of the node's membrane potential, which tracks recent input to the node, scaled by the node's current amplitude, which tracks the recent output of the node. The separate amplitude term serves to scale down the strength of output after prolonged activity. Membrane potential is driven up by the output of feed-forward and feedback connections, and down by a constant leak factor and in proportion to the total activity of nodes coding for the same space. Node amplitude can be conceptualized as the amount of available resources of a node or synapse. The amplitude of a recently inactive node is 1, having no effect on membrane potential. However, the amplitude scalar is driven down as a function of recent activity. This has the effect of decreasing output over time with sustained activation, a phenomena predicted from computational neuroscience (Grossberg, 1968, 1969) and later found in real neural cells known as synaptic depression (Abbott, Varela, Sen, & Nelson, 1997; Tsodyks & Markram, 1997). Without continued input, node output and membrane potential gradually fall to 0, and amplitude slowly recovers to 1.
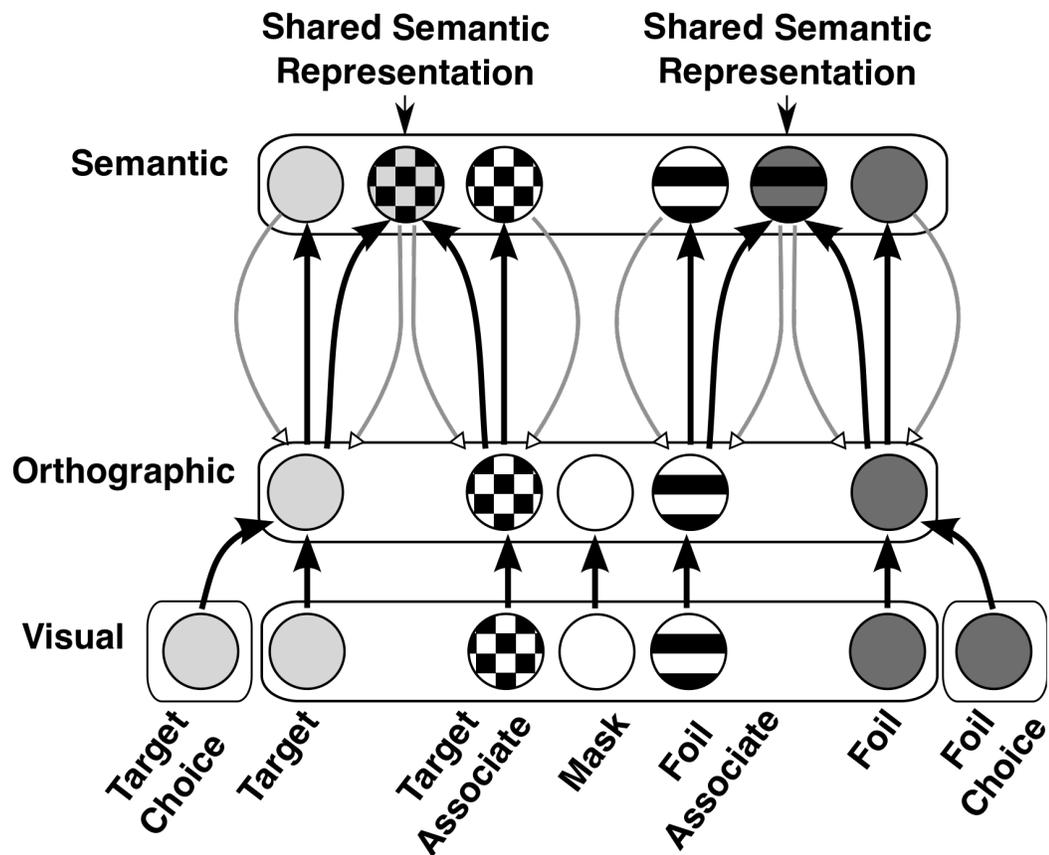
Figure 2.2. General model structure. Feed-forward links are black, feedback links are in gray. Activation of nodes within boxes inhibit other nodes within the box. The visual layer nodes code for both position and location. Thus there are separate nodes for centrally presented targets and foils and the peripherally presented choices, each inhibited only by other nodes coding for the same location. Higher levels of representation are location independent, and all inhibit one another. Shared nodes receive input from two orthographies, and represent the semantics invoked by associative or expectation relationships between items.

Finally, the activation of the model is mapped to behavioral accuracy through a race process. Initially, choice presentation will result in the output of the semantic representations of both choice alternatives rising. However, because of synaptic depression, after a period of sustained input to choice representations, the output of these nodes will drop off. We treat the point at which output drops as the identification latency. The choice with the shorter identification latency is the response. At this time, the decreases in amplitude from sustained activation balance the increases in membrane

potential. The identification latency of the semantic node corresponding to each choice is considered to be the mean of a distribution of possible latencies on any given trial. The variance of these distributions are proportional to the square of the mean identification latency. This way, distributions with longer average identification latencies have wider variability, and quickly identified items are less variable. Proportion correct is therefore the proportion of times the correct target node is identified first.

Together, these aspects of the model work to predict priming accuracy in the following manner. For a given condition, input is provided to lower layer nodes corresponding to the stimulus presented at each millisecond of the trial. Activation from short primes travels through the perceptual cascade, giving the primed responses a head start in the race to respond. Thus, when the prime is a target, identification latency is faster than it otherwise would be, and accuracy is enhanced due to the relative advantage in the race process. When the prime is a foil, the foil node gets a head start, and accuracy decreases. With longer prime durations, however, the sustained activation of the prime depletes the available resources of the prime representations, decreasing the responsiveness of the node to further activation from the target and/or choice presentation. This decrease in responsiveness has the effect of slowing the identification latency of a node, leading to decreases in accuracy for long target primes, and increases for long foil primes.

The nROUSE model provides a framework to understand the dynamics of priming and how they are influenced by occurrence and co-occurrence statistics. In this paper we model the influence of occurrence and co-occurrence by increases in the connection strengths from orthographic to semantic representations. For high frequency

items, the strength of output transmission between the corresponding orthography and semantics is higher than for low frequency items. Shared concepts between co-occurring items are represented by an additional semantic node representing the concept elicited by the co-occurrence of the words. Higher connection strengths to this shared node capture expectations for the shared concept from particular words.

Current Experiments

In this paper, we present three experiments. In Experiment 2.1 we compare words, pronounceable non-words, non-pronounceable non-words, and inverted words to test the effects of occurrence rates on immediate priming. In Experiment 2.2 we use word association as a proxy for co-occurrence, and investigate the influence of directional associations. In Experiment 2.3 we use co-occurrence statistics to directly test the effect of co-occurrence on priming dynamics, including a condition where word repetitions are actually expected. In combination, the results of these three experiments inform our understanding of the adaptations of the mind to the temporal occurrence and co-occurrence rates in word identification.

Experiment 2.1

Experiment 2.1 tests the influence of stimuli occurrence rate on the dynamics of word identification. Manipulating word frequency (Scarborough & Cortese, 1977), or comparing words to non-words (Bodner & Masson, 1997), affects reaction time in long term priming paradigms. The simple environmental statistic of rate of occurrence is captured by word frequency, and also correlated with familiarity of words. There are many known effects of familiarity and frequency on both language processing (Arnon &

Snider, 2010; Ellis, 2002; Sereno, Pacht, & Rayner, 1992; Siyanova-Chanturia, Conklin, & van Heuven, 2011), and memory (Balota & Spieler, 1999; Gregg, 1976). For example, in artificial grammar learning tasks, familiarity is a key component in implicit learning, predicting grammaticality judgments and confidence ratings (Scott & Dienes, 2008). As words are encountered repeatedly and become more familiar, the corresponding mental representations of those words change. One possible change is faster and stronger activation of the corresponding representation. A number of different learning schemes predict this. For example, low frequency stimuli might activate a large number of neurons to varying small amounts in a population code. After repeated exposure, these representations may become more finely tuned, representing the same stimuli with a smaller number of more strongly activating neurons in an exemplar code (for examples of possible computational changes with learning, see Palmeri, Wong, & Gauthier, 2004). Additionally, representations of frequently encountered stimuli may activate more quickly. This is found in reaction time modeling as a change in the accumulation rate of evidence for familiar items (Brown & Heathcote, 2005). Regardless of the implementation details, activity dependent habituation predicts that the peak priming effects and the transition from positive to negative priming should occur at shorter prime durations for known items than for unknown items. In line with these predictions, primes that are recognized, and presumably processed more completely, result in greater discounting (Huber, Shiffrin, Quach, et al., 2002). Furthermore, although upright and inverted faces contain the same visual information, they habituate differently and to varying degrees for different individuals (Rieth & Huber, 2010). However, faces are

often considered a 'special' stimulus (Yue & Tjan, 2006), and may be expected to have unique processing characteristics related to inverted faces (Yin, 1969).

An equally logical hypothesis is that words carry more information and are therefore processed longer and discounted less than non-words. Unfamiliar stimuli, like non-words, which contain no useful information, may be quickly discounted by the mind as noise.

Rather than test high versus low frequency words, to increase the power of the experiment, we compared words to non-words, and also to inverted words. Pronounceable non-words share low level properties with words, but have a frequency of 0. Any semantic content they do have will be through associations with actual words in their orthographic neighborhood. To further explore the influences of occurrence statistics at different levels of representation, we also tested non-pronounceable non-words and inverted words. Non-pronounceable non-words are defined as non-words with letter strings never appearing in English, e.g. 'SRR'. Stimuli from this condition have even less semantic content, and given that the configurations of letters are unfamiliar, may even have reduced orthographic representations. Inverted words provide a further comparison. Low-level content is disrupted in inverted words. However, this disruption can be overcome with time, revealing high-level semantic information. In summary, non-words contain elements which occur frequently, but no semantic information. In contrast, occurrence rates are low for visual components of inverted words, while semantic experience is rich. By contrasting these conditions we can compare the contributions of experience at a number of processing stages. If occurrence rates affect processing prior to letter processing (e.g. line segments), or not at all, then priming dynamics should not

differ between any of stimuli types. However, if occurrence rates do influence perceptual dynamics at many levels of representation in parallel (e.g. line segments, letters, letter groups, orthography, semantics) there will be differences in priming dynamics between all stimuli types.

In the nROUSE model, assuming that words have higher connection strengths than non-words leads to words saturating the corresponding representations more quickly, speeding the decrease of positive priming and reversal to negative priming with longer prime durations. Essentially, words are processed faster. Another possibility is that the dynamics of positive and negative priming are instead related to the exhaustion of information content. This predicts that the higher occurrence rates for words will actually cause the transition from assimilation to segregation to occur at a longer prime duration.

Methods

Participants. A total of 63 undergraduate psychology students participated in this experiment for course credit.

Equipment. The experiment was conducted on LCD monitors with a 60 Hz refresh rate. Presentation times were synchronized with the display refresh rate. All items were displayed on a gray background and the display resolution was set to 640 by 480 pixels.

Stimuli. There were four stimulus types in the experiment: words, pronounceable non-words, non-pronounceable non-words, and inverted words. Stimuli were created starting with a 746 word list of syllable marked two-syllable five or six letter words. Words were hand selected to eliminate words of very low frequency. Importantly, words for the non-word conditions were created from components of the actual words, so that

each letter appeared in each position the same number of times for each stimulus type over the course of the experiment. This ensured that differences between the conditions cannot be explained by differences in letter-position frequency (e.g. non-words being more likely to start with the letter 'i'). This is a critical control, as letter representations are sensitive to letter position (New & Grainger, 2011). Non-word stimuli were created by combining pairs of word syllables to create non-words and non-pronounceable non-words. From each pair of two words, two potential pronounceable non-words were created by switching syllables between the two words. For example, from the words 'rustic' and 'radish' the pronounceable non-words 'rudish' and 'rastic' were created. Potential non-pronounceable non-words were created by alternating letters between the original words. In the previous example, the corresponding non-pronounceable non-words would be 'rudtsc' and 'rasiih.' Non-word status was confirmed by ensuring that they did not appear in a 410,000 word dictionary (containing words as well as abbreviations, alternate spellings and common names), and that they were either five or six letters long. All tri-graphs (3 letter combinations) were present in words in the dictionary for pronounceable non-words, and non-pronounceable non-words contained at least one tri-graph not appearing in the dictionary. All non-words used in the experiment were additionally checked by hand to ensure they were suitable. Stimuli for the inverted word condition were simply the word stimuli rotated 180 degrees. This procedure resulted in 434 unique stimuli for each condition. The stimuli were randomly assigned to prime durations and priming type without repetition for each participant. Furthermore the targets and foils for each trial were matched to have the same number of letters, so that target length was not a reliable distinguishing feature. All stimuli were presented in

Courier. Table 2.1 presents example stimuli of each type for each stimuli type and primed choice condition.

Table 2.1.

Stimuli Type	Primed Choice	Example Prime	Example Target	Example Foil
Words	Target	MAPLE	MAPLE	VIXEN
Words	Foil	VIXEN	MAPLE	VIXEN
Pronounceable non-words	Target	PAYEN	PAYEN	SKIKA
Pronounceable non-words	Foil	SKIKA	PAYEN	SKIKA
Non-pronounceable non-words	Target	BNSEN	BNSEN	MATLL
Non-pronounceable non-words	Foil	MATLL	BNSEN	MATLL
Inverted words	Target	DOGMA	DOGMA	TAFFY
Inverted words	Foil	TAFFY	DOGMA	TAFFY

Procedure. Figure 2.1 presents an overview of the trial paradigm. The participants' task during each trial was to identify a briefly and centrally flashed target while ignoring a preceding prime word and match it to one of two choices. Primes were presented immediately prior to the target in the center of the screen for durations ranging from 17 ms (i.e., subliminal priming) to 2,000 ms (i.e., excessive priming). In this manner, the prime word also served as a forward mask to the target word. Target identification ability was tested by means of a forced choice between the target word and a foil word. Responses were not speeded and accuracy feedback was provided on every trial. Participants were explicitly informed that the prime itself would relate to the target as often as it would relate to the incorrect answer and therefore that there was no effective response strategy that used knowledge of the prime. They were instructed that the most effective strategy was to focus on the center, ignore the prime, and wait for the target.

Each trial began with a fixation cross presented for a duration such that the combined duration of the fixation cross and the prime word was 2,500 ms. In the target primed condition, the primes were identical to the target word, and therefore some cue

was needed to distinguish primes from targets. Two identical versions of the prime were presented, one directly above the other. Immediately after the prime, a target word was briefly presented in the center of the screen, and masked with a character mask consisting of random arrangements of the symbols '#' and '%'. The mask was presented for a duration such that the duration of mask and target combined was 500 ms. This maintained a constant duration between onset of the target and onset of choice alternatives. Following the mask, the target and foil words were presented on the left and right sides of the screen, separated by approximately 18 degrees of visual angle. Responses were collected through a button box, with the far left and right buttons corresponding to the selection of the left or right option on the screen. The two choices were presented until a response was made, after which accuracy feedback was presented for 1,500 ms. The choice position of the target word was randomly counterbalanced between the left and right sides.

After instruction, participants completed one or more blocks of 16 practice trials. For these trials, the duration of the target was gradually lowered from 300 ms to 100 ms to accommodate participants to the pace of presentation and the prime duration was fixed at 500 ms. Once they achieved over 75% correct, they continued to the experiment.

To maximize priming effects, accuracy rates need to be around 75%. If the task is too easy, accuracy will be very high and priming that aids performance cannot be observed. For priming that hinders performance, the same is true if the task is too hard and accuracy is at chance. Due to individual variability in basic aptitudes for this perception task, the target duration was tailored for each participant so that they accurately responded to approximately 75% of trials without an unrelated prime. To find

this threshold target duration, after completion of practice trials, participants completed a block of 80 trials where target duration was randomly varied between 33, 50, 67, and 83 ms (20 trials at each target duration). There were 20 trials of each stimulus type (word, pronounceable non-word, non-pronounceable non-word, and inverted word). On these trials the prime duration was fixed at 150 ms and the prime was unrelated to either choice. After the completion of this threshold block, participants were assigned the lowest target duration at which they answered 75% of the trials correct. If 75% identification was not achieved with the 83 ms target, a target duration of 100 ms was adopted for the remainder of the experiment.

During the main portion of the experiment, three variables were manipulated within subject: five levels of prime duration (17, 50, 150, 400 or 2,000 ms), primed choice with two levels (target primed and foil primed), and stimulus type with four levels (words, pronounceable non-words, non-pronounceable non-words, and inverted words). In the target and foil primed conditions, the prime itself was identical to either the target or foil, respectively. Additionally, there were two between subject variants of the experiment. In the first (N= 29 participants), the stimulus type was randomly counterbalanced for each trial (random presentation). In the second variant (N= 34 participants), the stimulus type was consistent throughout a block, with one stimulus type for each block of the experiment (blocked presentation). There were 10 trials for each condition, and the experiment was conducted in 4 blocks of 80 priming trials. Participants were given breaks between blocks of the experiment. In total, the experiment lasted approximately 35 to 45 minutes.
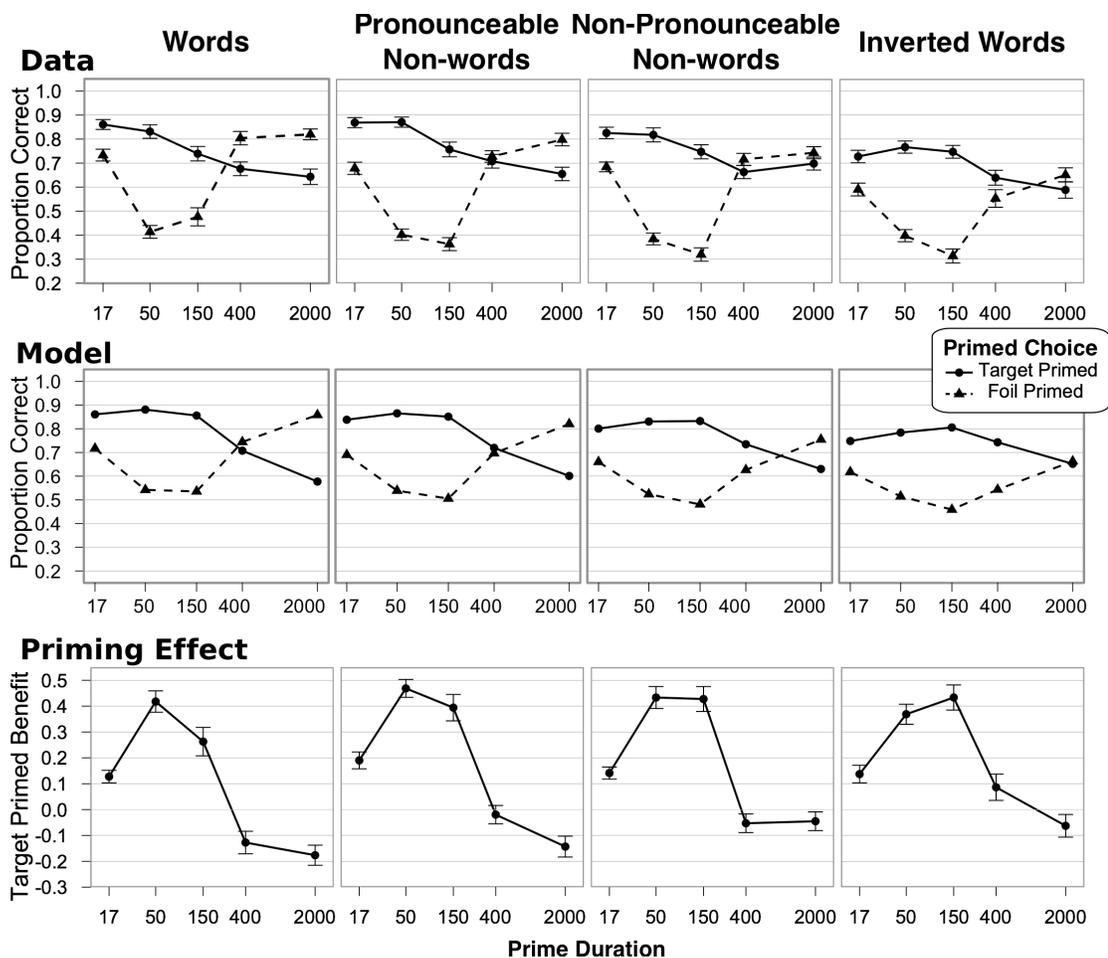
Figure 2.3. Experiment 2.1 results by stimulus type. Error bars are +/- 1 standard error of the mean or mean difference over participants. The top row shows the accuracy results from the experiment, the middle row shows the model results, and the bottom row shows the difference between the foil and target primed conditions in the behavioral data. The peak priming effect and transition from positive to negative processing occurs with the shortest prime durations for words, followed in order by pronounceable non-words, non-pronounceable non-word, and inverted words. Full details of the model are in the Appendix. The critical parameter varied in the model was the connection strength from the orthographic to the semantic later. This strength was set to 1 for words, .9 for pronounceable non-words, .75 for non-pronounceable non-words, and .55 for inverted words.

