Title
Beyond the Here and Now: Experimental Studies in How Action Dynamics Reflect Complex Cognitive Processes

Author
Tabatabaeian, Maryam

Publication Date
2017

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA, MERCED

Beyond the Here and Now: Experimental Studies in How Action Dynamics Reflect Complex Cognitive Processes

A dissertation proposal in Cognitive and Information Sciences

by

Maryam Tabatabaeian

Committee in charge:
Professor Rick Dale, Chair
Professor Ramesh Balasubramaniam
Professor Evan Heit

2017

The dissertation of Maryam Tabatabaeian is approved,
and it is acceptable in quality and form for publication on microfilm and electronically:

Professor Rick Dale, Chair

Professor Ramesh Balasubramaniam

Professor Evan Heit

University of California, Merced

2017

I dedicate this dissertation to:

my father who taught me how to dream high,

&

my mother who sacrificed a lot to see my dreams come true.

Contents

List of Figures

List of Tables

Acknowledgements

First and foremost, I would like to thank my primary advisor, Rick Dale, who was the best academic mentor I could have ever asked for. He was there for me in all the steps I learned to take in this road. He was always a spark of inspiration even in our shortest conversations. There was always something new and exciting coming out of our chats and that kept me going throughout these years. We had our own unique advisor-advisee relationship and now that the official part of it is coming to an end I know I have a mentor for life. I depend on having you in my life and for the rest of my career, Rick!

I would like to thank Ramesh Balasubramanium, as he brought neuroscience to the picture and nothing could have satisfied my enthusiasm and excitement more than knowing there is an expert with cool toys in the program. He accepted to be on my committee and never stopped inspiring me and opening new and exciting doors in front of me. I very much appreciated all the incredible opportunities and collaborations that Ramesh made possible as well as every single chat we've had about research, neuroscience, music, politics, and the whole world.

I would like to thank Evan Heit, for being the sound of reasoning and decision making in this research. From the days that I chose my first year project to the days that it blossomed to be the base for what I wanted to do in graduate school, I owe Evan for pointing me to the right resources and always asking the right big-picture questions that guided me through the ambiguous times.

I would like to thank every one of the faculty members in the Cognitive and Information Sciences program as they taught me how to believe in myself and my abilities while providing top-notch education and research culture. I am proud to be known as a Cogsci graduate from this program as I know what a respectable brand you have created. Special thanks to Professors Chris Kello, Michael Spivey, Dave Noelle, and David and Carolyn Jennings for they all offered me valuable insights one way or another throughout this journey. I would also like to thank my professors and mentors at the University of Louisiana, Lafayette who supported me, believed in me and let me start the journey in Cognitive Science under their supervision; special thanks to my mentors Mike Kalish, Subrata Dasgupta, and Ali Parzivand.

Finally, I would like to thank my collaborators, the graduate students and post docs who have overlapped with my tenure at UC Merced. Special thanks to Nicholas Duran, who generously helped me with my onboarding process to academic research in my first year at UC Merced. I also would like to thank National Science Foundation for partially funding the deception sections of this research by Grant BCS-0720322.

Maryam Tabatabaeian
Curriculum Vitae

Education

Ph.D. in **Cognitive Science** **2012-2017**
University of California Merced

M.Sc. in **Computer Science** **2011- 2012**
University of Louisiana at Lafayette

B.Sc. in **Computer Science** **2005- 2009**
Amirkabir University of Technology, Iran

Research Experience

2012-2017 **Graduate Researcher** at Cognitive and Information Sciences, UC Merced

2016 **Visiting Researcher** at Swartz Center for Computational Neuroscience, UC San Diego

2011-2012 **Research Assistant** at Center for Advanced Computer Studies, University of Louisiana, Lafayette

2011-2012 **Research Assistant** at Cognitive Science Program, University of Louisiana, Lafayette

Teaching Experience

2016-2017 **Graduate Stats Tutor,** University of California, Merced

2015-2016 **Modern Cognition,** at University of California, Merced

2014-2015 **Applied Ethics,** University of California, Merced

2013-2014 **Mind, Brain, Cognition,** University of California, Merced

2013-2014 **Cognitive Neuroscience,** University of California, Merced

2012-2013 **Introduction to Cognitive Science,** University of California, Merced

2011-2012 **Data Structures in Computer Science,** at University of Louisiana

Publications

Tabatabaeian, M., Dale, R., & Duran, N. D. (2015). Self-serving dishonest decisions can show facilitated cognitive dynamics. *Cognitive processing*, *16*(3), 291-300.

Tabatabaeian, M., Dale, R., & Duran, N. D. (2014). Facilitation in Dishonesty is Subject to Task Constraints. *Proceedings of the 36th Annual Conference of the Cognitive Science Society*, 1562–1567.

Vinson, D. W., Dale, R., **Tabatabaeian**, M., & Duran, N. D. (2015). Seeing and believing: Social influences on language processing. In *Attention and Vision in Language Processing* (pp. 197-213). Springer India.

Conference Posters & Presentations

Tabatabaeian, M., Dale, R., (2016). An action dynamics study of the onset of prediction. *Poster at 37th Annual Conference of the Cognitive Science Society*.

Tabatabaeian, M., Pilkington, S., Dale, R. (2015) Real-world implementation of Newcomb's thought experiment, using mouse-tracking techniques. *Poster at 37th Annual Conference of the Cognitive Science Society*.

Chernykhovskaya, Y., Jennings, C., **Tabatabaeian**, M. (2015), Implicit Association in Mathematics and Science. *Poster at 37th Annual Conference of the Cognitive Science Society*.

Tabatabaeian, M., Dale, R., & Duran, N. D. (2014). Facilitation in Dishonesty is Subject to Task Constraints. *Proceedings of the 36th Annual Conference of the Cognitive Science Society*, 1562–1567.

Tabatabaeian, M., Pilkington, S., Dale, R. (2014) Real-world implementation of Newcomb's thought experiment, using mouse-tracking techniques. *Poster at Society for Judgment and Decision Making. Long Beach, California.*

Tabatabaeian, M., Dale, R., & Duran, N. D. (2013). Is dishonesty an automatic tendency? *Poster at the 35th Annual Conference of the Cognitive Science Society*.

Awards & Memberships

Summer 2017 **Data Science Fellowship** at Insight Data Science

Spring 2017 **Dean's Career-Building Award** at UC Merced

Summer 2016 **Data Science Fellowship** at KOOSTAMA Inc.

2012-2016 **Summer Research Fellowship**, at UC Merced

Spring 2015 **Finalist Award-winner** in Graduate Speaker Competition, UC Merced

2015-2016 **Member of Dean's Diversity Advisory Council** at UC Merced

2015 – 2016 **International Liaison** at Graduate Student Association, UC Merced

Abstract

The aim of this dissertation is to take a journey into exploring more informative and rather continuous measures for assessing the cognitive processes experienced by humans. The traditional measures, such as Reaction Time, although incredibly helpful fail to provide any particular insight into what mental processes the participants undergo to produce a final result. In the recent years, it is becoming clear that a rich and semi-continuous set of measures is extractable from relatively implicit behaviors such as hand or eye movements. It has been shown that spatial and temporal dynamics of motor movements (i.e. *action dynamics*) can shed light on the progression of high-level cognitive tasks. This rich body of data could provide a nearly real-time translation of unfolding cognitive processes.

Throughout this document, I present a series of studies on how action dynamics could provide a window to observe the emergence of cognitive processes involved in deception and predictive learning. These are two areas in which action dynamics have not been explored nearly enough. I utilized a novel task to test two competing hypotheses concerning the cognitive processes involved in dishonesty. Moreover, a novel paradigm was provided for testing hypothesis regarding sequential learning and the role of prediction in implicit learning.

I close with a report of findings about the processes underlying deceptive behavior and predictive learning. The findings are followed by a discussion about implications of this work for the field of cognitive science and the limitations in action dynamics approach. On the one hand this work takes advantage of action dynamics to get under the hood as these complex cognitive processes unfold in the brain. On the other hand, it utilizes these cognitive phenomena to demonstrate the power of action dynamics in studying higher level complex cognitive processes.

Chapter 1

1. Introduction

1.1 Motivation and Background

For a very long time in the history of behavioral psychology and later on in cognitive psychology, discrete measures such as Reaction Time (RT) and Accuracy were dominating the field; useful dependent variables showing the final result of a process. Comparing reaction times and other similar measures between subjects was the basis for evaluating how different subjects perform in various tasks. These measures, although incredibly helpful, fail to provide any particular insight into what mental processes the participants undergo to produce different reaction times as a final result. Are there behavioral measures that could shed light on the process rather than emphasizing on the final result of that process? If so, what do these measures provide? How reliable are they and do they provide the same final results as suggested by traditional dependent variables such as Reaction Time?

The aim of this dissertation is to take a journey into exploring more informative and rather continuous measures for assessing the cognitive processes experienced by humans. Introducing concepts and measures to address the questions mentioned above, shape the core theme and motivation of this dissertation. Throughout this document, I will be exploring the *dynamics of human actions* to introduce continuous measures for studying cognitive processes. I will be presenting a series of studies of how action dynamics could provide a window to observe the emergence of cognitive processes involved in deception and prediction. The focus is on tracking hand-movements in two main studies: making a decision to lie, and making a prediction in sequence learning.

In the following sections of this Intro, I will provide a brief background on action dynamics and how they have come to be a convenient yet reliable set of tools for tracking the mind in real time.

1.2 Action Dynamics

In the recent years, it is becoming clear that a rich and nearly continuous set of measures is extractable from relatively implicit behaviors such as hand or eye movements. This rich body of data could provide a much better picture of different features in cognitive processing.

One way to observe gradual changes in a cognitive process is to explore action dynamics of the process. It has been shown that spatial and temporal dynamics of motor movements can shed light on the progression of high-level cognitive processes such as decision making and learning (for reviews see Spivey and Dale 2006; Song and Nakayama 2008; Freeman et al. 2011; Freeman and Ambady 2009). This growing line of research on action dynamics suggests that decision-making outcomes are not fully processed in the brain before being sent downstream to be expressed in motor subsystems. Instead, the ongoing competition between alternative options is captured concurrently in a person's overt movement dynamics. Thus, by studying the micro-behavioral properties of an unfolding decision, such as eye movements or reach movements, we might be able to get a sense of the underlying cognitive processes.

A pioneering work using continuous measures of movement trajectories for studying mental processes was offered by Ghez at al (1990). They analyzed "elbow force" trajectories to study the effect of target uncertainty on reach movements. The authors show how subjects set initial values for the amplitude and direction of their arm movement while they are anticipating the target to show up at a variant location. Subjects tended to initially set their default in the middle of the target range and after the target was presented they gradually tailored their arm movement more specifically towards that target in roughly 200ms. These elbow force trajectories provided a valuable insight into the flow of mental processes. Nonetheless, they could not support more spatially complex responses involving multiple joint motions. In a later study Ghez et al (1997), extended the measures by recording direction and amplitude of wrist movements and wrist forces. The results from these experiments also showed that most participants when presented with two competing targets to reach for, tend to move their hands to a middle ground around a spatially average point between the two targets.

Some studies go as far as claiming that the relation is bidirectional. While the internal process of decision-making is reflected in action dynamics, the actions and their restrictions indeed influence the decision making. Studies on "embodied choice" show this bidirectional influence between actions and mental processes. Lepora & Pezzulo (2015) designed a decision making task to compare three models for how actions and decisions interact: a serial decision-then-action model, a parallel action-and-decision model, and an embodied choice model where the action feeds back into the decision making process. The results confirmed that in an ecologically valid setting, the most proper theory of decision and action is offered by "embodied choice" suggesting that dynamics of actions can have a causal influence on cognition.

In more recent studies Gallivan and Chapman (2014) among many others show that when participants are presented with more than one option to choose, their initial hand movement trajectories are "spatially biased" corresponding to how probable each option is to be the target answer. These trajectories become more to the point when a cue is presented in favor of the target option. Action dynamics and more specifically rapid reach movements are simple yet powerful tools to provide a real time observation of the mental processes underlying these motor-response decisions. The results are in consistency with neurophysiological and computational research showing that there are neural populations in the brain offering parallel representations that compete to be chosen (See Gold & Shadlen, 2007 for review; Roitman & Shadlen, 2002; Murakami& Mainen, 2015; Lagzi & Rotter, 2015).

There are various tools for recording the dynamics of actions while people are involved in a cognitive task. Arm movements, wrist movements, and hand movements are common ways of tracking these dynamics. One of the well-known paradigms for collecting action dynamics data is mouse-tracking where the subjects are asked to complete a cognitive task using their computer mouse while the movements of the mouse is being recorded by the researcher. Next chapter gets into the details of this paradigm, as it is the main theme of the current document.

1.2.1 Mouse-tracking paradigm

The paradigm used in the studies presented here involves tracking the movement trajectories of participants' hand. The hand movements are captured by recording the movements of the computer mouse. The participants will move the mouse cursor to click on a desired location on the screen depending on the task they are participating in. The mouse trajectories will be recorded and used for extracting measures that demonstrate the gradual process (See Freeman, Dale, Farmer (2011) for more detail on mouse-tracking paradigm).