Results

Data from 12 participants who were outliers (with median accuracy or RT greater than 1.5 times the interquartile range from the median) for either the threshold or main portion of the experiment were excluded from the analyses, leaving data from 51 participants. Figure 2.3 presents the proportion of correct trials for each prime duration,

primed choice and stimulus type condition averaged over participants. The results for words replicated those of previous experiments using this paradigm (Huber, 2008). There was a significant three way interaction between the factors of primed choice, prime duration, and stimulus condition on proportion correct ($F(12, 600) = 3.80, p < .001$) in a repeated measures ANOVA. Priming with words resulted in positive priming effects (accuracy was higher for target primed trials) at prime durations of 17, 50 and 150 ms ($t(50) >= 4.77, p < .001$, all reported p-values are Holm adjusted for multiple comparisons). Significant negative priming effects were observed at 400 and 2,000 ms primes ($t(50) <= -2.92, p <= .03$). Non-words and inverted words showed similar, but increasingly delayed transitions from positive to negative priming. There was significant positive priming for pronounceable non-words at 17 through 150 ms ($t(50) >= 5.86, p < .001$), no priming effect at 400 ms ($t(50) = -.55, p = .586$), and significant negative priming at 2,000 ms ($t(50) = -3.55, p = .006$). There was significant positive priming for pronounceable non-words at 17-150 ms ($t(50) >= 5.86, p < .001$), no priming effect at 400 ms ($t(50) = -.55, p = .586$), and significant negative priming at 2,000 ms ($t(50) = -3.55, p = .006$). Non-pronounceable non-words also produced positive priming for prime durations of 17-150 ms ($t(50) >= 6.02, p < .001$), but no priming effect at 400 or 2,000 ms ($t(50) >= -1.47, p = .586$). The use of inverted words as stimuli similarly resulted in positive priming for prime durations of 17-150 ms ($t(50) >= 4.04, p <= .001$), no priming effects at 400 ($t(50) = 1.69, p = .478$) or 2,000 ms ($t(50) = -1.43, p = .586$). The size of the target primed advantage increased for all stimuli from 17 to 50 ms primes ($t(50) <= -5.17, p < .001$). When using words, the size of the priming effect decreased from 50 to 150 ms primes ($t(50) = 3.52, p = .007$), but there was no significant change for any of the

other stimuli types ($|t(50)| <= 1.65$, $p >= .522$). From 150 to 400 ms primes, the target

primed advantage decreased for all stimuli types, ($t(50) >= 6.43$, $p < .001$). Finally, there

was no significant change from 150 to 2,000 ms for words ($t(50) = 1.60$, $p = .522$),

pronounceable non-words ($t(50) = 2.61$, $p = .070$), or non-pronounceable non-words

($t(50) = -.20$, $p = 1$). However, for inverted words the target primed advantage continued

to decrease ($t(50) = 2.83$, $p = .046$).

Including presentation type (random versus blocked) in the repeated measures

model resulted in several presentation type interactions. There were significant three way

interactions between presentation type, primed choice, and prime duration ($F(12, 588) =$

2.11, $p = .015$), and presentation type, prime duration, and stimulus type ($F(4, 196) =$

3.53, $p = .008$). There were additionally two way interactions involving presentation type,

one with stimulus type ($F(3, 147) = 3.09$, $p = .029$), and another with primed choice ($F(1,$

$49) = 4.53$, $p = .038$). For words in random presentation, negative priming at 400 and

2,000 ms was not significant ($t(23) >= -2.62$, $p >= .122$) and the size of target primed

benefit did not increase between 50 and 150 ms prime duration ($t(23) = 1.39$, $p = .429$).

In contrast, when presented in blocks, the transition to negative priming occurred sooner,

such that there was not a significant positive priming effect for 150 ms primes ($t(26) =$

1.72, $p = .488$). Overall, the transition from positive to negative priming for words

occurred more rapidly for blocked presentation. These observations suggest a tendency to

view words and pronounceable non-words as more similar when stimulus types are

mixed in random presentation.

Statistical analysis of performance during the threshold portion of the experiment

revealed main effects of stimulus type ($F(3, 186) = 27.72$, $p < .001$) and target duration

(F(3, 186) = 102.00, p < .001). Averaged over target durations, word stimuli were perceived more accurately than all other stimuli types (all t(251) >= 2.43, p <= .031). Performance was higher on pronounceable non-words than both non-pronounceable non-words and inverted words (t(251) >= 2.03, p <= .043). Lastly, non-pronounceable non-words were more accurately identified than inverted words only (t(251) = 5.05, p < .001).

Discussion

Stimulus type strongly influenced the transition from positive to negative priming, as predicted by the nROUSE model through decreasing orthographic to semantic connection strengths for decreasing levels of frequency. Words produced positive priming for 17 through 150 ms, transitioning to negative priming for 400 and 2,000 ms primes. The peak priming effect for words occurred with 50 ms primes. For non-words, the transition slowed, showing negative priming only for 2,000 ms primes for pronounceable non-words, with no significant difference in the magnitude of priming between 50 and 150 ms. Non-pronounceable non-words and inverted words never reversed to negative priming, with the peak priming effect continuing to shift to longer prime durations. These differences were also reflected in perceptual identification from the trials determining the optimal target duration for each participant. Accuracy was highest for words, and progressively lower for pronounceable non-words, non-pronounceable non-words, and inverted words. These findings mirror other perceptual identification differences where familiarity was experimentally investigated through varied exposure to unknown items (Nelson, 2009).

Using face images instead of words in the same paradigm produced interesting individual differences not seen with words (Rieth & Huber, 2010). Participants who were

able to identify the target from shorter flashes showed similar priming dynamics to that of words in the current experiment, with a full reversal to negative priming with 2,000 ms primes. Participants who needed longer duration targets for accurate identification however showed priming dynamics more similar to the non-pronounceable non-words or inverted words, with no reversal to negative priming. The present results suggest that the individual differences in faces may be the result of differences in fluency between participants for the face stimuli, similar to the differences between words and non-words. The individual differences may reflect different degrees of familiarity with the types of face stimuli used, or possibly general differences in face representation.

The results also provided evidence that in the context of non-words, words show similar processing dynamics to non-words, i.e. semantic information is selectively activated. In masked and lexical decision priming contexts, effects related to the proportion of validly primed trials have been found (Bodner & Dypvik, 2005; Bodner & Masson, 2001; Pecher, Zeelenberg, & Raaijmakers, 1998; Weidemann, Huber, & Shiffrin, 2008). However this finding suggests the possibility that in a mixed context, the semantic content of words is ignored, and in doing so words are perceived similarly to non-words. When the stimulus type was randomly varied trial-to-trial, the transition from positive to negative priming for words was more similar to pronounceable non-words.

Experiment 2.2

Co-occurrence relationships between stimuli may also provide information for source separation. If it is known that some particular item is likely to follow the present one, a rational system should take this into account to aid temporal segregation.

Experiment 2.2 tests this hypothesis using directional associations. An associate pair is a set of two words where presentation of one word, the cue, frequently brings the second word to mind, the target. To minimize confusion between targets from associative pairs and priming targets, we refer to the targets from associative pairs as associative targets. Many associate pairs are bi-directional, meaning each brings the other to mind symmetrically (e.g., 'light' and 'dark'). However, in order to test the influence of one word creating a directional expectation for a second, in this experiment we used unidirectional associate pairs. In these associative cue-target pairs, the associative relationship was only in one direction (e.g, 'saber' – 'tooth').

While it may be assumed that forward associations are more critical, backward associations, e.g. from 'tooth' to 'saber', do play an important role in perception. Backwards transition probabilities (the probability of a word given the occurrence of the following word) have been found to be as important as forward probabilities for language processing (Pacton & Perruchet, 2008; Perruchet, 2004). When priming associative tasks, both forward and backward associations are required to observe priming effects (Zeelenberg, Shiffrin, & Raaijmakers, 1999). In long-term priming, semantic properties have been found to underlie both forward and backward associations, perhaps explaining why directional effects have not been observed; priming may occur at a bidirectional semantic, not associative, level (Lucas, 2000).

Prior work has demonstrated that associative priming effects wax and wane from shorter prime durations to longer durations or actively processed primes (Huber, 2008; Huber et al., 2001). Associative priming is captured in nROUSE by the inclusion of a shared node at the semantic level, activated by both the associative cue and target. Choice

output of the model is taken as a weighted combination of the output of the unique choice and the shared node. In the model, associative priming produces the waxing and waning of positive priming instead of the full transition to negative priming. Negative priming does not occur because the prime does not fully habituate the unique target or foil representation, only the shared node. In this case, shorter primes produce a positive priming effect through the shared node, and with longer prime durations the prime is appropriately, but not excessively, discounted through habituation of the shared node. Here we model the influence of directionality as an increase in the connection strength from the orthography of each associative cue to each shared node. This is similar to how occurrence was modeled for Experiment 2.1, but only applied to a single connection rather then for all stimuli in the trial. This models the properties of directional association as the result of stronger relationships between unidirectional associative cues and the shared cue-target concept.

Directional priming of the target and foil has been tested (Huber et al., 2001), finding no effect of prime duration. However, this prior experiment only used two prime strengths, raising the possibility that an effect may be seen with parametric variation of prime direction. Here we test primes at five different durations for a more sensitive test of directional effects. Experiment 2.2 consists of testing the two priming directions, forward associations and backwards associations, for target and foil primed choices and five prime durations.

Methods

Participants. A total of 63 undergraduate psychology students participated in this experiment for course credit.

Procedure. Except as noted, all procedures were identical to Experiment 2.1. In all trials there was some association between either the target or foil and the prime. In the forward association condition, the association was from the prime to the primed choice (i.e. the prime was the associative cue). The association was from primed word to prime for the backward association condition.

Stimuli. Selection of stimuli started with the initial word list used in Experiment 2.2 (the primary word list), which was cross referenced with the USF association norms. Word pairs were selected so that, as measured by the USF association norms (Nelson, McEvoy, & Schreiber, 1999), all associations were unidirectional (no backward association occurrences) and rated with a strength over .05 (mean = .16). All words were presented in upper case, and flanked by '#', so that the entire string was constantly 8 letters long. There were two priming direction conditions: forward associations and backward associations. Backward associations were created by switching the order of the associates. Table 2.2 presents example stimuli of each type for each prime condition.

Table 2.2.

Direction	Primed Choice	Example Prime	Example Target	Example Foil
Forward	Target	RAIL	ROAD	BONE
Forward	Foil	PORCH	TEETH	LIGHT
Backward	Target	WORD	CODE	SEND
Backward	Foil	WELL	FINAL	ALIVE

Results

Data from 15 participants who were outliers (with median accuracy or RT greater than 1.5 the interquartile range from the median) for either the threshold or main experimental portion of the experiment were excluded from the analyses, leaving 58 participants in Experiment 2.2. The average proportion of correct trials per participant for
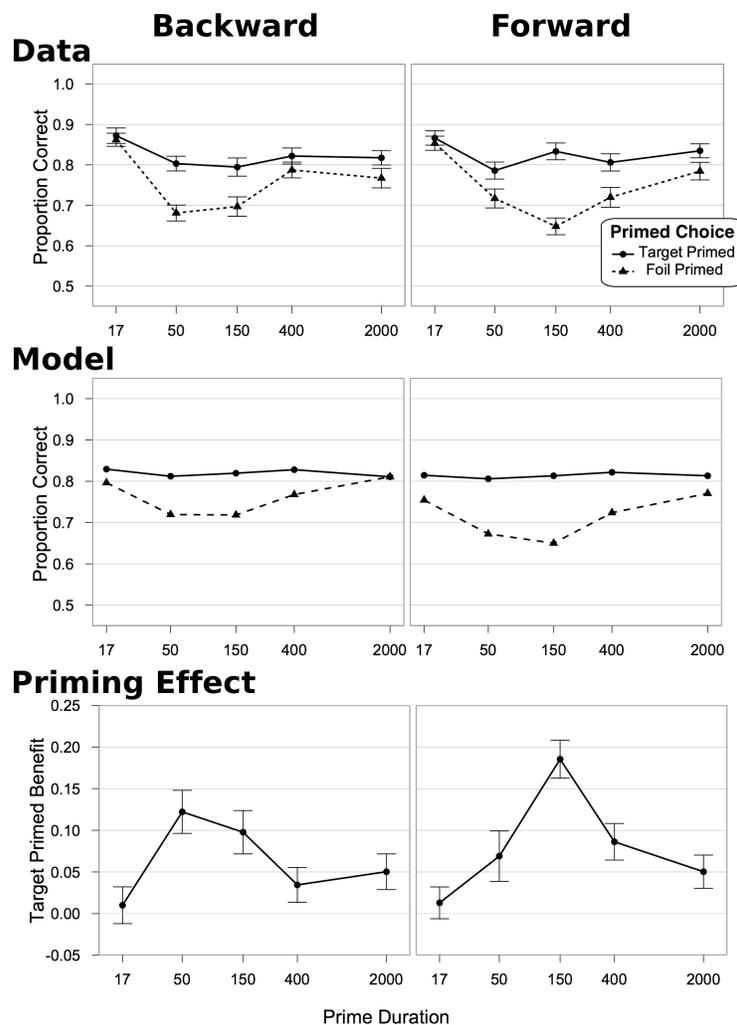
Figure 2.4. Experiment 2.2 accuracy by associative direction. Error bars are +/- 1 standard error of the mean or mean difference over participants. The top row shows the accuracy results from the experiment, the bottom row shows the model predictions for the best fit parameters. Both forward and backward associates produce priming effects, however the priming effects from forward associates peak more strongly and slightly later. Full details of the model are in the Appendix. The critical parameters varied here were the proportion of the representation that is shared between the prime and the primed word, set to .2, and the orthographic to semantic connection strength of directed association, set to 1.4. The directed association connection strength for forward priming was the connection from the prime to the shared representation. For backward priming it was the connection from the target and foil to the respective shared representations.

each condition in Experiment 2.2 is shown in Figure 2.4. There was a significant three way interaction between primed choice, prime duration, and associative direction $(F(4,228) = 3.00, p = .019)$. There were priming effects for both forward and backward priming. Overall, the positive effects of priming increase with longer durations to a point,

around 150 ms, after which the effects of priming shrink (significant interaction between prime duration and primed choice, ($F(4,228) = 8.65$, $p < .001$). However, there were also small, but significant, differences between prime directions. Forward associate primes had significant priming effects for 150 and 400 ms primes ($t(57) >= 3.92$, $p <= .002$). However, backward primes actually had priming effects at shorter 50 ms primes as well as 150 ms primes ($t(57) >= 3.76$, $p <= .002$), with the effect fading by 400 ms ($t(57) = 1.65$, $p = .313$). The target primed benefit increased for forward associative priming from 50 to 150 ms ($t(57) = -3.41$, $p = .009$), but decreased from 150 to 400 ms ($t(57) = 3.73$, $p = .003$). The only significant difference between adjacent prime durations for backward priming was from 17 ms primes to 50 ms primes, where the target primed benefit showed a significant increase ($t(57) = -3.23$, $p = .012$).

Discussion

Associative priming produced a waxing and waning positive priming effect for both forward and backward associative priming, as expected from previous results. However, parametric manipulation of prime duration revealed an asymmetry in the priming effect. The peak priming for unidirectional forward primes was later than the peak for primes that were backwards associates of the primed choice. The model captured the general pattern of data through the inclusion of a shared associative semantic node for each primed choice, and stronger connections from associative cues to this shared representation. Stronger connections from the forward associative cues produced larger priming effects by speeding the activation of the shared node. However, because of the lack of representational overlap between the prime and the target choices, the time course of priming remained relatively unaffected and negative priming did not occur.

Furthermore, in backward priming, both choices were associative cues, and activated separate shared representations with increased strength. This had the net effect of reducing the influence of the prime on peak times.

The results of Experiment 2.2 suggest that there is an effect of directionality on the time course of word identification. However, precise interpretation of this result is dependent on the theory of the generative process for associated word pairs. Based on post-hoc observations, there appeared to be two kinds of associates. The first consists of two semantically related items, usually with the associated target higher in word frequency than the cue, e.g. 'fling' and 'throw'. Using these pairs seemed to produce little effect of directional priming. The other type had a clear unidirectional relationship between two words, either forming a single word, e.g. 'screw' and 'driver', or a high frequency combination, e.g. 'metric' and 'system'. These associates showed stronger directional effects, possibly linked to the role of expectation. Overall, this suggests that unidirectional associative relationships span several different types of co-occurrence structures. Because of this, in Experiment 2.3 we used actual co-occurrence statistics to compare priming through expectation for a non-repeated word, expectation for a repeated word, or non-expected repetitions.

Experiment 2.3

Experiment 2.3 tested the role of expectation directly by using linguistic co-occurrence. Primes were selected that were either likely to precede the target or foil, or that never proceeded the target or foil. Several previous studies have explored the role of expectation in priming. In masked repetition priming, adaptation to prime statistics can

occur even without conscious knowledge (Bodner & Masson, 2003). In short term repetition priming, participants are able to learn general diagnostic relationships between primes and targets, even then the relationship is dependent on prime duration, but notably not prime color (Weidemann et al., 2008). Primes that create an explicit expectation of a particular category do enhance perceptual identification, although these effects emerge more slowly than priming due to repetition of a category (Neely, 1977). However unlike these past studies, which manipulated the predictive power of the prime indicating the correct response in general or by category, in this experiment we are interested in how expectations arise between specific prime and target words.

We used three types of priming to explore the role of expectation from co-occurrence in temporal segregation. The first type of priming was non-expected repetitions. This was similar to the previously tested immediate repetition priming, where the prime was a repetition of the target or foil, (e.g., 'cloth-of' priming 'CLOTH'). The words in this condition were chosen especially so they would not typically repeat. The second type of priming, expected repetition, used words which plausibly repeat on a short time scale , (e.g., 'back-to' priming 'BACK'). The third type of priming, expected non-repetition, tested the effect of expectation on the dynamics of priming by using primes that were distinct from, but that co-occur with, the target, (e.g., 'stake-a' priming 'CLAIM'). Alternative interpretations of priming dynamics predict very different time courses of priming for this condition. If discounting effects are indeed caused by the completion of processing of a prime (as would be predicted from Experiment 2.1), discounting effects should still be present when the prime is an expected repetition. However, an alternative account explains priming effects as the result of high level

expectation of change. After only a short prime, the prime might be assimilated with the target, because the prime was not expected to be so short. However, a longer prime may bias perception away from the prime, because the prime was not expected to repeat. If these effects are due to high level prediction of the likely next word, discounting of words which were expected to repeat would lessen. Thus, the expected repetition condition tested the level of processing at which the dynamics of expectation occur, and the degree to which discounting reflects processing speed versus prediction of likely future states.

To equate co-occurrence more accurately between these conditions, primes in Experiment 2.3 consisted of two words. Thus in the expected repetition and expected non-repetition conditions, the primed item is a grammatically valid completion. This experiment used repetition, so some manner of distinguishing between primes and targets was needed. Presenting the prime as two copies one below the other, as in Experiment 2.1, reduced the tendency to view the primes linguistically, and thus possibly reduced the tendency to form expectations. Therefore we used lower case primes. Lower case primes also allow us to test the degree to which these effects are entirely based on the precise orthography of the prime. Many previous long term (Meyer et al., 1974; Sereno, 1991) and short term (Huber et al., 2001) priming results have found strong effects of orthography. Lastly, this experiment includes a neutral prime condition, where the prime is unrelated to either choice to serve as an appropriate comparison for the target and foil conditions given the effects of masking.

This experiment was modeled in nROUSE in a similar manner to Experiment 2.2. The expectation of a co-occurrence, either for a non-repetition like 'stake-a-claim' or an expected repetition like 'back-to-back', was given a shared node activated by the

appropriate orthographies, either 'stake' and 'claim' or just 'back'. This node had a stronger connection strength from the word creating the expectation, here either 'stake' or 'back'. It may seem odd in the model that expected repetitions essentially have two semantic representations, but they are coding for two separate meanings of the word. One is the unique meaning of the word in isolation, and the other is the meaning of the word if repeated.

Table 2.3.

Type of Priming	Primed Choice	Example Prime	Example Target	Example Foil
Non-Expected Repetition	Target	cloth-of	CLOTH	FERRY
Non-Expected Repetition	Foil	cloth-of	FERRY	CLOTH
Non-Expected Repetition	Neither	cloth-of	WAIST	CHILL
Expected Non-Repetition	Target	stake-a	CLAIM	ALIEN
Expected Non-Repetition	Foil	stake-a	ALIEN	CLAIM
Expected Non-Repetition	Neither	stake-a	BRIDGE	SLEEVE
Expected Repetition	Target	back-to	BACK	FIST
Expected Repetition	Foil	back-to	FIST	BACK
Expected Repetition	Neither	back-to	FIRE	BOLT

Methods

Participants. A total of 119 undergraduate psychology students participated in this experiment for course credit.

Procedure. Except as noted all procedures were identical to Experiment 2.1. In this experiment primes were two word sequences. They were presented centered and in lower case, separated by a hyphen, see Table 2.3 for examples. Targets and choices were presented in upper case to distinguish them from primes. The type of priming was manipulated between expected repetitions, expected non-repetitions, and non-expected repetitions. For the expected repetition condition, the prime was a repetition of the primed word that would be expected given the prime. Expected non-repetitions referred

to word sequences where the primed word would be expected, but not a repetition. Finally, non-expected repetitions were word sequences where the prime word was a repetition of the first word in the sequence, but not one that would be expected. Within each type of priming, the primed choice was manipulated between word sequences that primed the foil choice, the target choice, and sequences which primed neither choice (foil, target, and neither primed conditions). Neither primed conditions were logically indistinguishable between the three type of priming conditions; in all cases the third word in the sequence was neither a repetition, nor an expected word. However, to control for any possible word frequency differences between the different type of priming conditions, separate neither primed conditions were tested for each type of priming, created from the same word stimuli used for the corresponding type of priming.

Stimuli. Word sequences were created from an initial list of 1560 four to six letter words (the primary word list), obtained from the MRC linguistic database (Coltheart, 1981), all with Kucera-Francis written word frequency above 10. Additional two to four letter words were selected to be used as the second word in the sequences (the secondary word list). These words were: a, an, and, are, as, at, by, for, from, go, he, in, is, it, my, no, of, on, one, or, so, that, the, them, this, to, up, was, we, and with. The Google-n-gram database, (Brants & Franz, 2006), which is a collection of the number of online occurrences of word sequences, was used to determine the expectancy of each possible three word sequence created by a word from the primary word list, followed by a word from the secondary word list, followed by each word on the primary word list. For each occurring triplet of words, the conditional proportion of times the third word occurred after the first two words out of all occurrences of the first two words was computed (i.e.

an estimate of the conditional probability of the triplet given the occurrence of the first two words). These measures were z-scored relative to all the word triplets, in the case of triplet frequency after a log transformation. All triplets were given an expectancy score defined as two times the z-score of the conditional proportion of the third word given the first two plus the z-score of the log frequency of triplet occurrence. This measure provided an acceptable balance between the overall frequency of the triplet, and the conditional probability of the third word given the first two.

Triplets for expected repetition priming were selected from triplets with expectancy scores over three where the third word was a repetition of the first word. Similarly, potential triplets for the expected non-repetition condition were selected from the pool of triplets so that the third word was not a repetition of the first and the expectancy score was over three, after eliminating triplets containing words from the primary list already used in the expected non-repetition condition stimuli. Expectancy scores over three were rare, and only observed for 0.25% of the triplets investigated. Potential triplets for the non-expected repetition condition were defined from word pairs for which the first word never repeated after the second (e.g. 'cloth' never followed 'cloth of') and the first word had not yet been used in either of the other conditions. The final triplets for each condition were selected by hand. For each word triplet, a choice alternative was chosen from the remaining words in the primary word list with the same number of letters as the third word in the triplet and that never occurred as a completion of the first two words, e.g. for the triplet 'crowd go wild', the word 'clue' was assigned as an alternative. This resulted in 90 word triplets for each of the three types of priming, each with an alternative third word, plus 80 additional expected non-repetition triplets for

use in threshold trials, and 20 further triplets mixed between the conditions for use in practice trials. For target primed conditions, the expected and/or repeated third word in the triplet was presented as the target and the alternative completion was the foil. So 'crowd-go wild' as a target primed condition would present 'crowd-go' as a prime in the center of the screen, 'WILD' as a target also in the center and both 'WILD and 'CLUE' as choices on the left and right sides of the screen. In foil primed conditions, the expected and/or repeated third word was the foil choice, and the non-occurring alternative was the target, in this case 'clue'. Over all conditions, each word that came from the primary word list was only used once for each participant, with the obvious exception of within trial repetitions. Stimuli for neither primed conditions were created by swapping the choice alternatives (the third word and its alternative) between pairs of triplets within conditions. For example, the pairing of 'crowd-go wild/clue' with 'kiss-of death/south' would result in stimuli for two neither primed condition trials, one with a prime of 'crowd-go', a target of 'death' and choices of 'death' and 'south', similarly for 'kiss-of', a target of 'wild' and the choices 'wild' and 'clue'. Triplets were paired so that the choice alternatives for each triplet never occurred following the first two words of the paired triplet. For each participant, each pair of word triplets was randomly assigned to one of the three primed choice conditions (neither, target or foil). This procedure ensures that for each participant the same targets appeared with the same foils, regardless of the condition in which it appeared, while randomly varying triplets between primed choice conditions, and not repeating stimuli.

Unfortunately, this procedure did not completely control for base word frequency between the expected and non-expected repetition groups (expected repetitions occurred

more often, the mean log frequency of expected repetitions was 17.15, while the mean log frequency for non-expected repetitions was 16.25 ($t(166) = 5.09$, $p < .001$). Based on the results of Experiment 2.1, this would work to increase discounting for expected repetitions (the opposite effect of what was found).

Results

Data from 20 participants who were outliers (with median accuracy or RT greater than 1.5 the interquartile range from the median) for either the threshold or main portion of the experiment were excluded from the analyses, leaving 99 participants in Experiment 2.3. The proportion of correct trials for each condition in Experiment 2.3 is presented in Figure 2.5. There was a significant three way interaction between the priming type (non-expected repetition, expected non-repetition, and expected repetition), primed choice, and prime duration ($F(16, 1568) = 8.87$, $p < .001$). Replicating the previous repetition priming results, when targets were non-expected repetitions of prime words (e.g. 'cloth-of' priming 'cloth'), there is a pattern of priming influencing perception in the direction of the prime for shorter prime durations which reverses at longer prime durations. Target primed accuracy was greater than foil primed accuracy for prime durations between 17 and 150 ms ($t(98) >= 5.17$, $p < .001$). Similarly, there was reversal of the priming effect at 2,000 ms prime durations ($t(98) = -2.71$, $p = .047$). However, this reversal was smaller than in Experiment 2.1. Given the strong influence of orthography on priming (Huber et al., 2001), we attribute this to the change of case between prime and target. For primes that created an expectancy for the target, but were not repetitions (e.g. 'crowd-go' priming 'wild'), primes biased perception toward the expected, primed target significantly for prime durations of 50 to 400 ms ($t(98) = 4.23$, $p < .001$). There was no

negative priming for expected non-repetitions at 2,000 ms ($t(98) = 1.97$, $p = .257$). As is clear in Figure 2.5, the influence of the prime on perception was smaller, the time course of discounting was slower, and there was no target priming deficit compared to foil priming in this condition. When the target was both expected based on the prime, and also a repetition of the prime (e.g., 'walk-the' priming 'walk') the general pattern of priming appears roughly similar to the non-expected repetition condition, with positive priming for prime durations between 17 and 150 ms ($t(98) >= 8.38$, $p < .001$) but no significant negative priming at 2,000 ms ($t(98) = -1.75$, $p = .335$). Considering only non-expected repetition and expected repetition conditions verifies the difference between conditions. There is a significant interaction between type of priming and primed choice ($F(2, 196) = 3.95$, $p = .021$), due to higher target primed accuracy for expected repetition ($t(977) = -2.48$, $p = .013$). In the expected repetition condition, the effects of expectation found in the expected non-repetition roughly add to those of repetition seen for non-expected repetitions. As noted, repeated words in this condition were higher frequency than in the non-expected repetition condition. However, given the results of Experiment 2.1, this frequency effect would be in the opposite direction to that observed here. As further verification, we median-split the expected repetition results based on word frequency. Confirming that the differences in frequency did not produce this effect, there were similar magnitudes of priming for both high and low frequency primes (the latter of which were marginally *less* frequent than non-expected repetition stimuli, ($t(128) = 1.91$, $p = .058$), both of which were distinct from the non-expected repetition words).
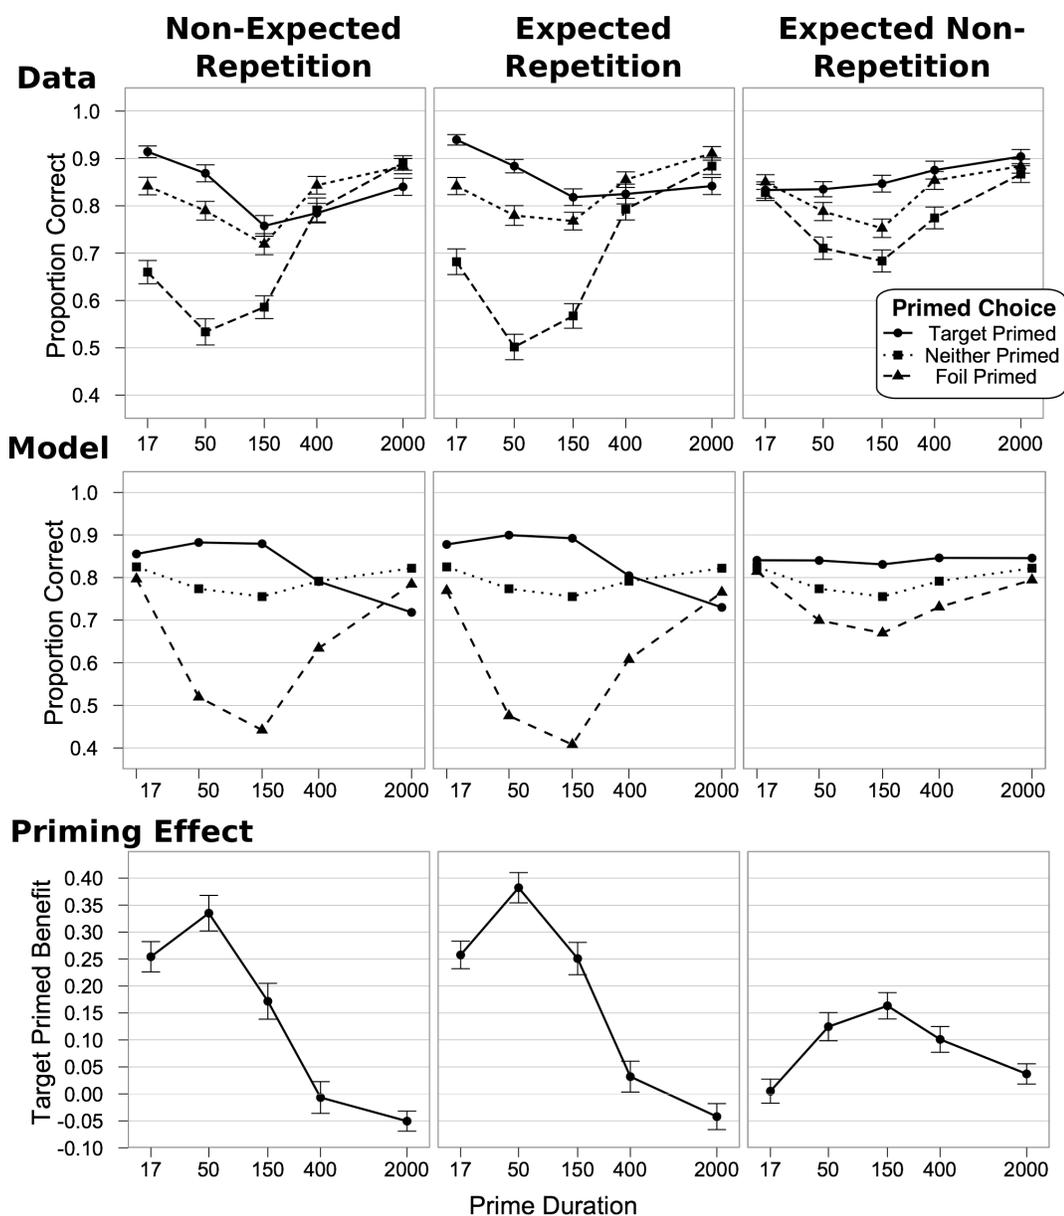
Figure 2.5. Experiment 2.3 accuracy by type of priming. Error bars are +/- 1 standard error of the mean or mean difference over participants. The top row shows the accuracy results from the experiment, the middle row shows the model results, and the bottom row shows the difference between the foil and target primed conditions in the behavioral data. Non-expected and expected repetition both show large positive priming effects for prime durations less than 400 ms, while priming effects for expected non-repetitions emerge and fade with increasing prime duration. Furthermore, expected repetition priming shows slightly larger priming effects as compared to non-expected repetition. Full details of the model are in the Chapter Appendix. Identical to Experiment 2.2, the critical parameters varied here were the proportion of the representation that is shared between the prime and the primed word, set to .2, and the orthographic to semantic connection strength of directed association, set to 1.4. Words in the expected repetition condition had two semantic representations, one corresponding to the semantics of the word when not repeated, the other corresponding to the word when it was repeated. This latter node was treated identically to a shared node for expected non-repetitions.

During threshold trials, the prime was always a non-repetition, however the prime words did sometimes create an expectancy to either the target word, foil word, or to neither choice. This expectancy relationship was found to produce a main effect on accuracy ($F(2, 236) = 140.14$, $p < .001$). When the target was expected from the prime, accuracy was higher than when neither choice or the foil was expected ($t(475) >= 11.85$, $p < .001$). When neither choice was expected, accuracy was higher than when the foil was expected ($t(475) >= 9.68$, $p < .001$). There was additionally an interaction between primed choice and target duration ($F(6, 708) = 10.93$, $p < .001$), reflecting a decreasing effect of primed choice with longer target durations, and a main effect of target duration ($F(3, 354) = 93.83$, $p < .001$); longer target durations improved accuracy.

Discussion

Experiment 2.3 demonstrated the impact of expected co-occurrence relationships between the prime and target in immediate repetition priming. Non-expected repetitions showed a similar pattern of priming to words in Experiment 2.1. The main differences were smaller overall priming effects, and less negative priming at longer prime durations. This difference is likely due to either the change in case between prime and target in Experiment 2.3, or possibly the more linguistic nature of the task due to primes being a phrase rather than a single word. Neither primed trials produced similar patterns of results as in past studies (Huber, 2008), with a dip in performance around 150 ms. Performance with neutral primes was not different between the expectancy conditions. This condition was logically identical between the different expectancy conditions. Expected non-repetitions, like "crowd-go" priming "WILD" produced similar waxing and waning positive priming effects to associative priming, demonstrating that expectations produce

positive priming effects. The effect of expected repetitions ('walk-the' as a prime for 'WALK') was generally similar to non-expected repetitions, but with slightly stronger priming effects in general. This is evidence that repetition causes the general pattern of facilitation and habituation, but expectation also factors into processing. In other words, these results show that the processing characteristics of these words follows a general pattern based on repetition, with the addition of learned adaptations to word co-occurrence.

In the model, the same mechanisms that were able to account for associative directionality effects in Experiment 2.2 successfully captured the results of Experiment 2.3. This modeling suggests that expectation in processing, either derived from association or co-occurrence, is enacted in enhanced feed-forward activation of mutual representations. To understand the general differences in the time course of priming in this experiment, future work will need to produce more constraining data for cases where the primes are visually different than the target and choices.

General Discussion

These three experiments answer questions about how the mind uses occurrence and co-occurrence statistics to identify and segregate items. Experiment 2.1 tested whether the pattern of assimilation and separation observed in the immediate repetition word priming is sensitive to occurrence statistics. The comparison between words and pronounceable non-words identified the effects of word frequency (pronounceable non-words being essentially 0 frequency words). The effects of orthographic frequency (in terms of groups of letters) were evaluated by comparing pronounceable non-words with

non-pronounceable non-words, with the latter defined as containing at least one three letter string not appearing in any word. Finally, experience with the occurrence of low level features was disrupted through comparison with inverted words. Although inverted words have semantic content like words, inverted letters are less frequently encountered than upright letters. The results of Experiment 2.1 show that the separation of perceptual items in the world is sensitive to frequency at multiple levels of processing. Priming for words peaked with shorter prime durations, and reversed more strongly than the other stimulus types. Differences between words and pronounceable non-words were only apparent in blocked presentation, suggesting that these priming effects are adaptable to experimental context. Reducing frequency further from pronounceable to non-pronounceable non-words resulted in shifts to priming effects as a function of prime duration. Finally, inverted words showed peak priming effects at the longest durations. This pattern of results was also reflected in the perceptual identification for each item. This is evidence that the time course of priming is influenced by the processing speed of individual items.

Experiment 2.2 tested the influence of co-occurrence through directional associations. We used unidirectional associates (where association is only from one to the other, not back) to prime either the target or the incorrect foil choice. The association direction between prime and primed word was randomly varied for each participant. As has been found in previous short and long term priming studies (Huber et al., 2001), associative priming was effective in both forward and backward directions, biasing responses to the prime. The effects of expectation direction were small but significant when considering multiple prime durations. Priming in a forward direction produced

more pronounced effects than backwards priming, with a later peak priming effect with 150 ms primes. These effects demonstrate that both non-directional and directional associative information is utilized to segregate perceptual events.

In the third experiment we tested if actual co-occurrence statistics produce effects similar to associative relationships. Furthermore, we tested a theory of the repetition effects as a high-level expectation for or against change by using primes that would be expected to repeat, e.g. 'walk the walk'. There were three types of priming: non-expected repetitions, expected repetitions, and expected non-repetitions. Non-expected repetitions replicated the general repetition priming pattern seen previously in Experiment 2.1, although the reversal to negative priming was not as strong. This is likely due to the case change between prime and target in Experiment 2.3 (Bowers, Vigliocco, & Haan, 1998; New & Grainger, 2011). It might also be due to the more linguistic nature of the primes in this experiment, which consisted of two words. Expected non-repetitions showed a similar priming pattern as seen in Experiment 2.2 for associative priming, exemplified by waxing and waning positive priming effects. Finally, the dynamics of expected repetitions were largely similar to non-expected repetitions, with increased positive priming effects. We attribute this pattern of results to the combination of strong effects of repetition with weaker expectation effects. In combination with previous results, this demonstrates that the dynamics of repetition priming is not a word-level expectation for or against repetition, but rather the result of the dynamics of orthographic repetition.

The results of Experiment 2.1 and Experiment 2.3 highlight the multiple levels of processing inherent in word processing dynamics. The observed discounting dynamics are not unique to word representations, linguistic processing, or word level expectations.

Rather they seem to occur in tandem at multiple hierarchical levels of representation.

This should be no surprise, as synaptic depression, the physiological component

hypothesized to underlie these effects in the brain, is found throughout cortex (Thomson

& West, 1993). These same assimilation and segregation principles, applied to different

types of representations, may underlie similar phenomena in cognitive psychology, for

example repetition blindness (Kanwisher, 1987), attentional blink (Raymond, Shapiro, &

Arnell, 1992), and inhibition of return (Posner & Cohen, 1984). Suggestions that

repetition blindness occur at retrieval rather than at the time of perception (Masson, 2004;

Whittlesea & Masson, 2005) are especially compatible with the conceptualization of

priming effects arising from source confusion. Habituation has been specifically

proposed to underlie inhibition of return effects (Dukewich, 2009), same difference

judgments (Davelaar, Tian, Huber, & Weidemann, 2011), and semantic satiation (Tian &

Huber, 2010). Furthermore similar ideas have been applied to in social psychology

(Irwin, Huber, & Winkielman, 2010).