Social psychology, more specifically implicit social biases and stereotypes, have repeatedly been the subject of mouse tracking studies. A reason could be the need for implicit measures to detect these personal biases. As pioneers of developing this paradigm, Freeman and colleagues, have utilized mouse-tracking in a variety of social cognition experiments (Johnson, Freeman, Pauker, 2012; Freeman, Penner, Saperstein, Scheutz, Ambady, 2011; Freeman, Pauker, Apfelbaum, Ambady, 2010; Freeman & Ambady, 2009). For instance, in a recent study the paradigm is applied to examine the process of interpreting ambiguous emotional stimuli consistent with internal biases (Mattek, Whalen, Berkowitz, Freeman, 2016). Participants were to assign facial expressions (happy, sad, and surprised) into positive or negative labels. During this forced choice task they used their computer mouse to choose "negative" or "positive" after seeing an image of a facial expression. The surprised facial expressions were the main targets as they were the ambiguous option. The results showed that people with a certain emotional bias tend to either assign the surprised face to the emotion they are biased towards or if they do the opposite their mouse trajectory shows an apparent attraction to the alternative option. For example if someone has a negative emotional bias they have a more difficult time assigning the ambiguous face to the positive category. This hesitation is captured in the nearly continuous movements of their mouse by exhibiting a drag towards the other category. The authors call this a non-modal vs. modal interpretation and demonstrate that regardless of the type of emotional bias non-modal choices take more effort. Using the mouse-tracking paradigm instead of the traditional reaction time helps them depict the process that leads to this extra effort. The authors also exhibit a critical finding by showing that increasing cognitive load lessens the difference between modal and non-modal response trajectories.

In a similar context, Lazerus et al. used the mouse tracking paradigm to test emotional bias towards in-group and out-group members (Lazerus, Ingbretsen, Stolier, Freeman, Cikara, 2016). Participants were randomly assigned to two competitive groups and were asked to first rate the negativity/positivity of facial expressions of in-group and out-group members. Further they categorized the facial expressions into positive and negative using a computer mouse. The authors found that people judged in-group faces as more positive, regardless of the facial emotion, both in deliberate rating and implicit mouse trajectory judgments. These findings validated the results from the mouse-tracking paradigm as it agreed with the explicit ratings.

In another social cognition setting Freeman, Pauker, and Sanchez (2016) used the mouse-tracking paradigm to study racial biases in people with different levels of interracial exposure. They used ambiguous mixed-race faces and asked people to categorize them to Black or White. The results confirmed that White individuals with limited interracial exposure exhibited a unique pattern of shifting between options for mixed-race faces. It was also shown that lower level interracial exposure could serve as the base for dynamics of racial judgments and social biases in evaluating mixed-race individuals. Moreover, Smeding and colleagues developed a mouse-tracking version of Implicit Association Test (IAT) in order to detect biases towards gender-domain stereotypical (e.g, male engineers) and counter-stereotypical (e.g., female engineers) social groups. By adding the continuous measures from mouse trajectories on top of reaction times used in traditional IAT tasks, they were able to depict the details of the decision competition in making stereotypical judgments (Smeding, Quinton, Lauer, Barca, Pezzulo, 2016). The authors show a competition at the abstract level and the sensorimotor level. They conclude that while self-congruency and previously learnt social biases have a huge impact on these decisions the role of sensorimotor apparatus and features of the experimental design could also play a role in the results we see in an Implicit Association Test.

Mouse tracking has also been applied to visual search scenarios (Quétard, et al., 2016) where the participants had to make a decision about the existence of a target in a real-world scene. The target could be placed in an expected location or an unexpected spot. Participants used their computer mouse to report whether the target is present in the scene or not. The results provide evidence that the detection and verification of an existing object is a process of accumulating information about both the existence and absence of a target as competing options throughout an evolving decision. Scenarios with targets in unusual locations showed slower decisions and made the mouse trajectories deviate towards the 'absence' response. Moreover, adding noise to the scene reduced mouse velocity.

In the field of perception, mouse tracking has revealed the ongoing competition between options before the perception is complete. It was used in a study on size congruity effect, which refers to responses being faster when the numerical size and physical size of the stimuli are consistent. A series of mouse tracking experiments were conducted to make a comparison between two models of size congruity effect: early interaction model which states that the inference happens early on in the representational stage and late interaction model indicating that inference happens dynamically through the process as a competition between responses. The findings provided evidence for the late interaction model by showing that the trajectories are more curved when there are incongruities between physical and numerical size. The competition between parallel and partially active responses throughout the perception process is evident by looking at the continuous measures offered by mouse trajectories.

Mouse tracking has been tested in studying various areas of cognitive science including attention (Xiao, Yamauchi, 2017), memory (Koop & Criss, 2016), dual cognitive processes (Freeman & Dale, 2013), embodied cognition and linguistics (Papesh, 2015), cognitive systems expectancies (Coco & Duran, 2016). It has also found its way into ecommerce and online user behavior research (Hibbeln et al, 2017; Horwitz, Kreuter, Conrad, 2016).

This dissertation is an attempt to shed light on the progress of two high level cognitive processes, namely deception and prediction using mouse tracking paradigm.

As stated by Freeman, Dale, and Farmer (2011) *hand in motion reveals mind in motion* and the current dissertation attempts to take advantage of action dynamics as a simple, elegant, arguably ecologically valid and yet very powerful tool to look into the dynamics of brain processes. Theoretical debates in deceptive behavior and predictive behavior will be assessed through experiments utilizing mouse tracking paradigm. Prior to presenting the experimental data, a comprehensive review of mental processes involved in deception, and predictive learning will be provided.

Chapter 2 presents a series of studies in deception. Major debates on the mental processes involved in deception will be discussed in this chapter. The experimental design and collected data will offer a dynamical systems explanation for the nature of the decision to engage in deceptive behavior. Chapter 3 starts with a review on the role of prediction in the most recent interpretations of how brains work. This role will then be explained in the context of statistical learning and how the structure in the environment could give rise to making predictions and therefore learning patterns. The chapter reports a series of experiments that are designed to study the emergence of predictive mind by tracking the dynamics of hand movements. In these experiments, action dynamics are used to track the mental processes as the participants learn the underlying patterns of sequences with various levels of structure. Each experiment introduces a specific modification to the main framework in order to address a certain aspect of this research.

Chapter 4 is an attempt to model the cognitive behavior observed in the experimental results. This chapter will provide the theory and implementation of two modeling approaches; a Recurrent Neural Network (RNN) and a more transparent model of memory and entropy. A background on Bayesian models and how they can be used in this context will also be discussed in this chapter. The final chapter is a general discussion with the hope to weave all the mentioned threads into a tapestry that will depict a more refined picture of action dynamics in studying cognitive processes and more specifically prediction. Further, this chapter briefly discusses the future direction and potential follow-up studies. Chapter 5, the final chapter, is a discussion on the findings of this dissertation and their implication for the field.

Chapter 2

2. Action Dynamics of Deception

We have all been the victim of deception, and given the pervasiveness of little lies (DePaulo, Kashy, Kirkendol, Wyer, & Epstein, 1996), on any given day most of us are also perpetrators. Deception appears to be a surprisingly common component of everyday social interactions (DePaulo et al. 1996; DePaulo, & Kashy, 1998). Inevitably, some individuals choose dishonesty over the truth because they find an advantage in lying. People are tempted to lie when it serves their self-interest, be it financial, social, or emotional. But even so, the question remains: In a tempting situation where being dishonest serves one's self-interest, is it easier for people to lie than stay honest, or is lying always more cognitively complex?