Experiment 2.2 found priming effects with unidirectional associate pairs. The

relationship between associative, semantic, and co-occurrence relationships has a long

history (Fischler, 1977). Associative relationships are defined through free association

responses with the presentation of a cue, and thus many cognitive factors influence these

production-based norms, including semantic information, cue and target frequency, the

number of semantic neighbors, co-occurrence relationships, and hierarchical

relationships. Meta-analysis of the priming literature (Lucas, 2000) and analysis of the

linguistic environment (Lund & Burgess, 1995, 1996) have found that the majority of

associative priming is mediated through semantic relationships. The current work is

agnostic as to the particular nature of associative priming. Consequently, our model treats both associative priming and expectation based on co-occurrence statistics as arising through shared semantic representations.

Huber et al. (2001) classified the locus of immediate priming effects into two types: preferential and perceptual. Perceptual effects occur during the perception of the target, while preference effects are priming influences when the response is collected. Similar distinctions in terms of bias and perception have also been made (Ratcliff & McKoon, 1997). In this classification, repetition and orthographic priming effects are largely the result of preference, while associative priming had a larger perceptual component (Huber et al., 2001). Similarly, masked priming produces changes to the decision process, not perception (Masson & MacLeod, 1996). The critical condition for discriminating these effects primes both the target and foil simultaneously. Priming effects on only the perception of the target should result in a both primed benefit, at least for prime durations where positive priming is observed. This is because the influences of priming both the target and foil can only aid the extraction of perceptual information from the target. However, both primed deficits (as well as foil primed effects in two alternative forced choice experiments) measure the presence of preference effects due to the prime (or at least that preference effects overwhelm the influences of priming on the target presentation). The current experiments do not include a both primed condition, and thus cannot make this distinction. However, in the nROUSE model, any activation of the prime will change the target presentation activation. A target prime will pre-activate the target representation, possibly habituating it with long target durations. An unrelated or foil prime activates its own representation, causing competing inhibition. The

measurement of priming effects occurs in response to the choice presentation, which is in fact influenced by priming of the foil. The computational account of stimulus frequency suggests that the effects of stimulus frequency does not change the ratio of priming's influence on preference to perception. Rather, the key change with frequency is in the activation strength, and in turn processing speed, of stimuli. Unlike repetition priming, associative priming does show both primed benefits (Huber et al., 2001), indicating changes to the perceptual process of extracting information from the target. In the nROUSE model, benefits to target perception occur through both co-activation of a shared semantic node, and feedback from the shared node to the target. Thus, in our account of these effects, associative direction and expectation predict greater perceptual effects. This makes the prediction that for forward associations, both primed benefits should be slightly larger, and for expected repetitions, both primed deficits should be smaller or absent. Future research will investigate these predictions.

In Experiment 2.3 we found that expected repetitions did not increase discounting. Other stimuli may be able to produce more pronounced effects of high level expectation, possibly offsetting or slowing discounting. For example, repetition of higher semantic levels of processing, e.g., topics (Griffiths, Steyvers, & Tenenbaum, 2007), or with emotional or attention grabbing words (Zeelenberg, Wagenmakers, & Rotteveel, 2006), may have different processing characteristics than observed here for repetitions or adjacency statistics. As it stands in the model, all semantic processing is considered to occur at a single processing level. In reality, this would not be the case. Several layers of semantic representation, both lateral and hierarchical, would be necessary to encode the rich relationships present in the world (e.g. semantics, grammar, topics, pragmatics).

Interactions between perceptual events in priming and other paradigms are likely to occur at multiple levels. This process is simplified in the current model because of the relatively small amount of constraining data, in contrast to the data collected by the human mind over years of experience. Another obvious shortcoming of the model is the over-simplified representations. The postulation of a single processing unit to represent concepts, with single shared nodes to process overlapping semantics is clearly deficient (Marsolek, Schnyer, Deason, Ritchey, & Verfaellie, 2006). Even more distributed representations to capture complete semantic content are warranted (Masson, 1995; Seidenberg & McClelland, 1989). Unfortunately, creating enough stimuli for such studies would be difficult, if not impossible, and effect sizes for the influence of higher semantic content would be small. Another approach is to expand the neural model to include richer and more nuanced representations by implementing learning and training on naturalistic stimuli. This requires more of a machine learning approach, possibly facilitated by the reconnection with the LEABRA architecture (O'Reilly & Munakata, 2000).

The present approach is similar to that of rational analysis, which has highlighted the importance of environmental statistics to memory (Anderson, 1990; Chater & Oaksford, 1999). Similar ideas have informed the understanding of semantic knowledge and higher cognitive function (Anderson, 1991; Landauer, 1997; Steyvers, Griffiths, & Dennis, 2006). This paper demonstrates that similar thinking may be warranted for perceptual processes such as the segregation of information in time. Here, we have demonstrated that more frequently encountered items are processed more quickly and that perception is biased by associations and expectations from past stimuli. These implicit memories are learned from the statistics in the world and influence perception.

While there are dissociable elements of recognition memory and priming effects, (Wagner & Gabrieli, 1998), similar computational processes may operate on different representational levels. Regardless, priming should continue to be understood in reference to perception, as well as memory (Bowers, 1999).

This paradigm may provide interesting insight into mental processes, for example hemispheric differences and dissociation of processing deficits in psychological and neurological disorders. In lexical decision priming tasks, presentations to the LVF (right hemisphere) were more sensitive to indirect relationships and context, while the RVF (left hemisphere) were more sensitive to the dominant features (Atchley, Burgess, & Keeney, 1999). These results also suggested that the time course of priming effects appears different in the two hemispheres. The paradigm and model presented here would provide a critical measure of processing differences in the two hemispheres. In visual search tasks, the presence of autistic traits is correlated with enhancements of perceptual identification and processing speed (Brock, Xu, & Brooks, 2011). One aspect of attentional search is the speed of processing and discounting of distracters. The current task may be adaptable to provide separate measures of these constructs. Finally, long term priming has provided insights into the changes in memory that occur with aging (Fleischman & Gabrieli, 1998). The ability of the mind to adapt to different environmental conditions, and to process and segregate information reliably, is a critical component of cognition. Diagnostic paradigms adopted from those used here could be an essential measurement tool to view changes in cognition over the lifespan.

Conclusions

In summary, we have shown that word priming dynamics are informed by occurrence and co-occurrence statistics. These dynamics are context-sensitive; in a non-linguistic mixed context, words show similar dynamics to non-pronounceable non-words. Measured behaviorally, associative priming and priming based on co-occurrence relationships show similar influences on processing. Lastly, the effects of expectation and repetition roughly combine when a word is expected to repeat. Mechanisms producing these effects were successfully incorporated into the nROUSE model: occurrence through connection strength, non-directional co-occurrence relationships through shared semantic representations, and co-occurrence expectations through increased connection strength from the orthography eliciting the expectation to shared semantic representations. As a whole, these experiments have demonstrated that the dynamics of word processing are not globally fixed across all stimuli, but are sensitive to implicit memories concerning the statistics of the environment.

Appendix

This Appendix contains the computational details of node activation in the nROUSE model, default model parameters, the mapping of activation to proportion correct responses, and the details of applying the model to Experiments 1 through 3. The general structure of the model is described in the main text, and illustrated in Figure 2.2.

The dynamics of each individual node are similar to that of the LEABRA model (O'Reilly & Munakata, 2000), the structure, dynamics, and parameters of which are motivated by neurophysiology. Equations 2.1-2.3 are the equations for the change in node i of level n's membrane potential, v, amplitude, a, and output, o. The inputs to a node (both feed-forward from level $n - 1$, indexed by j, and feedback from level $n+1$, indexed by k) are scaled by their strengths, w, and work to drive the membrane potential, v, upward. Feedback connections, when present, are scaled by an additional value, F. Previous implementations of the model used a fixed feedback scalar. However, with a fixed feedback scalar, weak feed-forward connections might result in greater feedback activation than feed-forward activation. For this reason, we implemented feedback as a scaled proportion of the feed-forward connection. The summed activation of neurons in the same layer push membrane potential down in proportion to an inhibition parameter I, along with a fixed leak term, L. Finally the integration speed, s, which is unique to every layer, scales the change of potential.

Equation 2 is the main addition to the LEABRA framework, and separately updates the available resources of the node governing spike amplitude, a. The available spike amplitude is driven up by a recovery parameter, R, and down, as a proportion of the output, D. This incorporates short term synaptic depression into the model. Equation 3

determines the output, *o*, by thresholding the membrane potential (subtracting θ, and setting negative results to 0) and scaling it by *a*. The majority of parameters are fixed to the default values used in Figure 8 of Huber & O'Reilly (2003), including: $L = 0.15$, $I = 0.30$, $R = 0.22$, $D = .324$, $\theta = .15$, $F = 0.25$, and except where indicated otherwise, $w = 1$, $s_1 = 0.054$, $s_2 = 0.046$, and $s_3 = 0.015$. Many of these parameters in this model are fixed to default values used across multiple experiments (Huber & O'Reilly, 2003, Rieth & Huber, 2010) and are motivated by neurophysiology. See Huber & O'Reilly (2003) for a more detailed discussion of the model and its derivation.

$$\frac{\Delta v_i^n}{S_n} = (1 - v_i^n)\left(\sum_{\forall j} w_{ij} o_j^{n-1} + \sum_{\forall k} F^n w_{ik} o_k^{n+1}\right) - (v_i^n)\left(I \sum_{\forall j} o_j^n + L\right) \qquad (2.1)$$

$$\frac{\Delta a}{S_n} = R(1 - a) - Do \qquad (2.2)$$

$$o = \begin{cases} (v - \theta)a & v > \theta \\ 0 & v \leq \theta \end{cases} \qquad (2.3)$$

As applied to accuracy data, a mechanism is required to map activation onto forced choice response. For each millisecond of a particular condition trial, the appropriate input is provided to the visual nodes of the model, (e.g., a foil prime active for 400 time steps, followed by the target for 50 times steps, the mask for 450 time steps, and finally the target and foil choices). These inputs travel through the network according to the appropriate connections. In the experiment, the target and foil were presented until the participant responded. Here, the response time is modeled as identification latency,

the point at which sustained input of the each choice presentation will eventually stop increasing due to habituation. We consider identification latency to be a normal distribution around the time of each activation peak, with variance proportional by a parameter, *m,* to the squared identification latency. Note that this parameter does not change the dynamics of the model, and can only change the relative accuracy of the conditions. The difference in identification latency distributions for both output alternatives has a mean equal to the difference in the identification latency, and a variance equal to the sum of the variances of the distributions. The cumulative area in this distribution of differences below 0 gives the modeled proportion correct.

Accounting for the effects of occurrence and co-occurrence in Experiment 2.1 through 2.3 requires extensions to this model. The connection strength between the orthographic and semantic layer were used to reflect implicit knowledge of occurrence and co-occurrence.

In Experiment 2.1, all stimuli in a trial had the same general occurrence frequency, therefore the connection strength from the orthographic layer to the semantic layer was varied by stimulus type. This was the same strategy used to model the difference between priming for upright and inverted faces (Rieth & Huber, 2010). Although they contain the same visual information, they may be less fluent, and more difficult to process into a configural representation (Yin, 1969). Here, words are similarly considered to have stronger connection strengths than non-words, due to the ability to form rich semantic representations from input. Connection strengths between the second and third layer for pronounceable non-words (.9), non-pronounceable non-words (.75), and inverted words (.55) were used to fit the data from Experiment 2.1, simply by

varying the parameters by hand. All other connection strengths, including those for words, were fixed at 1. Given that we are considering orthographic representations to be affected by using non-pronounceable non-words and inverted words, we could have freed the layer 1 to layer 2 connection weight for these conditions. However, this would make the fit parameters difficult to compare between conditions. The variance multiplier, m, was set to .015 times the reciprocal of the connection strength used. Thus, representations with weaker connection strengths were also more noisy.

The effects of co-occurrence from association and expectation (Experiments 2.2 and 2.3) are both accounted for using the same addition to the model. General associative and expectation priming are expected to occur through shared semantic content, implemented in the model through the addition of a shared semantic node between the associated stimuli. Activation of this node represents the activation of the shared semantic concept, while the activation of the other layer semantic nodes represents unique semantic content from each word. This shared node receives input from both associates, but its output is scaled by a parameter, P, the proportion of each individual word's total representation attributed to the shared semantic meaning. Output from the appropriate unique nodes are scaled by $(1 - P)$. The end result is that the total representational strength for the concept as a whole (unique plus shared) is still 1. In both Experiments 2.2 and 2.3, the proportion of shared activation is governed by the parameter P, which is set to .2.

In Experiment 2.2, both the target and the foil have separate shared nodes with separate associates. This is because in Experiment 2.2, the target and foil were either both associative cues, or both associative targets. The direction of expectation and association

from a cue to an expected stimulus is accounted for through enhanced feed-forward connections from layer two of the cue stimulus to layer three of the shared representation. Given activation of the cue, this connection activated the shared representation in anticipation. In essence, in this model, the reason the pairs are directional relationships instead of bidirectional pairs is that one member of the pair activates the shared representation more strongly than the other. This higher connection strength was set to 1.4. For forward priming this is the connection from the layer 2 prime word to the layer 3 shared representation with the primed choice. Both choices had and activated shared nodes because in the experiment choices were both either associative cues or associative targets. For backward priming, this is the connection between the choices, and their respective shared nodes. The variance multiplier was set to 0.015. The effects of associative priming were reasonably well fit by the model. The difference in peak priming between forward and backward priming is subtle in the model. It is difficult to asses how problematic this is, as the true peaks for forward and backward priming could be at prime durations intermediate to those tested.

Experiment 2.3 was fit similarly to Experiment 2.2. Like expected non-repetitions, expected repetitions included a second semantic node, to distinguish the semantic content in the repeated phrase. Effects of expectation for both expected repetitions and non-repetitions were captured with a stronger feed-forward strength (1.4) identical to how the directional effects in Experiment 2.2 were modeled. Because the general pattern of priming was different in Experiment 2.3 (likely due to the change in case between primes and targets), we also varied the speeds of each layer to capture the general differences in Experiment 2.3 from Experiment 2.1. These values were 0.023,

0.06, and 0.01 for S_1, S_2, and S_3, respectively. The variance multiplier was again set to 0.015. The model is able to capture the effects of expectation well. However, the overall pattern of priming is not captured as reliably. Currently, the model does not distinguish between upper and lower case items, which may limit its ability to capture data with case switches.

References

Abbott, L. F., Varela, J., Sen, K., & Nelson, S. (1997). Synaptic depression and cortical gain control. *Science*, *275*(5297), 221-224.

Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Earlbaum.

Anderson, J. R. (1991). The adaptive nature of human categorization. *Psychological Review*, *98*(3), 409.

Arnon, I., & Snider, N. (2010). More than words: Frequency effects for multi-word phrases. *Journal of Memory and Language*, *62*(1), 67-82. doi:10.1016/j.jml.2009.09.005

Atchley, R. a, Burgess, C., & Keeney, M. (1999). The effect of time course and context on the facilitation of semantic features in the cerebral hemispheres. *Neuropsychology*, *13*(3), 389-403.

Balota, D. A., & Spieler, D. H. (1999). Word frequency, repetition, and lexicality effects in word recognition tasks: beyond measures of central tendency. *Journal of Experimental Psychology: General*, *128*(1), 32-55.

Bodner, G. E., & Dypvik, A. T. (2005). Masked priming of number judgments depends on prime validity and task. *Memory & Cognition*, *33*(1), 29-47.

Bodner, G. E., & Masson, M. E. J. (1997). Masked repetition priming of words and nonwords: Evidence for a nonlexical basis for priming. *Journal of Memory and Language*, *37*, 268-293.

Bodner, G. E., & Masson, M. E. J. (2001). Prime validity affects masked repetition priming: Evidence for an episodic resource account of priming. *Journal of Memory and Language*, *45*(4), 616–647.

Bodner, G. E., & Masson, M. E. J. (2003). Beyond spreading activation: An influence of relatedness proportion on masked semantic priming. *Psychonomic Bulletin & Review*, *10*(3), 645-652.

Bowers, J. S. (1999). Priming is not all bias: Commentary on Ratcliff and McKoon (1997). *Psychonomic Bulletin & Review*, *7*(1), 83-99.

Bowers, J. S., Vigliocco, G., & Haan, R. (1998). Orthographic, phonological, and articulatory contributions to masked letter and word priming. *Journal of Experimental Psychology. Human Perception and Performance*, *24*(6), 1705-1719.

Brants, T., & Franz, A. (2006). Web 1T 5-gram Version 1.

Brock, J., Xu, J. Y., & Brooks, K. R. (2011). Individual differences in visual search: relationship to autistic traits, discrimination thresholds, and speed of processing. *Perception*, *40*(6), 739-742. doi:10.1068/p6953

Brown, S. D., & Heathcote, A. (2005). Practice increases the efficiency of evidence accumulation in perceptual choice. *Journal of Experimental Psychology. Human Perception and Performance*, *31*(2), 289-298. doi:10.1037/0096-1523.31.2.289

Chater, N., & Oaksford, M. (1999). Ten years of the rational analysis of cognition. *Trends in Cognitive Sciences*, *3*(2), 57-65.

Coltheart, M. (1981). The MRC psycholinguistic database. *The Quarterly Journal of Experimental Psychology*, *33*(4), 497-505.

Davelaar, E. J., Tian, X., Huber, D. E., & Weidemann, C. T. (2011). A Habituation Account of Change Detection in same/different judgments. *Cognitive, Affective, & Behavioral Neuroscience*, *11*(4), 608-626.

Dukewich, K. R. (2009). Reconceptualizing inhibition of return as habituation of the orienting response. *Psychonomic Bulletin & Review*, *16*(2), 238-251. doi:10.3758/PBR.16.2.238

Duñabeitia, J. A., Carreiras, M., & Perea, M. (2008). Are coffee and toffee served in a cup? Ortho-phonologically mediated associative priming. *Quarterly Journal of Experimental Psychology*, *61*(12), 1861-1872. doi:10.1080/17470210701774283

Ellis, N. C. (2002). Frequency effects in language processing. *Studies in Second Language Acquisition*, *24*(02), 143-188.

Fischler, I. (1977). Semantic facilitation without association in a lexical decision task. *Journal of Experimental Psychology. Human Perception and Performance*, *3*(1), 335-339.

Fleischman, D. A., & Gabrieli, J. D. E. (1998). Repetition priming in normal aging and Alzheimer's disease: A review of findings and theories. *Psychology and Aging*, *13*(1), 88-119.

Gregg, V. (1976). Word frequency, recognition and recall. In J. Brown (Ed.), *Recall and recognition.* (pp. 183-216). New York: Wiley.

Griffiths, T. L., Steyvers, M., & Tenenbaum, J. B. (2007). Topics in semantic representation. *Psychological review*, *114*(2), 211-44. doi:10.1037/0033-295X.114.2.211

Grossberg, S. (1968). Some physiological and biochemical consequences of psychological postulates. *Proceedings of the National Academy of Sciences of the United States of America*, *60*(3), 758-765.

Grossberg, S. (1969). On the production and release of chemical transmitters and related topics in cellular control. *Journal of Theoretical Biology*, *22*(2), 325-364.

Huber, D. E. (2008). Immediate priming and cognitive aftereffects. *Journal of Experimental Psychology: General*, *137*(2), 324-347. doi:10.1037/0096-3445.137.2.324

Huber, D. E., & Cousineau, D. (2003). A race model of perceptual forced choice reaction time. In R. Alterman & D. Kirsh (Eds.), Proceedings of the 25th Annual Conference of the Cognitive Science Society (pp. 687-692).

Huber, D. E., & O'Reilly, R. (2003). Persistence and accommodation in short-term priming and other perceptual paradigms: temporal segregation through synaptic depression. *Cognitive Science*, *27*(3), 403-430. doi:10.1016/S0364-0213(03)00012-0

Huber, D. E., Shiffrin, R. M., Lyle, K. B., & Quach, R. (2002). Mechanisms of source confusion and discounting in short-term priming 2: Effects of prime similarity and target duration. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*(6), 1120-1136. doi:10.1037//0278-7393.28.6.1120

Huber, D. E., Shiffrin, R. M., Lyle, K. B., & Ruys, K. I. (2001). Perception and preference in short-term word priming. *Psychological Review*, *108*(1), 149-182. doi:10.1037//0033-295X.108.1.I49

Huber, D. E., Shiffrin, R. M., Quach, R., & Lyle, K. B. (2002). Mechanisms of source confusion and discounting in short-term priming: 1. Effects of prime duration and prime recognition. *Memory & cognition*, *30*(5), 745-757.

Irwin, K. R., Huber, D. E., & Winkielman, P. (2010). Automatic affective dynamics: an activation–habituation model of affective assimilation and contrast. *Modeling Machine Emotions for Realizing Intelligence*, 17-34.

Kanwisher, N. (1987). Repetition Blindness: Type recognition without token individuation. *Cognition*, *27*(2), 117-143.

Landauer, T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, *17*, 660-665.

Lucas, M. (2000). Semantic priming without association: a meta-analytic review. *Psychonomic Bulletin & Review*, *7*(4), 618-630.

Lund, K., & Burgess, C. (1995). Semantic and associative priming in high-dimensional semantic space. *Proceedings of the 17th Annual Conference of the Cognitive Science Society*, *36*(3), 402-407.

Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods*, *28*(2), 203-208.

Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. New York: Freeman.

Marsolek, C. J., Schnyer, D. M., Deason, R. G., Ritchey, M., & Verfaellie, M. (2006). Visual antipriming: Evidence for ongoing adjustments of superimposed visual object representations. *Cognitive, Affective, & Behavioral Neuroscience*, *6*(3), 163-174.

Masson, M. E. J. (1995). A distributed memory model of semantic priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *21*(1), 3-23.

Masson, M. E. J. (2004). When words collide: facilitation and interference in the report of repeated words from rapidly presented lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*(6), 1279-1289. doi:10.1037/0278-7393.30.6.1279

Masson, M. E. J., & MacLeod, C. M. (1996). Contributions of processing fluency to repetition effects in masked word identification. *Canadian Journal of Experimental Psychology*, *50*(1), 9-21.

Meyer, D. E., & Schvaneveldt, R. W. (1971). Facilitation in recognizing pairs of words: Evidence of a dependence between retrieval operations. *Journal of Experimental Psychology*, *90*(2), 227-234.

Meyer, D. E., Schvaneveldt, R. W., & Ruddy, M. G. (1974). Functions of graphemic and phonemic codes in visual word-recognition. *Memory & Cognition*, *2*(2), 309-321.

Neely, J. H. (1977). Semantic Priming and Retrieval from Lexical Memory : Roles of Inhibitionless Spreading Activation and Limited-Capacity Attention. *Building*, *106*(3), 226-254.

Nelson, A. B. (2009). Examining the Co-Evolution of Knowledge and Event Memory. (Doctoral dissertation, Indiana University)

Nelson, D., McEvoy, C., & Schreiber, T. (1999). The University of South Florida word association norms. *Tampa, FL, University of South Florida*.

New, B., & Grainger, J. (2011). On letter frequency effects. *Acta Psychologica*, 1-7. doi:10.1016/j.actpsy.2011.07.001

Norris, D., & Kinoshita, S. (2008). Perception as evidence accumulation and Bayesian inference: insights from masked priming. *Journal of Experimental Psychology: General*, *137*(3), 434-455. doi:10.1037/a0012799

O'Reilly, R., & Munakata, Y. (2000). *Computational explorations in cognitive neuroscience: Understanding the mind by simulating the brain*. Cambridge, Massachusetts: The MIT Press.

Pacton, S., & Perruchet, P. (2008). An attention-based associative account of adjacent and nonadjacent dependency learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(1), 80-96. doi:10.1037/0278-7393.34.1.80

Palmeri, T. J., Wong, A. C.-N., & Gauthier, I. (2004). Computational approaches to the development of perceptual expertise. *Trends in Cognitive Sciences*, *8*(8), 378-386. doi:10.1016/j.tics.2004.06.001

Pecher, D., Zeelenberg, R., & Raaijmakers, J. G. W. (1998). Does Pizza Prime Coin? Perceptual Priming in Lexical Decision and Pronunciation. *Journal of Memory and Language*, *38*(4), 401–418. doi:10.1006/jmla.1997.2557

Pecher, D., Zeelenberg, R., & Raaijmakers, J. G. W. (2002). Associative priming in a masked perceptual identification task: Evidence for automatic processes. *The Quarterly Journal of Experimental Psychology: Section A*, *55*(4), 1157–1173. doi:10.1080/0272498024400014

Pecher, D., Zeelenberg, R., & Wagenmakers, E.-J. M. (2005). Enemies and friends in the neighborhood: orthographic similarity effects in semantic categorization. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*(1), 121-128. doi:10.1037/0278-7393.31.1.121

Perruchet, P. (2004). The exploitation of distributional information in syllable processing. *Journal of Neurolinguistics*, *17*(2-3), 97-119. doi:10.1016/S0911-6044(03)00059-9

Posner, M. I., & Cohen, Y. (1984). Components of visual orienting. *Attention and performance X: Control of language processes*, *32*, 531–556.

Ratcliff, R., & McKoon, G. (1997). A counter model for implicit priming in perceptual word identification. *Psychological Review*, *104*(2), 319-343.

Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: an attentional blink?. *Journal of Experimental Psychology. Human Perception and Performance*, *18*(3), 849-860. doi:10.1037/0096-1523.18.3.849

Rieth, C. A., & Huber, D. E. (2010). Priming and habituation for faces: Individual differences and inversion effects. *Journal of Experimental Psychology. Human Perception and Performance*, *36*(3), 596-618. doi:10.1037/a0018737

Rumelhart, D. E., & McClelland, J. L. (1982). An interactive activation model of context effects in letter perception: II. The contextual enhancement effect and some tests and extensions of the model. *Psychological Review*, *89*(1), 60-94.

Scarborough, D., & Cortese, C. (1977). Frequency and repetition effects in lexical memory. *Journal of Experimental*, *3*(1), 1-17.

Scott, R. B., & Dienes, Z. (2008). The conscious, the unconscious, and familiarity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(5), 1264-1288. doi:10.1037/a0012943

Seidenberg, M. S., & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, *96*(4), 523-568.

Sereno, J. A. (1991). Graphemic, associative, and syntactic priming effects at a brief stimulus onset asynchrony in lexical decision and naming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*(3), 459-477. doi:10.1037/0278-7393.17.3.459

Sereno, S. C., Pacht, J. M., & Rayner, K. (1992). The effect of meaning frequency on processing lexically ambiguous words: Evidence from eye fixations. *Psychological Science*, *3*(5), 296-300.

Siyanova-Chanturia, A., Conklin, K., & van Heuven, W. J. B. (2011). Seeing a phrase "time and again" matters: the role of phrasal frequency in the processing of multiword sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *37*(3), 776-784. doi:10.1037/a0022531

Steyvers, M., Griffiths, T. L., & Dennis, S. (2006). Probabilistic inference in human semantic memory. *Trends in Cognitive Sciences*, *10*(7), 327-334. doi:10.1016/j.tics.2006.05.005

Thomson, A., & West, D. (1993). Fluctuations in pyramid-pyramid excitatory postsynaptic potentials modified by presynaptic firing pattern and postsynaptic membrane potential using paired intracellular recordings in rat neocortex. *Neuroscience*, *54*(2), 329–346.

Tian, X., & Huber, D. E. (2010). Testing an associative account of semantic satiation. *Cognitive Psychology*, *60*(4), 267-290. doi:10.1016/j.cogpsych.2010.01.003

Tsodyks, M. V., & Markram, H. (1997). The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proceedings of the National Academy of Sciences*, *94*(2), 719-723.

Wagner, A. D., & Gabrieli, J. D. E. (1998). On the relationship between recognition familiarity and perceptual fluency: Evidence for distinct mnemonic processes. *Acta Psychologica*, *98*(2-3), 211-230.

Weidemann, C. T., Huber, D. E., & Shiffrin, R. M. (2008). Prime diagnosticity in short-term repetition priming: is primed evidence discounted, even when it reliably indicates the correct answer? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *34*(2), 257-281. doi:10.1037/0278-7393.34.2.257

Weldon, M. S. S. (1993). The time course of perceptual and conceptual contributions to word fragment completion priming. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *19*(5), 1010-1023.

Whittlesea, B. W. A., & Masson, M. E. J. (2005). Repetition blindness in rapid lists: activation and inhibition versus construction and attribution. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *31*(1), 54-67. doi:10.1037/0278-7393.31.1.54

Yin, R. K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology*, *81*(1), 141-145.

Yue, X., & Tjan, B. (2006). What makes faces special? *Vision Research*, *46*(22), 3802-3811. doi:10.1016/j.visres.2006.06.017

Zeelenberg, R., Shiffrin, R. M., & Raaijmakers, J. G. W. (1999). Priming in a free association task as a function of association directionality. *Memory & Cognition*, *27*(6), 956-961.

Zeelenberg, R., Wagenmakers, E.-J. M., & Rotteveel, M. (2006). The impact of emotion on perception: bias or enhanced processing? *Psychological Science*, *17*(4), 287-91. doi:10.1111/j.1467-9280.2006.01700.x

CHAPTER 3: Adaptation to spatiotemporal contingencies in spatial cueing

Effectively foraging the perceptual environment requires a balance between the expected rewards of continued processing at the current focus and the potential reward of shifting to a new location. Extended maintenance of attention at its current location may be advantageous if there is reason to expect the appearance of new information there (e.g., watching a stoplight to quickly detect when it turns green). More typically, spatial attention is not maintained at a specific location and is instead automatically captured by the appearance of salient cues (Wolfe & Horowitz, 2004). The Posner cueing paradigm (e.g. Posner, Snyder, & Davidson, 1980) has provided a great deal of information regarding attentional capture and the dynamics of attentional shifts and maintenance. In a typical Posner cueing experiment, the participant is instructed to detect a target at one of two or more locations. Prior to the target's appearance, one of these locations is made salient using a cue, such as a flash, which may be valid (in the same location as the impending target), or invalid (at another location). In the current study, we investigated whether implicit learning of spatial temporal contingencies can change the typical pattern of results found with this task.

There are two key results found with the Posner cueing paradigm, and they are differentiated by the delay between the onset of the cue and the onset of the target. If the delay is short (e.g., 100 ms), reaction times (RT) to a target appearing at a validly cued location are faster than to targets appearing at other locations (Posner et al., 1980). This constitutes a 'positive cueing' effect and represents sensible foraging under the expectation that a 'salient' location (e.g., as signified by an abrupt onset) is likely to contain important, novel information. However, for longer delays (e.g., 350 ms), RTs to

142

targets at cued locations are often *slower* compared to targets appearing at uncued

locations (Posner & Cohen, 1984). This second result is commonly termed 'inhibition of

return', and has been extensively studied over the last several decades. Because this name

presupposes a cause that has been discounted in the literature, we refer to it simply as

'negative cueing' (Berlucchi, 2006). In terms of foraging behavior, negative cueing is

also sensible—once all the information has been gained from a salient location, there is

no reason to expect anything new at that location in the immediate future.

Negative cueing has been found in a wide range of situations. It occurs with both

saccadic eye movements (Rafal, Calabresi, Brennan, & Sciolto, 1989) and when the eyes

remain fixated while spatial attention covertly shifts (Posner & Cohen, 1984). Not only

does negative cueing occur in simple detection tasks, but it can also occur with target

discrimination tasks (Cheal, Chastain, & Lyon, 1998; Lupiáñez, Milán, Tornay, Madrid,

& Tudela, 1997; Pratt, Kingstone, & Khoe, 1997). Negative cueing can be initiated not

only by exogenous cues, like a flash at a specific location, but also endogenous cues, such

as expectation of cue validity (Lupiáñez et al., 2004), or a central arrow or verbal cue in

object directed attention (Weger, Abrams, Law, & Pratt, 2008). Furthermore, negative

cueing can arise from an exogenous auditory cue to a visual target (Spence & Driver,

1998), and from an auditory cue to an auditory target (Mondor, Breau, & Milliken, 1998).

The mechanisms of negative cueing effects across these different tasks likely differ

depending on the nature of the cues and the motor effectors involved. Nevertheless, the

ubiquitous nature of negative cueing suggests that it arises from a common computational

constraint related to information processing. In support of this claim, increasing the

perceptual processing demands of the task by changing the task from detection to

discrimination (Lupiáñez et al., 1997), or by increasing the difficulty of target detection (Castel, Pratt, & Drummond, 2005), has been found to delay the onset of negative cueing. Furthermore, negative cueing effects are not observed in rapidly changing visual environments (Wang, Zhang, & Klein, 2010).

In terms of foraging behavior and information processing, shifts from positive to negative cueing as a function of delay between cue and target may reflect learned responses to the spatiotemporal contingencies of the typical visual environment. A salient cue indicates that something has just changed in the environment, signaling that there is something new to perceive at that location (i.e., positive cueing immediately following the cue). However, once the information at that location has been collected, it is unlikely that something new will subsequently appear (i.e., because nothing new is expected at the previously cued location, negative cueing is observed). We are not concerned with the mechanistic processes underlying spatial cueing, but rather the visual environment that may have sculpted those processes. These ideas are similar to the 'foraging facilitator' hypothesis (Klein, 2000; Maylor, 1985; Posner & Cohen, 1984; Taylor & Klein, 1998) except we suppose that these effects are a rational adaptation to the statistical regularities of the visual environment. Therefore, we refer to this theory as the 'rational forager' hypothesis. According to the rational forager hypothesis, negative cueing occurs for experimental conditions that are in general incongruent with the spatiotemporal contingencies that exist in the visual environment (e.g., information appearing at a cued location after a long delay). In a typical spatial cueing experiment, the location and timing of the cue is fully counterbalanced such that all delays occur equally often for all locations. However, we hypothesize that the more typical visual environment primarily

consists of short, but not long, delays between salient visual events and the appearance of novel information.

According to the rational forager hypothesis, if the spatiotemporal contingencies of the visual environment change, then spatial cueing behavior should change in a similar manner. To test this claim, we performed experiments that trained participants on different spatiotemporal contingencies and then assessed whether learning generalized to spatial cueing in an unstructured environment (i.e., one in which all cue to target delays are tested equally often). We did not inform participants of these manipulations and most participants did not realize that the spatiotemporal contingencies were manipulated. Therefore it appears that these contingencies are implicitly learned. Previous studies found evidence of spatial contingency learning at either long or short target delays. However, these studies only examined behavior while participants were still experiencing the newly learned contingencies, often with explicit awareness because the manipulation entailed an obvious bias to specific locations (e.g., cues typically indicated, or did not indicate the target location). For example, in an anti-cueing paradigm, participants were informed that targets would appear in locations opposite to the cue. Under these conditions the positive cueing effects typically seen with short delays were reversed (Warner, Juola, & Koshino, 1990). Similarly, the negative cueing effects typically seen with longer target delays also adapts with experience. For example, negative cueing effects gradually decreased the longer participants were tested in an unstructured environment in which targets appeared equally often in all locations (Weaver, Lupiáñez, & Watson, 1998). Finally, in a saccadic inhibition of return study (i.e., using eye movement latencies rather than RT) the magnitude of negative cueing changed as

expected, depending on whether targets at cued locations were likely, unlikely, or equally probable (Farrell, Ludwig, Ellis, & Gilchrist, 2010).

If the rational forager hypothesis is the correct explanation of both positive and negative cueing effects, then spatiotemporal contingencies in the everyday structured visual environment must transfer to the laboratory, causing the typical effects seen even though targets are presented equally often at all locations. However, the studies reviewed above only examined adaptation to ongoing contingencies rather than investigating whether a newly learned contingency generalizes to subsequent testing. This is a critical distinction; in the presence of an ongoing contingency, observed adaptation effects might only reflect what was experienced on the last trial. Demonstrating the plausibility of this alternative explanation, last trial effects are large in negative cueing experiments that test all delays and positions equally (Dodd & Pratt, 2007). Thus, cueing behavior may appear to be globally adaptive because the last trial reflects the local statistics of the different trial types. It has not yet been established that spatiotemporal contingency learning lasts to the degree that is required under the rational forager hypothesis. One exception comes from studies examining the role of reading direction. For English readers, the negative cueing effect is larger when the cue and target appear on the left of the screen whereas for readers of Hebrew, the opposite is true (Spalek, 2004; Spalek & Hammad, 2005). This is a compelling result and yet it remains unclear whether generalization from a previous contingency can result from a relatively short training period (e.g., in a single training session), or whether it only occurs with a life-time of experience, such as with adult readers. Furthermore, a limitation of all the studies reviewed above is that they only

examined adaptation of positive cueing with short delays or negative cueing with long delays, rather than assessing whether both effects are jointly adaptive.

To address these limitations, the current experiment used complicated contingencies that required combinations of location and delay, followed by testing where all target locations and delays were equally likely. An advantage of including delays appropriate to both positive and negative cueing is that it becomes difficult to explicitly detect that anything has been manipulated. By including both short and long delays with opposing contingencies, it becomes possible to run an experiment in which the target is equally likely to appear at the cued and uncued locations even though there is a specific spatiotemporal contingency that can be implicitly learned.

Experiment 3.1

We trained three groups of participants in a spatial cueing task with each group receiving a different spatiotemporal contingency. All three contingencies entailed an equal proportions of valid and invalid cue trials (i.e., the target was equally likely to appear on the cued side as the uncued side throughout the experiment). During training, one group received only trials where the target appeared on the same side as the cue after a short cue-target delay or on the opposite side from the cue after a long delay (short valid/long invalid cues). Another group was trained with the opposite trial types: the target only appeared on the opposite side from the cue after a short delay or on the same side as the cue after a long delay (short invalid/long valid cues). A third group was trained with equal proportions of all four trial types. Note that for all three training conditions, neither location nor timing was independently predictive of the target--

learning could occur only if spatial attention is sensitive to combinations of location and delay. After training, generalization was tested in the same manner for all three groups, with equal proportions of valid and invalid trials at a range of delays, including both trained and untrained delays.

Methods

　　Participants. A total of 58 participants completed the experiment. Each participant was randomly assigned to one of three possible training conditions: 20 to the short invalid/long valid training, 19 to the short valid/long invalid training, and 19 to the all cueing conditions training. The study was approved by the UCSD IRB and all participants gave informed consent.
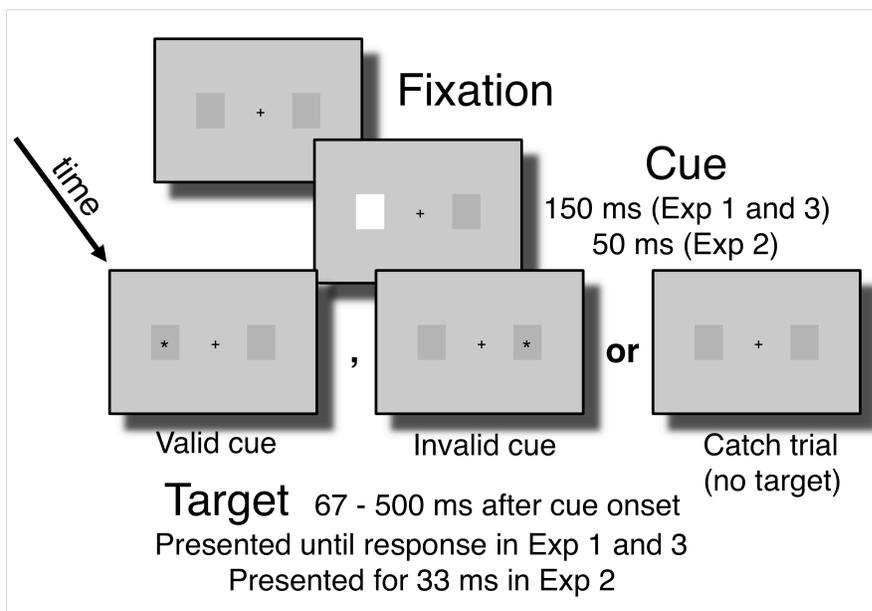

Figure 3.1. The basic trial sequence for Experiments in Chapter 3. For delays that were longer than the cue duration, a display identical to the initial fixation screen was presented between the offset of the cue and the onset of the target. For target delays that were shorter than the cue duration, the target appeared on top of the cue.