The answer to the above is not obvious and has been the target of some debate, touching on issues in judgment and decision-making about quick and intuitive processes versus slower and deliberative ones.[1] There are at least two relevant theories that would seem to address the underlying processes of dishonest decisions. One view, stemming from Spinoza's hypothesis about an inevitable truth bias in human belief system (Gilbert, 1991), suggests that honesty is the grounded process and is therefore more accessible and immediate. From this perspective, in order to act dishonestly, one first has to overcome a truth bias, resulting in more time and effort (Duran, Dale, & McNamara, 2010; Duran & Dale, 2012; McKinstry, Dale, & Spivey, 2008; Spence, et al., 2001). Indeed, most theories of the mechanisms underlying deception imply that lying is more cognitively costly (Vrij, Mann, Fisher, Milne, & Bull, 2008; Verschuere, & De Houwer, 2011). Those who espouse this view do not often make explicit claims regarding precise underlying mechanisms of deception (e.g., whether more implicit or deliberative). Yet many accounts referenced above, involve a more or less subtle implication that cognitive processes underlying deception are strategically (and thus more slowly) deployed in order to serve self-interest.[2]

There is, however, another theoretical possibility. A more recent line of research argues that dishonesty can be greatly facilitated, perhaps even be "automatic," and therefore will be the first and easier choice in any tempting situation where lying pays. This hypothesis predicts that people will need more time and self-control while being honest and refraining from cheating (Shalvi, Eldar, & Bereby-Meyer, 2012). Shalvi and colleagues (2012) have presented a study showing that dishonesty can in fact be the facilitated, rather than the more cognitively costly, response. Given that under time pressure people are forced to act as dictated by their more accessible tendencies, Shalvi

---

[1] We do not have space to discuss this extensively, but these two notions of fast vs. slow processes have had broad influence on social cognitive for decades (e.g., recently, see Kahneman, 2011), and have also been discussed in the domain of deception by Seymour and colleagues (Seymour, 2001; Seymour & Schumacher, 2009).

[2] Space limitations restrict our discussion of this tendency here, but many theories imply the involvement of cognitive control and related resources in order to deceive, which implies more deliberative processes (see: Walczyk, Roper, Seemann, & Humphrey, 2003; Seymour, & Schumacher, 2009).

et al. suggest that, in a situation where one is tempted to lie and is also under high time-pressure, one will more likely choose to lie. The authors conclude that when lying pays, people will automatically choose dishonesty over truth unless they have enough time to deliberately refrain from lying. It is also consistent with the literature concerning how depletion of self-control can increase the chance of performing dishonestly (Gino, Schweitzer, Mead, & Ariely, 2011).

2.1 Experiment 1

Variables often measured in the behavioral studies mentioned above (e.g. reaction time or the final answer provided by participants) reflect only the end point of a long cognitive process. In order to achieve more information about this longer process, action dynamics could provide measures to “peel back” the cognitive processing that gives way either to an honest or dishonest response.

In the study presented here (Tabatabaeian, Dale, & Duran 2015), I used action dynamics to investigate people’s behavior when they are naturalistically tempted to act dishonestly (as opposed to being explicitly instructed to do so). This allows for testing the two competing hypotheses concerning the cognitive processes involved in dishonesty. The novel task utilized in the current experiment provides a unique opportunity to apply action dynamics to natural act of dishonesty. Inspired by Greene et al. (2009), in this study participants were rewarded for their self-reported accuracy in predicting a virtual coin flip. The movements of their mouse cursor while reporting the accuracy by clicking on Correct or Wrong were recorded and used to illustrate the differences between honest and dishonest decisions. If honesty is the default process, people who behave dishonestly are expected to have mouse trajectories that, although end at the deceptive answer, are curved towards the competing truth response--demonstrating the hesitation and extra effort needed for overcoming the truth bias. On the other hand, if dishonesty is specially facilitated in a self-serving situation, dishonest people will directly choose the rewarding option even though their prediction was not accurate.

2.1.1 Methods

The experiment was programmed as an online Adobe Flash-based game posted on Amazon Mechanical Turk (AMT), a web-based crowdsourcing platform. Experimenters (i.e., "Requesters") can post their designed task on AMT where participants (i.e., "Workers") can sign up and take part in the uploaded task. Workers are paid for completing the task. We used a pilot version of the experiment to estimate the effective sample size for the study. Moreover, a rough power analysis was run to find the required number of observations. However, the work reported here is novel enough that it was difficult to run a standard power analysis from past work.

2.1.2 Participants

Ninety-seven participants were recruited online through Amazon's Mechanical Turk. They were paid $0.40 for their time. Researchers used a numeric code on the server to ensure that participants had actually completed the task, and approved their payment on AMT.

2.1.3 Procedure

Participants were instructed to predict the outcome of 20 consecutive coin flips with a certain pattern to the sequence of heads and tails, which they may or may not notice. They were asked to report their accuracy after each coin flip and were informed that they would win a bonus for each correct prediction. Participants were led to believe that we were interested in how receiving a reward while making private guesses can influence implicit learning of the underlying patterns. However, after their participation, they were debriefed that the study was about response movements of people who tend to cheat when lying serves self-interest.

Following the instructions, participants saw a page that asked them to make their prediction and preferably write it down on a piece of paper so that they will not forget. This page was repeated before each coin-flip. After making the prediction, they were directed to a page where they clicked on a "Flip" button to see an animated coin-flip. Once the coin landed, they could go to the next page and report the accuracy of their prediction. On this page, they saw two boxes on the top left and top right of the screen labeled as "Correct" and "Wrong." The assignment of labels to left or right side of the screen was counterbalanced between subjects. Participants were told to click on one of the boxes based on the accuracy of their prediction. If they chose "Correct" they would see a message indicating that they had received a bonus. On the other hand, if they chose "Wrong" they would see a message indicating that they had received no extra bonus. It is worth noting that this procedure assured subjects that they would not be caught cheating since their predictions were private and experimenters would never know their actual predictions. Figure 1 illustrates the sequence of events in the task.

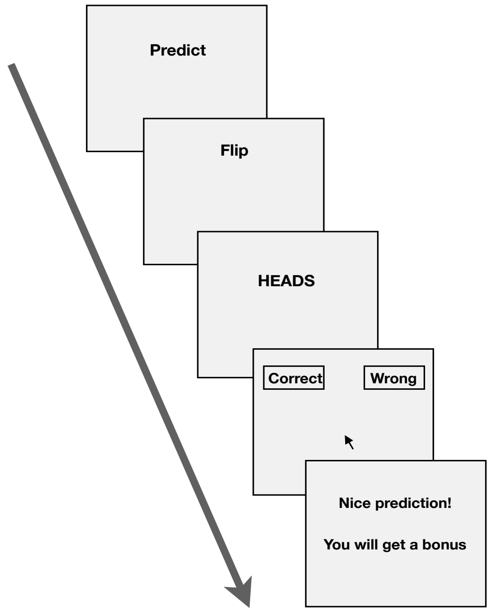

Figure 1. Task sequence: subjects 1) make a prediction, 2) flip the virtual coin, 3) see the outcome and 4) evaluate their prediction by clicking on one of the two boxes on top of the screen (i.e. Correct and Wrong) which were assigned to left or right on a counterbalanced order. 5) They will be informed that they got a bonus (or not) if they report Correct (or Wrong).

Each participant got 20 trials through which the coin outcome was determined using a randomized list. This list was made by a random number generator to guarantee randomness, with the stipulation that heads and tails appeared an equal number of times (50% probability of heads / tails throughout), for each subject. Following the last trial participants were prompted to describe any patterns they might have noticed in the sequence of flips. At the end, every participant received the same bonus payment ($0.25 total). All the mouse movements, where participants clicked on "Correct" or "Wrong" to report their accuracy, were recorded for further analysis.

2.1.4 Results

Since participants' predictions were private, we detected lying by comparing the distribution of self-reported accuracy with the expected distribution of fair coin-flips. Consistent with previous studies in the field (Greene & Paxton, 2009; Shalvi et al., 2012), the distribution of reported correct predictions in the current study (M = 11.78, SD = 2.7) significantly differs from a fair distribution of random coin-flips (M = 10), ($t(96)$ = 6.42, p < .001). This suggests that people have actually exhibited dishonest behavior. Figure 2 shows the distribution of self-reported accuracy for Experiment 1. Obviously, researchers cannot tell whether individuals did or did not lie. However, by the logic of this self-serving task, dishonest responses are more likely among participants in the rightmost portion of the histogram than on the left.

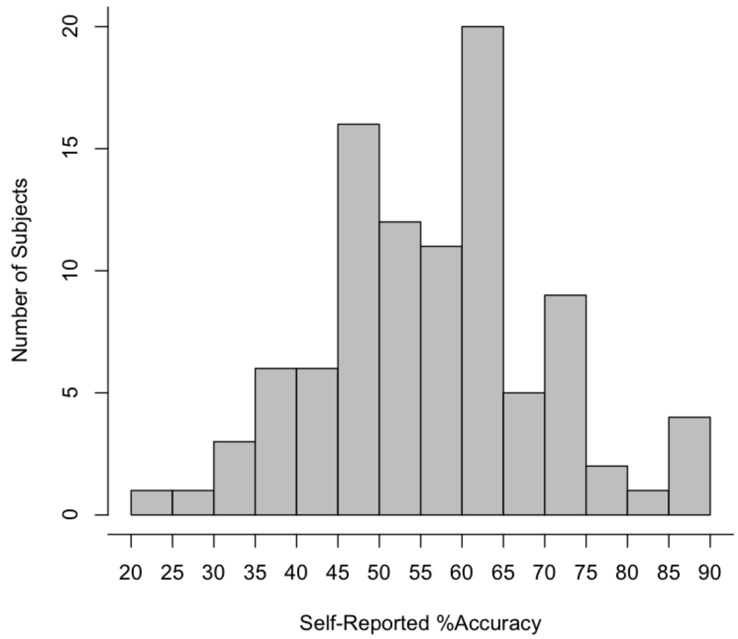

Figure 2. Distribution of percentage of self-reported correct predictions, in Experiment 1

Mouse-trajectory shape. Trials with motion times greater than 5000ms were removed as outliers (19 out of 1960 trials, approximately 1% of the data). For the sake of visualizing behavioral patterns in mouse movements, participants were labeled as "Dishonest" (more than 70% Correct) and "Honest" (less than 55% Correct). The mouse trajectories of "Honest" and "Dishonest" participants were interpolated to 101 time steps (see Spivey, Grosjean, Knoblich, 2005), and superimposed to produce average trajectories, which are depicted in Figure 3. The average trajectories of "Honest" and "Dishonest" participants show more division while reporting "Correct" versus "Wrong". As shown in Figure 3 the average trajectory for "Dishonest" subjects when choosing Correct is shorter and more direct compared to "Honest" subjects; suggesting that on average they experienced less hesitation while choosing the deceptive answer.

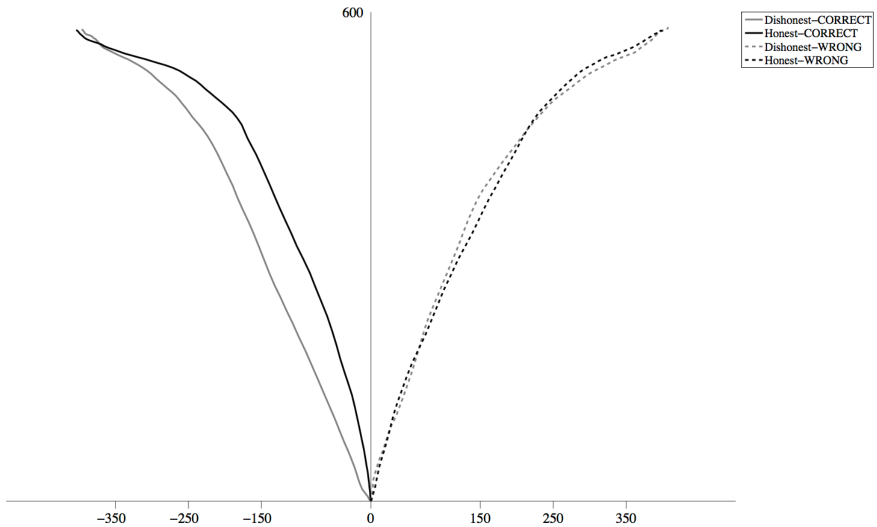

Figure 3 The average trajectories of "Honest" vs. "Dishonest" subjects while choosing Correct (solid line) or Wrong (dashed line) in Experiment 1

Mouse-trajectory properties. Mouse-movement trajectories allow a wide variety of dependent variables that can simply be extracted through analysis of the (x,y) coordinates across time. Even though we analyze each independently, they are not interpreted independently: The measures should point to similar patterns regarding cognitive processes. Put differently, this allows us to "triangulate" from many variables the effects on cognitive processing in "Honest" vs. "Dishonest" subjects. Below are the definitions of eight variables that have been computed and used to characterize the temporal and trajectory behavior for each trial.

Temporal measures*.* The overall time of one trial from the moment a subject sees the page containing two choices (Correct, Wrong) to the moment they click on one of the two is denoted as *total time* (msec). Moreover, *latency* (msec) is the amount of time that the mouse cursor stays in a region with 80-pixel radius from the initiation point. The region is defined as a latency region, reflecting the time period before participants have initiated their decision. *Motion time* (msec) was calculated as the amount of time after the cursor leaves the latency region until the participant clicks on the final answer (Correct or Wrong).