Stimuli and Apparatus. The experiment was conducted using computers with a 60 Hz refresh rate and 640 x 480 display resolution. The background of the display was a neutral gray. The two potential target location boxes were defined by lighter gray 48 by 56 pixel regions located equally spaced from the center fixation point (separated by approximately 16 degrees of visual angle). Both the fixation and the two potential target location boxes were present for the entire trial. Cues were white rectangles (40 by 48 pixels) presented in the center of one of the potential target location boxes. Targets were black asterisks presented in the center of one the potential target location boxes.

Procedure. Figure 3.1 presents the basic trial sequence. Each trial started with a fixation point in the center of the screen. Participants were instructed to maintain their gaze at the fixation and to respond (press a key) as quickly as possible after an asterisk appeared on the screen. Prior to the target presentation, the cue appeared at one of the potential target locations for 150 ms. Cues were equally likely appear at the right or left potential target location boxes. On each trial there was a 50/50 chance that the cue validly indicated the target location, which also appeared at either the left or right location with equal probability. During final testing, the target followed the appearance of the cue after a delay ranging from 67 to 500 ms. Note that for trials with 67 or 100 ms target delays, the cue was still present when the target appeared such that the target appeared on top of the cue. To reduce anticipatory responses, no target was presented on some trials. Participants were instructed not to respond to these catch trials. Following each trial, accuracy and RT feedback were briefly reported to the participant to promote interest and motivation for the task. If the participant made a response prior to the target on a target present trial, or responded on a target absent trial, or if a response was not made within

1,500 ms of the target's appearance, the trial was recorded as incorrect. If participants

responded in the absence of an actual target, either prior to the targets appearance or on a

target absent trial, a brief message appeared on screen to encourage accuracy. Trials were

divided into four blocks. The first two blocks each consisted of 176 training trials. The

second two blocks were each 200 testing trials. Participants took two minute video game

breaks between blocks. Participants' correct RTs and target detection accuracy were the

dependent measures. Participants were instructed to keep their eyes fixated on the center

position although eye movements were not monitored. We do not view this as serious

limitation considering that the rational forager hypothesis supposes that adaptation to

spatiotemporal contingencies should occur for both covert shifts of attention and for

shifts of attention that involve eye movements.

The experiment took place in two phases: a training phase and a test phase. This

phase distinction was not disclosed to participants and, as reported in debriefing, was

rarely noticed. There were three different randomly assigned groups of participants. The

groups differed only in the types of trials occurring in the training phase. For all training

groups, only 100 and 350 ms target delays occurred during training. However, the

proportion of valid and invalid cues at each target delay was manipulated between

groups. Two experimental groups were trained in environments with a fully deterministic

spatiotemporal relationship between the validity of the cue and the target delay. For

participants in the short invalid/long valid training group, training targets appeared at the

uncued location after a short (100 ms) delay or at the cued location after a long (350 ms)

delay. For participants in the short valid/long invalid training group, training targets

appeared at the cued location after a short (100 ms) delay or at the uncued location after a

long delay (350 ms). The third training group served as a control. For this all condition

training group, all four cueing conditions were equally likely (100 ms delay valid cue,

100 ms delay invalid cue, 350 ms delay valid cue, or 350 ms delay invalid cue). All trials

were randomly ordered. Each training phase consisted of 288 target present trials and 72

target absent catch trials. The target present trials were divided evenly between the

trained target delay and cueing conditions.

The testing phase was identical for all three training groups. Participants were

tested with 67, 100, 200, 350, or 500 ms target delays for both valid and invalid cues. In

the testing phase, all target delay and cueing conditions were equally likely. There were

32 target present trials at each combination of target delay and cue type (valid and

invalid) and 80 target absent trials. At the end of the experiment, participants were asked

if they noticed any relationship between the cue and the appearance of the target, and if

so to describe the relationship. Any participant who was able to describe any aspect of the

training relationship or reported noticing a difference between the training and test phase

of the experiment was eliminated from the analyses.

Results

A total of 11 participants were excluded from further analyses because they either

indicated insight into the nature of training (n=3) or had an average accuracy lower than

1.5 times the inter-quartile range from the median in either training or test phases (n=8).

Trials with RTs of less than 100 ms were also excluded from further analysis (1.2% of

trials over all participants). Data from 15 participants in the short invalid/long valid

training, 15 in the short valid/long invalid training, and 17 trained with all cueing

conditions remained. Participants responded incorrectly (i.e. before the appearance of the

target or on target absent trials) on an average of 1.9% of training trials and 2.8% of test

trials. All subsequent analyses only report results from the training phase, which was

identical for all three groups. Thus, our focus was not on whether participants adapted to

a specific training regime while that regime was still in effect, but rather whether there

was a carryover effect of adaptation from training upon subsequent testing.
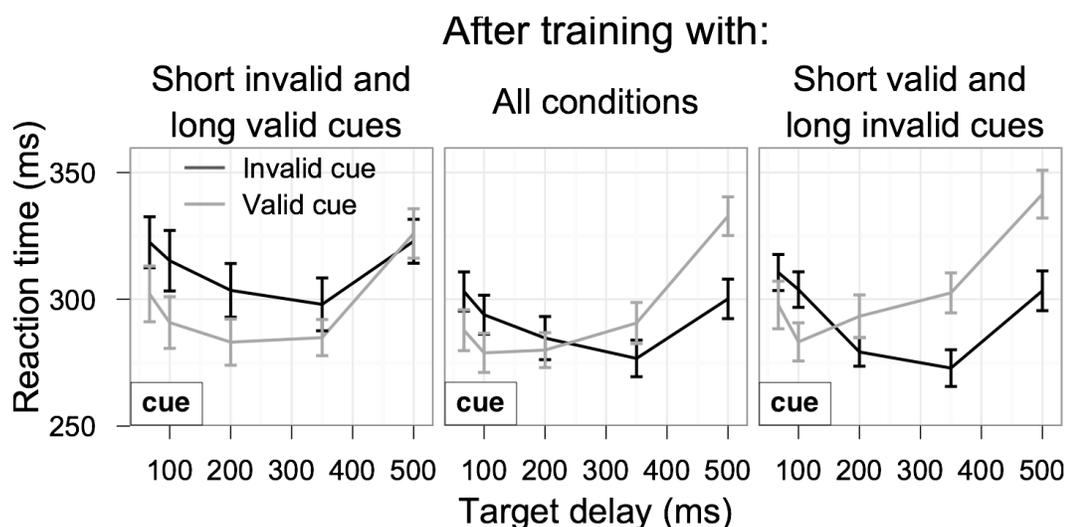

Figure 3.2. Average median correct RT across trial types for Experiment 3.1. These RTs
are from test phase of the experiment, which was identical for each group. The box labeled
'cue' in the lower left corner of each plot indicates the target delays for which the cue
remained on screen during the presentation of the target. Error bars are +/- 1 standard error
of the mean of participant medians.

First we consider accuracy. Incorrect anticipatory responses occurred more often

for longer target delay trials (main effect of target delay, $F(4,176) = 55.75$, $p < .001$; t-

tests comparing accuracy at the 350 ms delay to shorter delays, all $t(93) > 4.74$, $p < .001$;

t-tests comparing accuracy at the 500 ms delay to all other delays, all $t(93) > 7.61$, $p <

.001$). There were no differences in accuracy between training conditions ($F(2,44) = 1.35$,

$p = .27$), nor interactions with the training conditions (All $F <=1.29$, $p >= .251$).

Correct median RTs were averaged across participants and are presented in Figure 3.2. Statistical analysis with a repeated measures ANOVA revealed an interaction between training type, target delay, and target location ($F(8,176) = 4.72, p < .001$). Results for participants in the all condition training replicated the typical positive and negative cueing effects (delay by target location interaction, $F(4,64) = 34.12, p < .001$) found in previous studies. We used planned comparisons to test for negative cueing at long target delays (350 and 500 ms) and positive cueing at short target delays (67 and 100 ms). Indicating negative cueing, RT was faster to invalid cue trials at 350 and 500 ms target delays (350 ms $t(16) = -3.09, p = .007$; 500 ms, $t(16) = -5.74, p < .001$). Positive cueing was found at 100 and 67 ms target delays (100 ms $t(16) <= 3.92, p = .001$; 67 ms $t(16) <= 3.50, p = .003$).

For both experimental groups there were interactions between target delay and cueing condition (both $F(4,56) >= 5.54, p < =.001$). At long target delays these interactions were consistent with training. Participants in the short valid/long invalid training group were trained that following a long delay, the location of the cue was likely to be an invalid indicator of the target location. Consistent with this training, negative cueing was observed at 350 and 500 ms target delays (350 ms $t(14) = -6.98, p < .001$; 500 ms, $t(14) = -7.75, p < .001$) during the test phase. Furthermore, this negative cueing was larger than that for the control training group. When comparing the short valid/long invalid and all conditions training groups, the three way interaction between training group, cueing and target delay was significant ($F(4,120) = 2.99, p = .022$). Additionally, the short valid/long invalid had greater negative cueing at the 350 ms target delay (training by cueing interaction at 350 ms for those groups $F(1,30) = 6.29, p = .018$).

Unlike these results, participants in the short invalid/long valid training group displayed no significant negative cueing effects. In fact, at a 350 target delay there was a significant positive cueing effect (350 ms, $t(14) = 2.76$, $p = .015$). These positive cueing effects were also present at the 200 ms target delay (200 ms $t(14) = 3.91$, $p = .002$) . In comparison, no effect was observed at the 200 ms target delay for the all conditions training group ($t(16) = 1.00$, $p = .333$), and negative cueing was observed for the short valid/long invalid group ($t(14) = 2.24$, $p = .042$).

In contrast to the robust training effects at long delays, there were no training effects at short delays. For the 67 and 100 ms target delay conditions there was no main effect of training and no interactions of training with delay and/or cueing condition (all $F(2,44) <= 1.08$, $p >= .349$). Nevertheless, there were cueing effects for these delays: Positive effects of cueing were found for both experimental training conditions (short valid/long invalid, 67 ms $t(14) <= -4.36$, $p < .001$; 100 ms $t(14) <= -2.77$, $p = .015$; short invalid/long valid, 67 ms $t(14) <= -5.03$, $p < .001$, 100 ms $t(14) <= -5.18$, $p < .001$).

Discussion

In Experiment 3.1, participants learned spatiotemporal contingencies without explicit awareness and this learning generalized to a subsequent testing session that removed these contingencies. Compared to participants trained with both valid and invalid cues at long delays, training with only invalid cues at long delays increased the size of the negative cueing effect. Conversely, participants trained with only valid cues at long delays not only decreased the negative cueing effect but actually produced a reversal: a *positive* effect of cueing was observed at long cue-target delays. For both experimental training groups the changes in behavior at the trained 350 ms target delay

generalized to other long target delays. Unlike prior work that examined adaptation to spatial cueing contingencies, these results cannot be explained by priming from the last few trials—prior studies examined performance during training whereas we examined generalization to subsequent testing that removed the spatiotemporal contingency.

This generalization from training confirms a key assumption of the rational forager hypothesis, providing evidence that the local spatiotemporal cueing statistics of one environment can affect a subsequent test environment that alters the statistics by making cue position and timing non-diagnostic. This supports the claim that the typical pattern of results reflects adaptation to the everyday visual environment. However, one aspect of the rational forager hypothesis was not confirmed—the results of Experiment 3.1 suggest that generalization from training only exists for long but not short cue-target delays. The next two experiments tested two potential explanations of this difference.

Experiment 3.2

In Experiment 3.1, training effects were observed at long but not short target delays. However, only following short target delays did the cue overlap with the presentation of the target. For valid cue trials this overlap creates greater visual contrast for the target, enhancing processing of targets at the sensory level. If targets are easier to detect at the cued locations for short target delays, this might reduce the need for shifts of spatial attention because with higher visual contrast, spatial attention is less important for rapid target detection. If there is less demand for shifts of spatial attention, then there may be less learning of the cue contingency for short delays. To test this possibility, Experiment 3.2 replicated Experiment 3.1 while using a shorter duration cue, such that

the cue never overlapped with the target. Thus, there was no cue overlap for both short and long cue-target delays. Unlike Experiment 3.1, Experiment 3.2 did not test for generalization of training to other delays that did not appear during training. Instead, the question asked in Experiment 3.2 was whether training effects existed for both the trained short and long cue-target delay when there was no cue overlap for either delay. The delays used in Experiment 3.2 were chosen to represent a prototypical short and long delay and were set at 83 and 400 ms, respectively.

Methods

The procedures and conditions used in Experiment 3.2 were identical to those of Experiment 3.1 except where stated otherwise.

Participants. A total of 30 participants completed the experiment. Each participant was randomly assigned to one of three training conditions: 9 to the short invalid/long valid training, 11 to the short valid/long invalid training, and 10 to the all cueing conditions training.

Procedure. Experiment 3.2 used 50 ms cues, and target delays of 83 and 400 ms for both the three training conditions as well as for testing. With these timings, the presentation of cue and target did not overlap. Furthermore, in Experiment 3.2 the target was a brief 33 ms flash. A total of 360 target present training trials were divided into the two or four trial types, plus 72 target absent trials. The testing phase of the experiment consisted 40 trials of each of the 4 combinations of validity and target delay and 32 catch trials. To reduce the effects of anticipation with only a two possible target delays, target presentations were randomly jittered from the target delay condition by single screen refreshes (17 ms), so true target delays varied from 67-100 ms and 383-417 ms. This was

done during both the training trials as well as during test trials. The experiment was

conducted in three blocks. The first two were training blocks of 144 trials, and the last

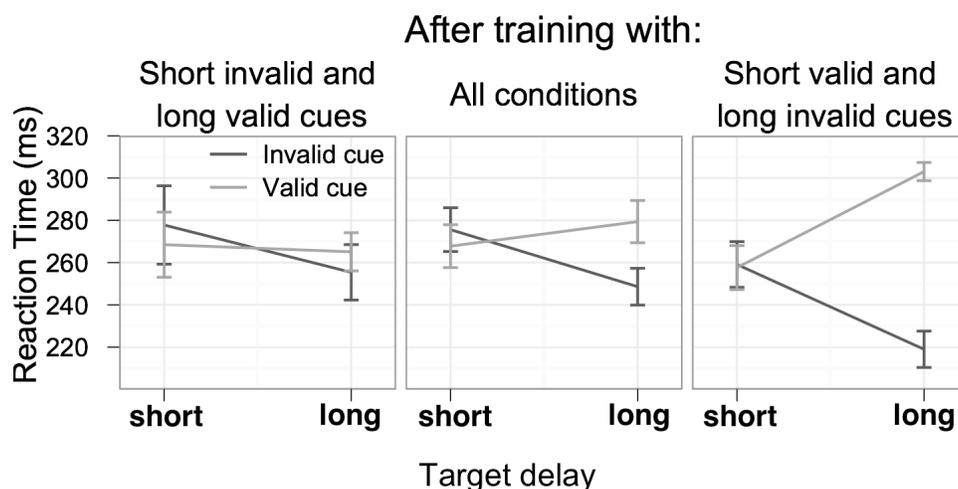consisted of 44 training trials followed by 192 trials in the testing phase.

Figure 3.3. Average median correct RT for Experiment 3.2. These data are from the testing phase of the experiment, which is identical for each group. In this experiment the cue and target never overlapped. Target delays were jittered between from 67 to 100 ms for short target delays, and between 383 and 417 for long target delays. Error bars are +/- 1 standard error of the mean of participant medians.

Results

As in Experiment 3.1, participants were deemed to be outliers and eliminated

from further analyses if they were beyond 1.5 times the IQR from the median accuracy or

RT for either the training or test phase (n = 6), or gave any indication of noticing the

training manipulation (n = 2). Overall error rates were 3.5% for the training phase, and

2.6% in the test phase. Figure 3.3 presents the average of each participant's median RT

for each condition during the test phase. In the testing phase of the experiment, there was

a three way interaction between training group, cue validity and target delay ($F(2,19) =$

14.94, $p = .001$), such that the negative cueing effect was absent in the short invalid/long

valid training group ($t(4) = -1.13$, $p = .32$), present in the all cueing conditions training

group ($t(8) = -4.63$, $p = .002$), and largest for those in the short valid/long invalid training

group ($t(7) = -11.81$, $p < .001$). There was no effect of training at short target delays. In

fact, there were no positive cueing effects at all in the test phase of the experiment (short

invalid/long valid training, $t(4) = 1.16$, $p = .310$; all cueing conditions, $t(8) = 1.32$, $p =$

.224; short valid/long invalid training, $t(7) = 0.24$, $p = .819$).

Discussion

Replicating Experiment 3.1, training effects were present only for the long target

delay, indicating that the lack of a training effect for the short delay in Experiment 3.1

was not a consequence of the temporal overlap between the target and cue for the short

delay. However, there was one interesting difference between Experiment 3.1 and

Experiment 3.2: Experiment 3.2, which eliminated overlap of the cue and target in the

short delay condition, failed to reveal *any* significant positive cueing effects for the short

delay condition. Based on this result, it may be tempting to hypothesize that positive

cueing effects with short delays may in general reflect visual overlap of cue and target

and thus enhanced visual contrast only for valid trials. However, this hypothesis would be

inconsistent with other experiments finding positive cueing effects for short delays in the

absence of temporal overlap of the cue and target (Maruff, Yucel, Danckert, Stuart, &

Currie, 1999; McAuliffe & Pratt, 2005; Pratt, Hillis, & Gold, 2001). Instead, it may be

that the positive cueing effect of Experiment 3.1 for short delays reflected the

combination of a true attentional effect combined with a visual contrast effect—by

eliminating the visual contrast effect, perhaps Experiment 3.2 contained insufficient

power to reveal the smaller attentional effect. Regardless of the explanation for this

difference between Experiment 3.1 and Experiment 3.2, it is clear that eliminating cue overlap failed to enable training effects for the short target delays. An alternative explanation of the failure to observe training effects for short target delays is tested in Experiment 3.3.

Experiment 3.3

Unlike previous experiments that found training effects with short delays (Warner et al., 1990), Experiments 1 and 2 used a mixture of both short and long target delays. This was done to examine whether the attentional system is sensitive to combinations of delay and location (i.e., spatiotemporal contingencies). The results from the long target delay conditions suggest that this is indeed the case, but the lack of effects with short target delays are puzzling. One possibility is that attention is sensitive to spatiotemporal contingencies for both short and long delays, but that it lacks the capacity to take advantage of these contingencies in a situation that presents both delays in an intermixed fashion. If participants are unable or unwilling to make two attention shifts during a trial (i.e., a first shift under the assumption that the trial is a short delay trial, but then a second shift if the trial turns out to be a long delay trial), then attention would be unable to utilize both the short and long delay contingencies even if both were implicitly learned. If attention only makes one shift of attention, the optimal strategy is to wait until such time as it can be determined that the trial is a long delay trial. In contrast, if that one shift is made as soon as the target appears, there is still a 50/50 chance that the trial be valid versus invalid, depending on whether the target delay is short versus long. For the experimental training regimes of Experiments 1 and 2, if that one shift is made only after

such a delay that the trial is known to be a long delay trial, then the shift can be made to the target location with 100% certainty.

To test this alternative, Experiment 3.3 used training that only included short target delays. Thus, if the system can learn spatiotemporal contingencies even for short target delays, and if the participant only experiences short delays, the optimal strategy is to shift attention as soon as the cue appears. However, running such an experiment poses a problem—with only a single target delay, it is likely that participants will become explicitly aware of the trained contingency if that contingency is 100% reliable. Therefore, unlike Experiments 1 and 2, the validity of the spatiotemporal contingency in Experiment 3.3 was probabilistic rather than deterministic. This is likely to reduce the magnitude of any training effects, although this change is necessary to disguise the nature of the training. Analogous to Experiments 1 and 2, Experiment 3.3 trained participants that the cue was likely to be valid, likely to be invalid, or that both cue types were equally likely. Participants were then tested with equal proportions of valid and invalid cues at three short target delays, including the trained delay. At first glance, this may appear to be a replication of the anti-cueing paradigm (Warner et al., 1990). However, unlike that paradigm: 1) we did not explicitly inform participants about these contingencies; 2) we tested generalization of learning in an environment without contingences; and 3) training was completed in a single session.

Methods

The procedures and conditions used in Experiment 3.3 were identical to those of Experiment 3.1 except where otherwise stated.

Participants. A total of 105 participants completed the experiment. Each participant was randomly assigned to one of three training conditions: 35 to invalid common training, 34 to valid common training, and 36 to all cueing conditions training.

Procedure. All target-present training trials in Experiment 3.3 used a 100 ms target delay. Participants in the all cueing conditions training group completed 192 valid cue trials, 192 invalid cue trials, and 96 target absent trials. The valid and invalid common training conditions consisted of 336 trials of the common target location condition, 48 trials of the complementary target condition and 96 target absent trials. For example, a participant in the valid common condition completed 336 trials where the target appeared at the same location as the cue, 48 trials where the target appeared in the opposite location from the cue, and 96 trials with a cue but no target. The testing phase of the experiment was identical for all groups and consisted of 48 target absent trials and 32 trials at each of six conditions: target delays of 67, 100, and 200 ms following either valid or invalid cue presentations. There were three blocks in the experiment, two training and one testing.

Results

As in both previous Experiments, participants were excluded if their median accuracy or RT for either the training or test blocks were beyond 1.5 times the IQR from the group median (n = 14) or if they gave any indication of noticing the manipulation (n = 11). On average, participants responded in the absence of the target on 2.8% on the training block trials, and 4.6% in the test blocks. Figure 3.4 presents the average median correct RT data from the test phase of this experiment. Probabilistic training did have an effect on RT in the test phase of the experiment. There was a three way interaction
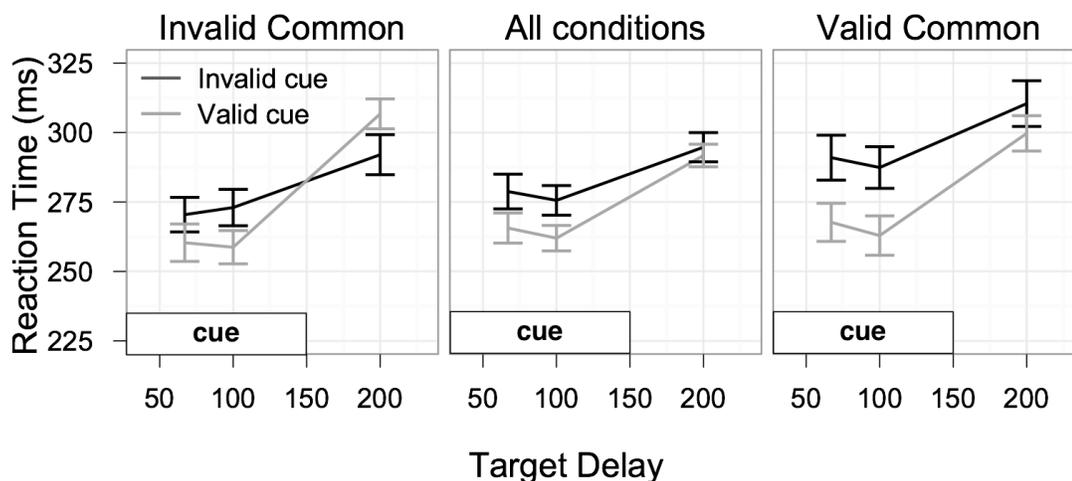
Figure 3.4. Average median correct RT for Experiment 3.3. These data are from the testing phase of the experiment, which was identical for each group. The box labeled 'cue' in the lower left corner of each plot indicates the target delays for which the cue remained on screen during the presentation of the target. Error bars are +/- 1 standard error of the mean of participant medians.

between the factors of cue validity, training group, and the delay duration ($F(4,154) =$ 2.47, $p = .047$). This interaction manifested as a stronger delay duration and cue validity interaction for participants trained in the invalid common condition ($F(2,46) = 26.72$, $p <$ .001) than the valid common ($F(2,48) = 3.99$, $p = .025$) compared to those trained in the all cueing condition ($F(2,60) = 4.28$, $p = .018$). Of greater theoretical relevance, there was a large interaction between cue validity and training condition ($F(2,77) = 8.36$, $p < .001$). This was due to a significant effect of valid cueing for those in the valid common condition ($t(74) = 7.74$, $p < .001$), a smaller effect for training in all cueing conditions ($t(92) = 4.98$, $p < .001$), and no effect for invalid common condition ($t(71) = 1.21$, $p =$ .231).

Discussion

Experiment 3.3 demonstrated that probabilistic training using only short target durations produces training effects that generalize to subsequent testing during which all

conditions appear equally often. Exposure to diagnostic contingencies with primarily valid cues resulted in larger positive cueing effects than training with equal proportions of valid and invalid cues. Exposure to primarily invalid cues also affected subsequent testing, however this difference was smaller and seemed to be driven primarily by the longest target delay.

General Discussion

We report evidence that spatial attention implicitly adapts to contingencies that entail combinations of time and location, and, furthermore, that this training continues to affect subsequent test performance that removes these spatiotemporal contingencies. This finding supports the rational forager hypothesis, which assumes that positive and negative cueing effects as commonly observed reflect adaptation to the statistical regularity of the everyday visual environment. These adaptations to spatiotemporal contingencies occurred implicitly over the course of a single experimental session and generalized to blocks of non-diagnostic trials and untrained target delays. At long target delays, where negative cueing is typically observed, the trained spatial contingencies readily affected RT behavior in the expected manner. For instance, in Experiment 3.1, training that targets were likely to appear at a cued location after a long target delay produced positive cueing effects at the trained long target delay during subsequent testing rather than the more typical negative cueing effects (i.e., IOR) that are seen at long target delays. For short target delays, positive cueing was found in Experiment 3.1, but was not influenced by training. In Experiment 3.1, short target delays differed from long target delays in that only for short target delays did the cue visually overlap with the target. Experiment 3.2

used cues that never overlapped with targets and yet the same pattern of results was found (i.e., there were training effects only for long target delays), This ruled out differential cue overlap as an explanation of the differential training effects. Experiment 3.3 tested whether the training difference between short versus long target delays was due to the use of both delays in an intermixed fashion. Intermixed delay contingencies necessitates two rapid shifts of attention to make full use of the spatiotemporal contingencies. In Experiment 3.3, only short target delays were trained and tested, making it possible to take advantage of the trained contingency with a single shift of attention. With this change, training effects were found for short target delays.

The finding of training effects for both short and long target delays that generalize to untrained delays during subsequent testing without contingencies supports the rational forager hypothesis. These results demonstrate that spatial attention can readily adapt to spatiotemporal contingencies even as adults, and even with just a short exposure period. Thus, both positive and negative cueing effects as commonly observed may reflect adaptation to the statistical regularities of the everyday visual environment.

Whether the everyday visual environment actually has the structure that corresponds to typical spatial cueing behavior awaits further study. Nevertheless, these experiments demonstrate that if the everyday environment has these properties, spatial attention should naturally adapt to these regularities and this adaptation should affect performance as typically tested in the lab. Study of the everyday visual environment is difficult because it requires computational definitions of visual saliency and task-relevancy. However, initial attempts provide some evidence in support of the statistical regularities hypothesized by the rational forager account. This was achieved by first

analyzing the spatial content of natural image patches using component analyses, which yields Gabor-like receptive fields (Bell & Sejnowski, 1997; Olshausen & Field, 1996). By analyzing patches in both space and time, components with motion are found. Notably, a large portion of these dynamic components flip polarity in time (Hyvärinen, Hurri, & Hoyer, 2009). That is, many components are initially primarily light intensity, but switch to dark intensity after a period (or vice versa). This indicates that the visual environment can be decomposed into pieces that tend to remain similar for a relatively short time period, but change after a delay. However, these findings are at the wrong scale to explain behavioral spatial cueing effects (e.g., too fast and too small), which is to be expected considering that these analyses were based on image pixels rather than visual objects. Nevertheless, this is suggestive that something similar may arise if an algorithm is developed for defining saliency and relevancy at the level of objects.

Beyond spatial cueing, a number of non-spatial paradigms find that increasing exposure time/delay between an initial presentation and a subsequent target produces a behavioral pattern that similarly changes from facilitation for compatible trials to a deficit for compatible trials. Furthermore, the time course of these transitions is strikingly similar to that found with spatial cueing. These paradigms include attentional blink (Raymond, Shapiro, & Arnell, 1992), immediate repetition priming (Huber, 2008; Huber, Shiffrin, Quach, & Lyle, 2002), repetition blindness (Kanwisher, 1987), and negative priming (Neill, 1977), which all show variants of this pattern. For the case of repetition priming, Huber and O'Reilly (2003) proposed that this dynamic time course is due to synaptic depression, which exists as a mechanism for parsing events in time to avoid source confusion between previously viewed objects versus the currently viewed object.

If this same dynamic exists within a spatial attention map, it might explain the time course of spatial cueing. If this account is correct, the current results imply that synaptic depression readily adapts to capture different profiles for the rise and fall of spatial attention as dictated by environmental regularities. Further investigations may allow a unification of these paradigms as learned responses to diagnostic information, and thus provide greater insight into domain general perceptual processing. For instance, an investigation of negative cueing and negative priming concluded that these two paradigms reflect the same underlying proceses (Milliken, Tipper, Houghton, & Lupiáñez, 2000).

One challenge for the rational forager hypothesis is to show that these training effects are generated by the same mechanism as spatial cueing effects in general as opposed to different mechanisms in each case that happen to produce similar median RT results. One way to get at this issue is to examine other attributes of choice behavior besides median correct RT to determine whether the attributes that underlie typical spatial cueing results are similar to the attributes that adapt to spatiotemporal contingencies. Providing evidence against the claim that the same mechanism underlies both phenomena, a study measuring saccadic RT in the presence of different spatial contingencies examined attributes of RT distributions from both correct and error trials as described by the linear ballistic accumulator RT model (LBA, Brown & Heathcote, 2008). This study found a dissociation in the attributes of perceptual decision making that best explained different cueing conditions versus those that best explained training effects. Negative cueing was best modeled as an increase in the rate of accumulation of evidence for invalid cue trials (Ludwig, Farrell, Ellis, & Gilchrist, 2009), but the effect of

training was related to changes in the evidence threshold for responding (Farrell et al., 2010). In other words, negative cueing was due to slower detection of the target whereas training affected median RT by changing the bias to report targets. However, Farrell et al. used single long target delay and the RT data came from blocks of trials while participants were still experiencing the spatial contingency. Thus, it is possible that the training effects in their study primarily reflected the short-term effect of the last few trials. For instance, it may be that the last trial affects the response threshold while long-term learning, such as examined in the current experiments, affects the mean accumulation rate. To explore this issue in greater detail, the Appendix reports the results of applying the LBA model to the data from Experiment 3.1. In contrast to the results of Farrell et al., and in support of the rational forager hypothesis both training effects and cueing effects were best accounted for by changes in the mean accumulation rate. Thus, based on the long lasting training effects examined here, we find support for the claim that a common mechanism underlies both cueing effects in general as well as adaptation to the visual environment.

The current study provides additional evidence that our perceptual and attentional systems are not fixed, and can change in fundamental ways after just a brief training session. This implies that attempts to exploit the spatiotemporal properties of attention, (e.g., flashing advertisements to get our attention, centralized brake lights, Kahane & Hertz, 1998) may ultimately be futile as attention adapts and updates based on recent experience. Our results may also be relevant to patient populations. For instance, patients with Alzheimer's disease (Faust & Balota, 1997) and patients with schizophrenia require longer delays between cues and targets than typical individuals to find negative cueing

effects (Gouzoulis-Mayfrank et al., 2004; Larrison-Faucher, Briand, & Sereno, 2002).

Rather than interpreting this as slowed processing, our results suggest an intriguing

alternative: it may be that these patients have adapted to a different set of environmental

spatiotemporal regularities that occurs when one's definition of a task relevant target is

changed due to cognitive level deficits (i.e., perhaps these patients are viewing different

aspects of the visual environment, and for this different set of preferred objects there is a

different set of spatiotemporal regularities).

In conclusion, these experiments demonstrated that participants implicitly and

rapidly adapt to novel spatiotemporal contingencies in spatial cueing, and this adaptation

generalizes to untrained delays and generalizes to the typical testing situation that uses all

combinations of delay and cueing equally. These results verify key assumptions of the

rational forager hypothesis, which proposes that typical spatial cueing effects reflect the

statistical regularities of the everyday visual environment. The rational forager hypothesis

assumes that spatial attention can adapt to spatiotemporal regularities for both short and

long target delays, that this adaptation is sufficiently long lasting, and that it can

generalize to the typical testing situation. The current experiments verified these

assumptions.

Appendix

Here we present fits of the results from Experiment 3.1 using the linear ballistic accumulator (LBA) model of RT (Brown & Heathcote, 2008). The LBA model fits RT distributions using 5 parameters: b, the response threshold; v, the mean accumulation rate; s, the standard deviation of the accumulation rate; A, the upper limit of the starting point; and T_{err}, the non-decision time. RT of a single trial is modeled as an accumulation process with a response occurring when the accumulator reaches the threshold, b. The accumulation starts at a random point between 0 and A, and increases linearly at a rate randomly drawn from a Gaussian distribution with mean v and standard deviation s. A fixed time for non-decision processes (T_{err}) is also added. Figure A3.1 provides a graphical representation of this process. The parametric nature of the model allows for analytic solutions, making it easier to apply the model as compared to similar sequential processing models (Ratcliff, 1978).
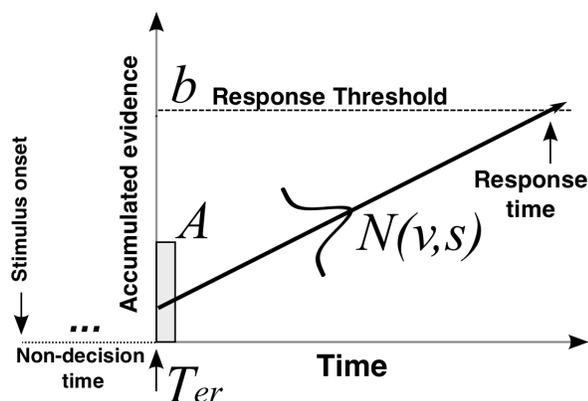

Figure A3.1. Mechanics and parameters of the LBA model. After stimulus onset, there is a fixed duration unrelated to the decision process (non-decision time, T_{err}). After this time, an accumulator starts accruing evidence linearly (the dark line pointing to the upper right), from a starting point randomly chosen from a uniform distribution between 0 and the parameter A. Evidence accumulates linearly with a slope randomly drawn from a normal distribution with mean v and standard deviation s, until reaching a response threshold, b. Integrating over the random values in this process, a complete RT distribution is formed.

The LBA model usually consists of separate accumulators for each possible competing response, and models accuracy as the response crossing the threshold first. However, in the current application we have only one response, and therefore we only used one accumulator to model RT. The model was fit the data as follows. We first computed 5 quantiles of the cumulative RT distribution at 10, 30, 50, 70, and 90%. The cumulative distribution of the LBA accumulator was evaluated at these 5 RTs and one corresponding to the response deadline of 1,500 ms. The LBA cumulative proportions at these RTs were compared to the observed quantile proportions (.1, .3, .5, .7, .9, and 1) by with two measures of fit, G^2 and AIC:

$$G^2 = 2\sum_i \left(N_i \sum_j p_{ij} \log\left(\frac{p_{ij}}{\hat{p}_{ij}}\right) \right)$$

$$AIC = G^2 + 2(n_{params})$$

The index variable i increments over conditions, while j increments over quantiles. N_i is the number of data points for condition i, p_{ij} is the true quantile proportion, $\hat{p}_{ij}$ is the estimated quantile proportion from the model, and n_{params} is the number of parameters. Best fit parameters were selected by minimizing G^2 with numerical optimization.

Given the large space of possible mappings of the five model parameters onto the five levels of target delay, and valid/invalid cueing conditions, we used a nested modeling procedure to determine which factors are best fit by which parameters. Figure

A3.2 presents an illustration of this process with more complex models (i.e., models with more free parameters) appearing at the top and less complex models appearing at the bottom. We started by fitting all conditions with a single set of parameters (i.e., the 'base model' at the bottom of the figure). In this model, the parameters $s, A, v,$ and T_{err} were allowed to take on any value to maximize the likelihood of the observed data across all conditions, while b was fixed at 1 (fitting all parameters results in an under-constrained model, Donkin, Brown, & Heathcote, 2009). In other words, for this base model, all conditions were fit with the same parameters as if there were no differences between conditions. Moving up in the nested model hierarchy on the left-hand side in the figure, we next fit each participant's data allowing one of the parameters to take on different values (i.e., an 'open' parameter) for each of the ten conditions (or nine in the case of b), while the other parameters were forced to take on the same value for all conditions similar to the base model. This results in a fit with thirteen total parameters for each participant. This was repeated for each parameter, resulting in five separate model fits of thirteen parameters, one with each parameter open to vary over all conditions. Moving up the hierarchy from this point, there were two options for more complex models by allowing another parameter to vary over either the five target delay conditions or the valid and invalid cueing conditions, for a total of seventeen or fourteen parameters respectively. Nested above these two points in the hierarchy were models with one parameter open across all conditions, one parameter open for target delay, and a third parameter open for the cueing conditions. These models had eighteen parameters to fit the six quantiles for each of the ten conditions. Finally, we allowed all parameters (with the exception of b which was fixed to 1.0 across all conditions) to vary openly for each

condition, for a total of forty parameters. For additional comparisons, we also fit models with one parameter free over target delay conditions, and another over valid and invalid cueing conditions (nine parameters, as seen in the lower right hand portion of the figure). After fitting each individuals data, the fit statistics were summed over individuals, and evaluated by AIC.
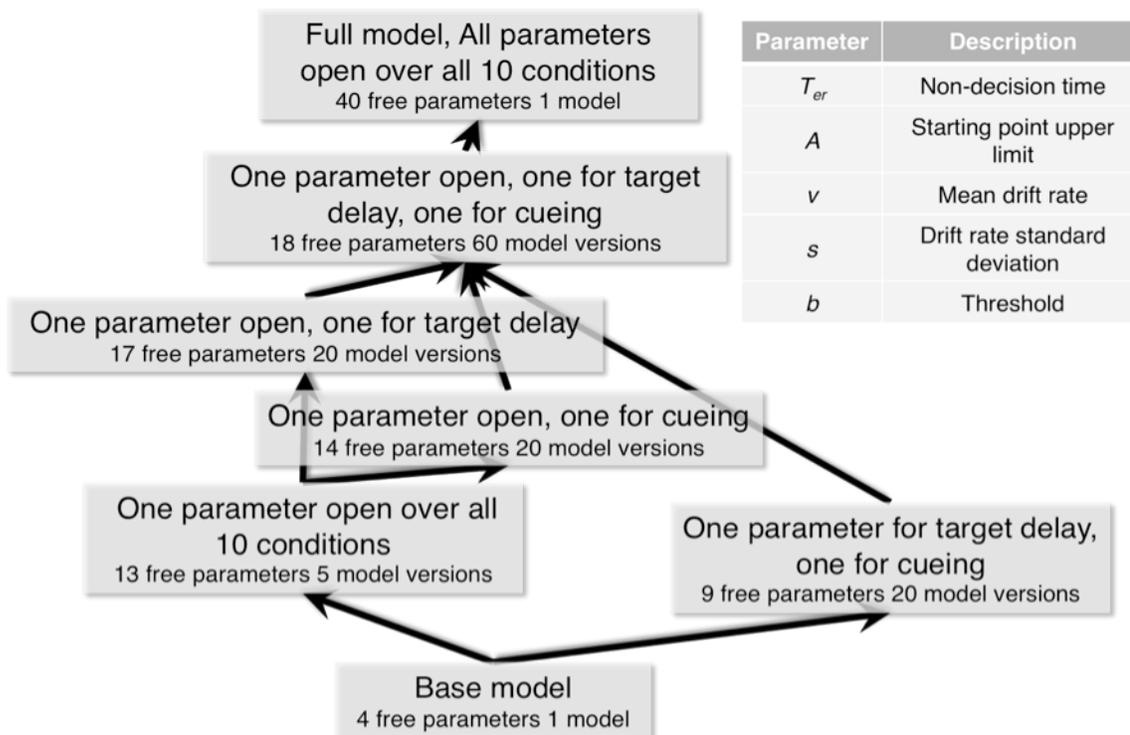

Figure A3.2. The nested modeling procedure. Each model type was fit for each possible combination of different parameters to make a number of different model versions. The base value of the threshold parameter was fixed to 1.