Trajectory measures**.** *Distance* (pixels) is the Euclidean distance traveled by the trajectory from the initiation point until clicking on the final answer (Correct or Wrong). *Distance in motion* (pixels) is the Euclidean distance traveled by the trajectory after leaving the latency region until clicking on the final answer (Correct or Wrong). *x-flips*, a measure of complexity, are the number of times the mouse cursor changes direction along the *x*-axis (i.e., the axis of decision). *x-range* is defined as the absolute difference between the smallest and largest x-coordinate that the mouse reached in transitioning towards the chosen answer. This measure can capture the pull toward the alternative response (relative amount of attraction). Finally, *x-range in motion* is the same concept as *x*-range but calculated merely in the motion time.

The mean value and standard deviation for each variable for Correct and Wrong trajectories are provided in Table 1. It is immediately evident in almost all measures that Correct responses by "Dishonest" subjects tend to show more facilitation: faster times, shorter trajectory distances, simpler trajectories, etc.

We conducted a linear mixed effects model with a fully specified random effects structure for each of the eight dependent variables. As fixed effects we used total accuracy, response type (Correct vs. Wrong), and the interaction term between them. As random effects, we had intercepts for subjects, as well as by-subject random slopes for the fixed effects. A summary of results is provided in Table 2.

The model indicated that total time is significantly predicted by response type, (B = -60.0, p = .014). The model suggests that a trial reported as correct will be about 60ms faster than a trial reported as incorrect. Neither the total accuracy nor the interaction term was significant. The results held the same pattern for motion time, as the response type was a significant predictor (B = -52.1, p = .016).

The analysis revealed that total accuracy had a significant effect on distance (β = -15.8, p = .003). That is, the total distance along the trajectory was shorter for people with a higher percent correct. This suggests that participants who are more likely to be dishonest had shorter and more direct trajectories indicating less hesitation and more confidence. Similarly, total accuracy had a significant effect on distance in motion (β = -13.8, p = .012).

Table 1
Means and Standard Deviations of the Mouse Trajectory Variables by Honesty and Response Type
Experiment 1

	Dishonest				Honest			
	Correct		Wrong		Correct		Wrong	
Variable	M	SD	M	SD	M	SD	M	SD
Total time (msec)	1175.31	584.50	1165.76	641.70	1183.82	495.10	1268.44	630.70
Motion time (msec)	909.00	444.29	919.02	529.50	938.86	425.60	1004.58	542.45
Distance (pixels)	879.24	280.00	870.17	203.50	1012.16	537.90	915.71	342.30
Distance in mot.	739.90	275.63	720.03	219.00	866.35	544.24	776.3	348.18
x-range	462.97	132.37	467.66	110.35	510.65	192.01	472.33	142.34
x-range in mot.	401.02	161.28	407.21	140.53	458.22	221.80	419.15	169.68
x-flips	0.95	0.97	0.98	0.97	1.30	1.31	1.15	1.21

Table 2
Coefficient estimates from mixed-effects models predicting variables in Experiment 1

Variable	Total Accuracy			Response type			Interaction		
	Coeff	Std Err	t.value	Coeff	Std Err	t.value	Coeff	Std Err	t.value
Total time (msec)	-1.92	13.75	-0.14	-60.09	24.63	-2.43*	1.18	10.77	0.10
Motion time (msec)	-3.23	10.91	-0.30	-52.16	21.82	12.39*	-0.45	9.23	-0.05
Distance (pixels)	-15.84	5.49	-2.9**	32.30	25.68	1.26	-16.98	9.39	-1.80
Distance in motion(pixels)	-13.88	5.58	-2.48*	25.91	25.81	1.00	-13.07	9.71	-1.34
x range	-5.12	2.32	-2.2*	11.48	10.23	1.12	-8.86	3.66	-2.42*
x range in motion	-6.64	2.66	-2.5*	9.36	12.44	0.75	-8.47	4.66	-1.81
x flips	-0.04	0.02	-2.4*	0.001	0.06	0.02	-0.01	0.02	-0.85

* p<.05. ** p<.01.

In the case of x-flips (using Poisson distribution), total accuracy also had a significant effect (β = -0.04, p = .016), suggesting that number of x-flips was slightly lower in participants who are more likely to be dishonest. Thus, changes in direction of mouse movements happened less often in “Dishonest” participants, indicating they had more confidence in their decision. We conducted the same model on x-range and x-range in motion. The analysis showed that total accuracy was also a significant predictor of x-range (β = -5.1, p = .027) and x-range in motion (β = -6.64, p = .012). This shows that “Dishonest” subjects illustrated more direct and less curved trajectories. Moreover, we obtained a significant interaction between the total accuracy and response type in predicting x-range. The interaction between subjects' total accuracy and whether Correct was reported as the outcome of the trial provided a significant effect on x-range (β = -8.8, p = .015). Thus, as it is also evident in average trajectories (Figure 3) “Dishonest” participants showed less deviation towards the truthful alternative while choosing the deceptive answer.

While the variety of dependent variables in this experiment might make the audience wonder if we had to run multiple tests for comparing each pair of variables, it is important to note that we conducted one primary model to assess the general trend of differences. All of the defined metrics are pointing to the same direction and the analysis show that they agree on what we capture as the over-all picture.

2.1.5 Discussion

The results from this experiment indicated that people show less complexity in their mouse trajectories when they are being dishonest in incentivized tasks. “Dishonest” participants showed shorter and relatively direct mouse trajectories with no significant signs of hesitation. While being deceptive, their trajectories were more direct with less deviation towards the alternative response, illustrating less effort and time. “Honest” participants on the other hand, were demonstrating hesitation with longer trajectories and more attempts to change the direction of the mouse cursor. Consistent with the second theory (Shalvi et al., 2012; Greene & Paxton, 2009), the results suggest that cognitive processes can indeed be facilitated during dishonest decisions. However, motivation/self-interest is a key task variable that determines which theory will make the right relevant predictions. In this experiment, self-interest was provided by monetary bonus which seems to have triggered temptation and therefore the observed effects in "Dishonest" participants. Next section presents an experiment where there is no monetary reward offered. The main objective of this follow-up experiment is to assess the influence of self-interest in the existing effects.

2.2 Experiment 2

The previous experiment introduced a task in which participants were not explicitly asked to act dishonestly. Rather, the task tempted them to cheat by offering a bonus payment, under the impression that they are detecting a pattern in the coin flips. In a follow-up experiment, I wished to assess if offering extra bonus does indeed induce temptation to lie in participants. I used the same set up as Experiment 1, only without rewarding participants with a bonus for their accuracy.

2.2.1 Method

In this experiment, participants were asked to report their accuracy after each coin-flip and whether they noticed any pattern in the sequence of heads and tails overall. We expected the effects from the previous experiment to be diminished to some extent.

2.2.2 Participants and Procedure

Ninety-five subjects were recruited through AMT. The procedure was the same as Experiment 1, except no bonus was given. Participants were led to believe that the main objective of the study was to assess whether implicit learning of underlying patterns is possible through the process of guessing privately and receiving feedback. Participants were debriefed about the actual goal of the study after completing the experiment.

2.2.3 Results

The number of "Dishonest" participants (who claimed more than 70% correct) was smaller compared to Experiment 1 (9 vs. 16) as well as the maximum self-reported accuracy (85% vs. 90%). Nevertheless, this distribution still showed significant difference from a fair distribution of coin flips ($t(94)=5.04$, $p < .001$). This bias could be caused by people's desire to show successful performance in the task and learning the underlying patterns. Figure 4 shows the distribution of percentage correct reported by all 95 subjects.

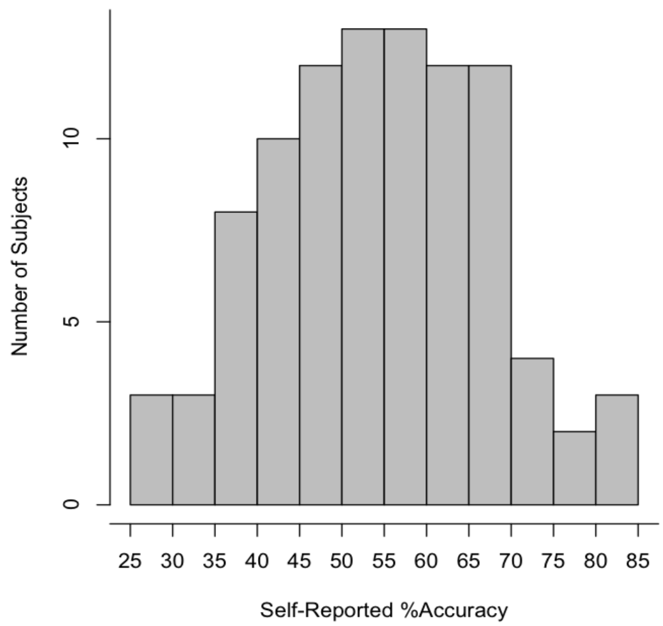

Figure 4. Distribution of percentage of self-reported correct predictions, in Experiment 2

Mouse-trajectory shape. Figure 5 shows the average trajectories of "Honest" subjects compared to "Dishonest" subjects. It appears that, indeed, incentive was driving the cognitive facilitation, as there is not much difference between the "Honest" and "Dishonest" average trajectories while choosing Correct.

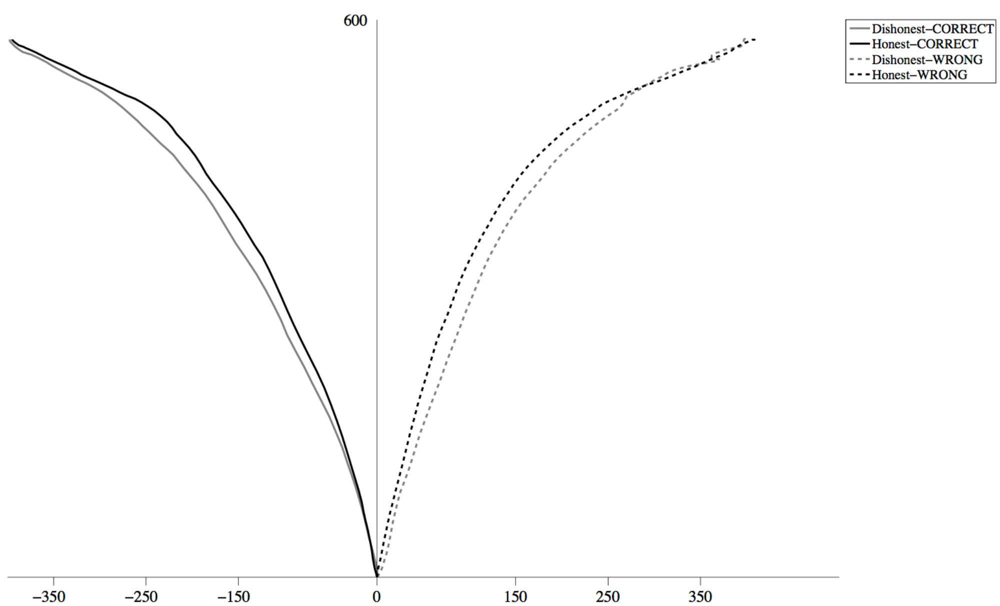

Figure 5. The average trajectories of "Honest" vs. "Dishonest" subjects while choosing Correct (Solid line) or Wrong (dashed line) in non-incentivized condition.

Mouse-trajectory properties. Trials with motion times greater than 5000ms were discarded prior to analysis (0.4%). Mean values and standard deviations for all dependent variables in Correct and Wrong trajectories are listed in Table 3.

The same linear mixed effects model with total accuracy and response type as fixed effects was used to analyze the data from the second experiment. Total accuracy and the response type (Correct, Wrong), did not have any significant effect on total time or motion time. Similarly, I did not find any significant effect on x-range, x-range in motion, x-flips, distance, or distance in motion. Detailed results are provided in Table 4.

Table 3
Means and Standard Deviations of the Mouse Trajectory Variables by Honesty and Response Type in Experiment 2.

	Dishonest				Honest			
	Correct		Wrong		Correct		Wrong	
Variable	M	SD	M	SD	M	SD	M	SD
Total time (msec)	1062.3	506.22	1102.75	595.38	1075.27	501.11	1065.10	517.33
Motion time	795.26	358.58	824.54	423.35	852.38	399.55	848.26	422.70
Distance	893.37	313.83	963.40	388.83	922.17	381.27	962.33	361.73
Distance in mot.	743.66	314.38	814.85	402.81	783.22	381.93	825.72	366.53
x-range	466.43	130.83	487.52	141.46	475.02	156.25	509.40	169.25
x-range in mot.	410.11	158.15	432.56	176.67	425.87	182.98	460.87	197.75
x-flips	1.19	1.27	1.5	1.36	1.02	1.13	1.06	1.05

Table 4
Coefficient estimates from mixed-effects models predicting variables in Experiment 2.

Variable	Total Accuracy			Response type			Interaction		
	Coeff	SE	t.value	Coeff	SE	t.value	Coeff	SE	t.value
Total time (msec)	12.14	11.77	1.03	-1.62	20.97	-0.08	1.40	8.83	0.15
Motion time (msec)	4.4	8.11	0.54	-2.3	17.83	-0.12	2.19	7.43	0.29
Distance (pixels)	-4.85	5.71	-0.85	-20.75	20.17	-1.03	-2.21	9.33	-0.23
Distance in motion	-6.62	5.65	-1.17	-22.40	20.24	-1.11	-2.10	9.37	-0.22
x range	-3.13	2.45	-1.27	-16.77	10.19	-1.64	1.44	4.28	0.33
x range in motion	-3.66	2.98	-1.23	-19.26	12.40	-1.55	1.87	5.13	0.36
x flips	0.02	0.01	1.24	-0.02	0.06	-0.40	-0.03	0.02	-1.20

* p<.05. ** p<.01.