The best fit model as evaluated by AIC allowed v, the mean accumulation rate to vary openly over the ten conditions, and varied s, the accumulation rate standard deviation, between valid and invalid cue conditions. The best fit parameters and average reconstructed data are presented in Figure A3.3. The reconstructed data consists of the median of the fit distributions for each condition, averaged over the fits to each

participant. As can be seen in the figure, changes in the mean accumulation rate captured both the effect of training (three way interaction between training, cueing, and target delay, $F(8,188) = 5.62$, $p = .009$), and the cueing effects (two way interaction between target delay and cueing, $F(4,188) = 37.34$, $p < .001$). Neither the standard deviation of the accumulation rate, nor the other parameters, varied between training or cueing conditions (all $F <= 1.17$). Based on this model fit, which includes data across spatial and temporal manipulations of cueing, both cueing effects and training are the result of changes in mean accumulation rate.
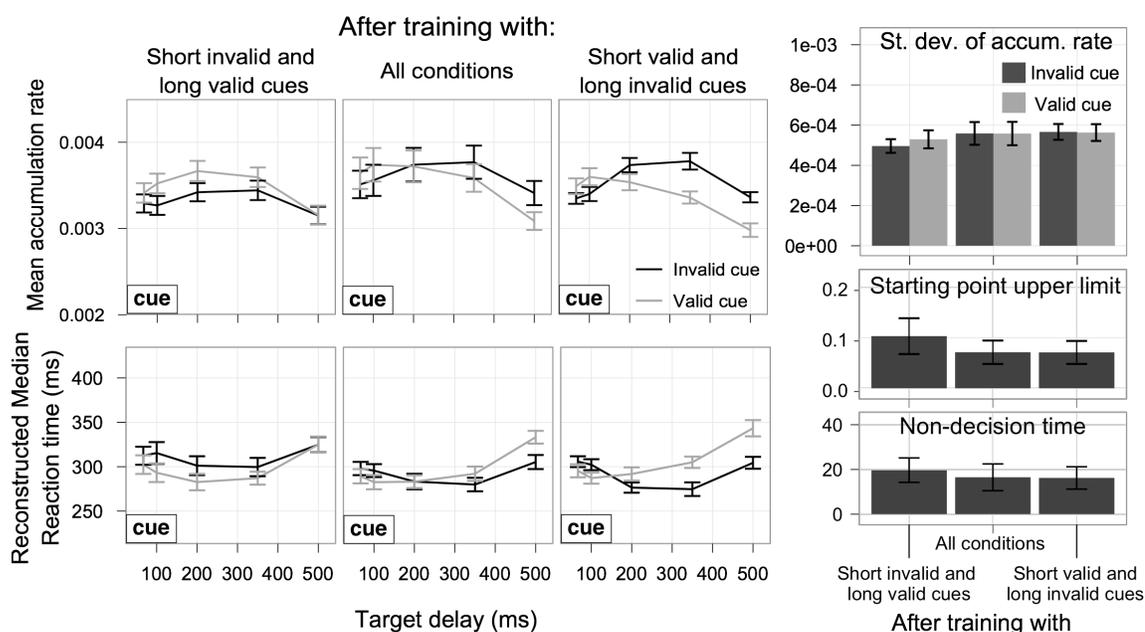

Figure A3.3. LBA fitting results. For the best fitting model, the mean accumulation rate was allowed to vary between all conditions and the standard deviation of the accumulation rate was allowed to vary between cueing conditions. The other parameters took on the same value for the different cueing and target delay conditions and the threshold was fixed at 1 for all conditions. The three graphs in the lower left present the reconstructed data from the best fit model. To reconstruct the average data, the parameters for each participant's data were used to obtain a median RT from the fit model, which was then averaged across participants. Error bars are +/- 1 standard error of the mean.

References

Bell, A. J., & Sejnowski, T. J. (1997). The "independent components" of natural scenes are edge filters. *Vision Research*, *37*(23), 3327–3338.

Berlucchi, G. (2006). Inhibition of return: A phenomenon in search of a mechanism and a better name. *Cognitive Neuropsychology*, *23*(7), 1065–1074.

Brown, S. D., & Heathcote, A. (2008). The simplest complete model of choice response time: linear ballistic accumulation. *Cognitive Psychology*, *57*(3), 153-178. doi:10.1016/j.cogpsych.2007.12.002

Castel, A. D., Pratt, J., & Drummond, E. (2005). The effects of action video game experience on the time course of inhibition of return and the efficiency of visual search. *Acta Psychologica*, *119*(2), 217-230. doi:10.1016/j.actpsy.2005.02.004

Cheal, M., Chastain, G., & Lyon, D. R. (1998). Inhibition of return in visual identification tasks. *Visual Cognition*, *5*(3), 365–388.

Dodd, M. D., & Pratt, J. (2007). The effect of previous trial type on inhibition of return. *Psychological Research*, *71*(4), 411-417. doi:10.1007/s00426-005-0028-0

Donkin, C., Brown, S. D., & Heathcote, A. (2009). The overconstraint of response time models: rethinking the scaling problem. *Psychonomic Bulletin & Review*, *16*(6), 1129-1135. doi:10.3758/PBR.16.6.1129

Farrell, S., Ludwig, C. J. H., Ellis, L. A., & Gilchrist, I. D. (2010). Influence of environmental statistics on inhibition of saccadic return. *Proceedings of the National Academy of Sciences*, *107*(2), 929-934. doi:10.1073/pnas.0906845107

Faust, M. E., & Balota, D. A. (1997). Inhibition of return and visuospatial attention in healthy older adults and individuals with dementia of the Alzheimer type. *Neuropsychology*, *11*(1), 13-29.

Gouzoulis-Mayfrank, E., Heekeren, K., Voss, T., Moerth, D., Thelen, B., & Meincke, U. (2004). Blunted inhibition of return in schizophrenia–evidence from a longitudinal study. *Progress in Neuro-Psychopharmacology and Biological Psychiatry*, *28*(2), 389–396.

Huber, D. E. (2008). Immediate priming and cognitive aftereffects. *Journal of Experimental Psychology: General*, *137*(2), 324-347. doi:10.1037/0096-3445.137.2.324

Huber, D. E., & O'Reilly, R. (2003). Persistence and accommodation in short-term priming and other perceptual paradigms: temporal segregation through synaptic depression. *Cognitive Science*, *27*(3), 403-430. doi:10.1016/S0364-0213(03)00012-0

Huber, D. E., Shiffrin, R. M., Quach, R., & Lyle, K. B. (2002). Mechanisms of source confusion and discounting in short-term priming: 1. Effects of prime duration and prime recognition. *Memory & Cognition*, *30*(5), 745-757.

Hyvärinen, A., Hurri, J., & Hoyer, P. O. (2009). *Natural Image Statistics: A probabilistic approach to early computational vision.* London, U.K.: Springer.

Kahane, C. J., & Hertz, E. (1998). *The long-term effectiveness of Center High Mounted Stop Lamps in passenger cars and light trucks*. (NHTSA DOT HS 808 696). Retrieved from http://www.nhtsa.gov/cars/rules/regrev/evaluate/808696.html

Kanwisher, N. (1987). Repetition Blindness: Type recognition without token individuation. *Cognition*, *27*(2), 117–143.

Klein, R. M. (2000). Inhibition of return. *Trends in Cognitive Sciences*, *4*(4), 138-147.

Larrison-Faucher, A., Briand, K. A., & Sereno, A. B. (2002). Delayed onset of inhibition of return in schizophrenia. *Progress in Neuro-Psychopharmacology and Biological Psychiatry*, *26*(3), 505–512.

Ludwig, C. J. H., Farrell, S., Ellis, L. A., & Gilchrist, I. D. (2009). The mechanism underlying inhibition of saccadic return. *Cognitive Psychology*, *59*(2), 180-202. doi:10.1016/j.cogpsych.2009.04.002

Lupiáñez, J., Decaix, C., Sieroff, E., Chokron, S., Milliken, B., Bartolomeo, P., & Siéroff, E. (2004). Independent effects of endogenous and exogenous spatial cueing: inhibition of return at endogenously attended target locations. *Experimental Brain Research*, *159*(4), 447–457. doi:10.1007/s00221-004-1963-5

Lupiáñez, J., Milán, E. G., Tornay, F. J., Madrid, E., & Tudela, P. (1997). Does IOR occur in discrimination tasks? Yes, it does, but later. *Perception & Psychophysics*, *59*(8),, 1241-1254.

Maruff, P., Yucel, M., Danckert, J., Stuart, G., & Currie, J. (1999). Facilitation and inhibition arising from the exogenous orienting of covert attention depends on the temporal properties of spatial cues and targets. *Neuropsychologia*, *37*(6), 731-744.

Maylor, E. A. (1985). Facilitatory and inhibitory components of orienting in visual space. *Attention and performance XI*, 189–204.

McAuliffe, J., & Pratt, J. (2005). The role of temporal and spatial factors in the covert orienting of visual attention tasks. *Psychological Research*, *69*(4), 285-291. doi:10.1007/s00426-004-0179-4

Milliken, B., Tipper, S. P., Houghton, G., & Lupiáñez, J. (2000). Attending, ignoring, and repetition: on the relation between negative priming and inhibition of return. *Perception & Psychophysics*, *62*(6), 1280-1296.

Mondor, T. A., Breau, L. M., & Milliken, B. (1998). Inhibitory processes in auditory selective attention: evidence of location-based and frequency-based inhibition of return. *Attention, Perception, & Psychophysics*, *60*(2), 296–302.

Neill, W. T. (1977). Inhibitory and facilitatory processes in selective attention. *Journal of Experimental Psychology: Human Perception and Performance*, *3*(3), 444-450. doi:10.1037/0096-1523.3.3.444

Olshausen, B. A., & Field, D. J. (1996). Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, *381*(6583), 607–609.

Posner, M. I., & Cohen, Y. (1984). Components of visual orienting. *Attention and performance X: Control of language processes*, *32*, 531–556.

Posner, M. I., Snyder, C. R., & Davidson, B. J. (1980). Attention and the detection of signals. *Journal of Experimental Psychology: General*, *109*(2), 160-174.

Pratt, J., Hillis, J., & Gold, J. M. (2001). The effect of the physical characteristics of cues and targets on facilitation and inhibition. *Psychonomic Bulletin & Review*, *8*(3), 489-495.

Pratt, J., Kingstone, A., & Khoe, W. (1997). Inhibition of return in location- and identity-based choice decision tasks. *Perception & Psychophysics*, *59*(6), 964-971.

Rafal, R. D., Calabresi, P. a, Brennan, C. W., & Sciolto, T. K. (1989). Saccade preparation inhibits reorienting to recently attended locations. *Journal of Experimental Psychology. Human Perception and Performance*, *15*(4), 673-685.

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

Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: an attentional blink?. *Journal of Experimental Psychology. Human Perception and Performance*, *18*(3), 849-860. doi:10.1037/0096-1523.18.3.849

Spalek, T. M. (2004). Supporting the attentional momentum view of IOR: Is attention biased to go right? *Perception & Psychophysics*, *66*(2), 219-233.

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. doi:10.1111/j.0956-7976.2005.00774.x

Spence, C., & Driver, J. (1998). inhibition of return following an auditory cue. *Experimental Brain Research*, *118*(3), 352–360.

Taylor, T. L., & Klein, R. M. (1998). On the causes and effects of inhibition of return. *Psychonomic Bulletin & Review*, *5*(4), 625–643.

Wang, Z., Zhang, K., & Klein, R. M. (2010). Inhibition of return in static but not necessarily in dynamic search. *Attention, Perception, & Psychophysics*, *72*(1), 76-85. doi:10.3758/APP.72.1.76

Warner, C. B., Juola, J. F., & Koshino, H. (1990). Voluntary allocation versus automatic capture of visual attention. *Perception & Psychophysics*, *48*(3), 243-251.

Weaver, B., Lupiáñez, J., & Watson, F. L. (1998). The effects of practice on object-based, location-based, and static-display inhibition of return. *Attention, Perception, & Psychophysics*, *60*(6), 993–1003.

Weger, U., Abrams, R. A., Law, M., & Pratt, J. (2008). Attending to objects: Endogenous cues can produce inhibition of return. *Visual Cognition*, *16*(5), 659-674. doi:10.1080/13506280701229247

Wolfe, J. M., & Horowitz, T. S. (2004). What attributes guide the deployment of visual attention and how do they do it? *Nature Reviews Neuroscience*, *5*(6), 495–501.

DISCUSSION

The preceding work characterized adaptations of the segregation of information and allocation of attention to the temporal properties of the world. Chapter 1 demonstrated that similar duration-dependent priming effects occur with face images as do with words. These results are consistent with the habituation account of word priming data, which suggests that activation of a representation is initially integrated with previous stimuli, but eventually habituates in response to the strength and speed of activation (Huber, 2008; Huber & O'Reilly, 2003). As a result, short primes produce positive priming effects, while longer primes produce negative priming effects. The similar results with both words and faces demonstrate that these priming dynamics are a more general phenomena, and not governed purely by linguistic attributes of the stimuli. Interestingly, individual differences in the time course of priming were found with faces that were not found with words. Participants who needed longer (67-100 ms) presentations to identify a face, showed larger peak positive priming effects and no negative priming effects. This is in contrast to participants requiring only a short target duration (33-50 ms) for accurate perception. These fast participants showed significant negative priming effect for 2,000 ms primes. For inverted faces, both groups of participants showed priming results similar to the participants needing a longer target threshold. The differences in between upright and inverted faces hint that the same visual stimulus can produce different patterns priming. This may be due to simple familiarity with upright faces, or because of distinct processing for upright faces (Downing, Chan, Peelen, Dodds, & Kanwisher, 2006; Yin, 1969).

Chapter 2 tested the sensitivity of temporal segregation to occurrence and co-occurrence statistics. The first experiment tested the effects of stimulus frequency in short term priming. Priming dynamics were measured for words, pronounceable non-words, non-pronounceable non-words, and inverted words. The prime duration producing the peak priming effect increased as familiarity decreased from words to non-words, and additionally to inverted words. The existence and strength of negative priming increased with greater stimulus frequency. A second experiment investigated the influence of co-occurrence statistics by using unidirectional associative word pairs. Association direction was found to have an effect on the time course of priming, unlike previous studies only examining two prime strengths. Priming was found for both forward and backward direction associations. However, priming through forward associations produced slightly larger priming effects, with the prime effects peaking at later priming durations than backward primes. A final experiment tested the effects of expectation using actual co-occurrence rates. In this study primes sometimes created expectations for either that target or foil. Three types of priming were compared, non-expected repetitions (identical to the previous immediate repetition priming experiments), expected non-repetitions, and expected repetitions. The results for expected non-repetitions were similar to those for associative priming, suggesting that association and expectation may prime in a similar manner, possibly through the same mechanism. Expected repetitions tested the hypothesis that the typical pattern of priming dynamics was due to high level expectations for repetitions after short prime durations, and against repetitions after long prime durations. Expected repetitions produced similar results as non-expected repetitions, with actually slightly larger priming effects. It is possible that this effect is

produced by the combination of expectation priming seen for expected non-repetitions and repetition priming effects.

While only small effects of the order of co-occurrence were found in Chapter 2, Chapter 3 tested whether general spatial-temporal relationships are utilized by perception. To test the idea that environmental statistics influence general spatial-temporal processing in spatial cueing effects, participants were implicitly trained in one of several diagnostic cueing relationships. After a short amount of training with diagnostic trials, participants demonstrated biases congruent with training in a non-diagnostic environment. Spatial cueing effects at long cue target onset asynchrony adapted quickly, even with opposite relationships at short cue target onset asynchrony. Adaptation to diagnostic relationships at short cue-target onset asynchronies occurred in situations where reliable expectations requiring only one shift of attention could develop. Reaction time modeling attributed both the changes after training and the typical cueing effects to the rate of accumulation of target information.

These studies show that the temporal properties of perception are adapted to the structure of the world, and have demonstrated the relative strength of these adaptations. Using multiple types of stimuli revealed that the dynamics of priming are adapted to the occurrence of individual stimuli for both linguistic and non-linguistic stimuli. Adaptations in habituation were also found to occur at multiple levels of processing. Implicit learning of spatial-temporal regularities between the location of a cue and the appearance of a target resulted in generalization of these regularities as biases in non-diagnostic environments. Adaptations to basic properties of stimuli, or general temporal patterns were large in comparison to adaptation to co-occurrence relationships between

specific stimuli. While large priming effects of association were found, there were only small effects of associative direction. This indicates minimal influence of directional relationships between specific items in perceptual segregation. Additionally, contrary to a pure high-level expectation account of habituation effects, words which are likely repeat show largely similar priming effects to words which are not likely to repeat. There was only a small effect of expected repetition. The obtained results can still be explained though expectation at multiple levels, i.e. even if the full word is likely to repeat, letters in general do not. However, in this case the relative contributions of low-level expectation are much greater. In Chapter 3, adaptations to short cue target relationships were only found for situations where the diagnostic relationships did not require shifts of attention during the trial. This indicates that participants were unable or unwilling (at least with a small amount of uninformed training) to learn contingencies requiring different predictions at different points in time. Simplifying training by removing this requirement resulted in generalization to non-structured tasks.

The dynamics of mental processes are complicated, and these investigations have only scratched the surface. One difficulty is the consideration of interactions between multiple levels of processing. The same visual stimuli can activate different semantic content depending on what preceded it, what orientation it was in, and in what context it is perceived in. The present results suggest that responses to repeated stimuli are mostly, but not entirely, determined by habituation to the visual content of the stimuli itself. Habituation to processed stimuli is seemingly a beneficial response. Once something is sufficiently perceived, it is rational to seek novel, useful information. Similarly, if something has lasted longer than a typical duration, subsequent feature activations are

likely to come from something new. As shown in these experiments, the dynamics of habituation are adjusted for items occurring with different frequency. Attention can also learn to adapt to general relationships between salient events. However, perceptual dynamics are not strongly tailored to relationships between individual items. These results supports the idea that habituation is a general mechanism, and that tuning of habituation occurs mostly in response to the activation of individual items. While general habituation dynamics largely guide responses in these paradigms, it still may be important to understand the smaller effects, as these small differences might reflect large distinctions in subjective experience. Unfortunately, these small differences are difficult to study behaviorally. Instead of trying to untangle the many confounding factors in the stimuli available to test small differences, another approach is to study how the dynamic representations of stimuli are learned, and apply these models to information from the world.

The preceding chapters describe a number of investigations into the temporal activation of perceptions in the mind. Segregation of items in time through habituation of activated representations was measured by repetition priming. Priming dynamics were shown to be sensitive to stimulus occurrence and co-occurrence, but only weakly influenced by ordered relationships between individual items. The learning of general spatial-temporal relationships was investigated in a spatial cueing paradigm. Participants implicitly learned diagnostic regularities in training and generalized them to non-diagnostic blocks. In total, these experiments demonstrate ways that perception and attention is adapted to temporal statistics.

References

Downing, P. E., Chan, A. W.-Y., Peelen, M. V., Dodds, C. M., & Kanwisher, N. (2006). Domain specificity in visual cortex. *Cerebral Cortex, 16*(10), 1453-1461. doi:10.1093/cercor/bhj086

Huber, D. E. (2008). Immediate priming and cognitive aftereffects. *Journal of Experimental Psychology: General*, *137*(2), 324-347. doi:10.1037/0096-3445.137.2.324

Huber, D. E., & O'Reilly, R. C. (2003). Persistence and accommodation in short-term priming and other perceptual paradigms: temporal segregation through synaptic depression. *Cognitive Science, 27*(3), 403-430. doi:10.1016/S0364-0213(03)00012-0

Yin, R. K. (1969). Looking at upside-down faces. *Journal of Experimental Psychology*, *81*(1), 141-145.