2.2.4 Discussion

As expected, these results suggest that the effects from Experiment 1 are no longer present after discarding the reward and therefore the temptation to lie. The current experiment provides a framework to focus on studying the role of self-benefit in participants' performance. The clear difference observed in the behavior of "Dishonest" and Honest" participants is no more evident after discarding the monetary benefit. The results from these two experiments clearly reiterates the recent findings by Shalvi et al. indicating that lying is only facilitated or "automatic" when there is a self-interest or benefit in doing so.

This, however, raises an additional question: to what extent are the results dependent on the nature of the task? Are the effects observed here driven by the details in task setting rather than a more general underlying cause? The level of difficulty in this task is another concern that needs to be addressed as it could directly affect the performance of participants and the mental processes involved in producing the observed effects. The next experiment is an attempt to increase the difficulty of the task while testing the influence of task properties on the results.

2.3 Experiment 3

In the previous experiments, participants were instructed to click on the box labeled "Correct" if their prediction matched the actual outcome of the coin flip and click on "Wrong" otherwise. One possible criticism of this procedure is the ease of the task; one can automatically choose "Correct" all the time without getting engaged in decision-making processes throughout the experiment. In other words, participants can decide to always immediately click on "Correct" regardless of their actual accuracy. In order to address this concern, Experiment 3 adds more difficulty to the task. Instead of clicking on boxes labeled as "Correct" and "Wrong," participants were instructed to click on the box with a label consistent with their prediction: "Heads" or "Tails." It is hypothesized that by doing so, participants cannot decide about the answer at the beginning of a trial; instead they need to actively track the outcome of the coin flips in order to cheat. Thus, it is impossible to win extra money if they are not actively engaged in the task. If dishonesty is always facilitated in a tempting situation, independent from the task nature, we expect to replicate the results from the previous work.

2.3.1 Participants

Ninety-one participants were recruited online through Amazon's Mechanical Turk (AMT). They were paid $0.40 for their time. A numeric code on the server was assigned to the participants to ensure that they had actually completed the task, and approved their payment on AMT.

2.3.2 Procedure

The procedure is very similar to Experiment 1. Participants were instructed to privately predict the outcome of a computerized coin flip 20 times. When ready, the participants could click on a button labeled as “Flip,” which triggered a computer animated coin flip with an outcome determined by a random generator to have equal probability for heads and tails. Once the coin landed, participants were directed to a page that contained two boxes labeled as “Heads” and “Tails” on the top left and top right of the screen. The mouse cursor was automatically placed on the bottom center of the screen and the participants were expected to move the mouse towards the desired option. The trial only finished after the participants finalized their decision by clicking on one of the responses. If their prediction was consistent with the actual outcome of the coin flip, they were shown a message indicating that they have won a bonus. Otherwise, they were informed that no bonus was awarded. The mouse trajectory and final response for each trial were collected for further analysis. Figure 6 displays the task sequence.

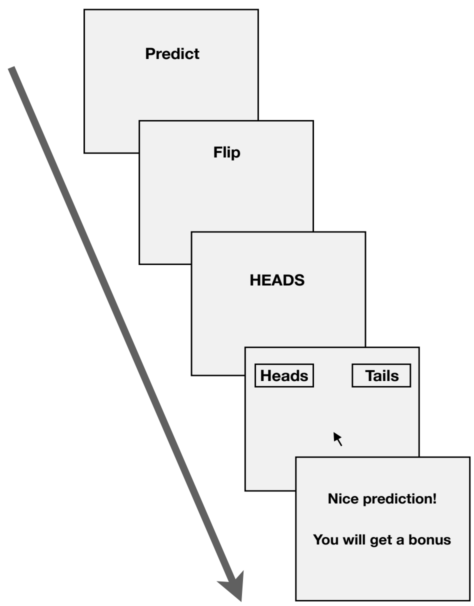

Figure 6. Task sequence: subjects 1) make a prediction, 2) flip the virtual coin, 3) see the outcome and 4) report their prediction by clicking on one of the two boxes on top of the screen (i.e. Heads and Tails) which were assigned to left or right side counterbalanced. 5) They will be informed that they got a bonus (or not) if their report was consistent (or inconsistent) with the actual outcome.

2.3.3 Results

Ten subjects failed to complete all 20 trials and were excluded from the analysis. Data from 81 remaining participants were used to run the statistical analysis. Trials with total times greater than 5000ms were discarded prior to analysis (0.4% of data). As participants' predictions were private, we detected lying by comparing the distribution of self-reported accuracy with the expected distribution of fair coin-flips. The distribution of reported correct predictions (*M*= 12.70, *SD* = 3.1) was significantly different from a fair distribution of random coin-flips (*M* = 10), $t(80) = 7.70$, $p < .001$. The analysis suggests that participants have been dishonest at the group level. Figure 7 shows the distribution of self-reported percent correct. Even though the task set up makes it impossible to distinguish dishonesty on an individual basis, dishonest responses are expected to be more prominent in the rightmost portion of the distribution. Here we investigate the shape and properties of the mouse trajectories for potential honest and dishonest decisions.

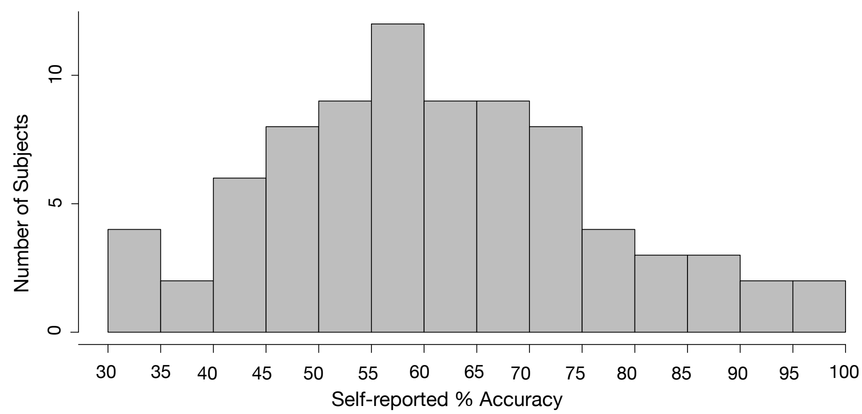

Figure 7. The distribution of percentage of self-reported correct predictions.

Mouse-trajectory shape. Participants were labeled as "Honest" and "Dishonest" based on their performance in the experiment. 31 participants with more than 70% accuracy were classified as "Dishonest" while other participants were considered as "Honest." 70% was chosen as a point where is significantly beyond what was expected from a binomial distribution considering conventional .05 level probability cutoffs.

In order to investigate the average performance of each group, 100 time steps were interpolated to produce average trajectories, which are represented in Figure 8. Trials in which the reported prediction (heads or tails) was consistent with the outcome of the coin flip were considered "Correct" and trials of the opposite kind were marked as "Wrong." Showing a different pattern from the previous findings, average trajectories for correct do not illustrate any significant disparity between "Honest" and "Dishonest" participants. However, surprisingly, the difference appears to be captured by *wrong* trajectories. Figure 8 demonstrates the difference between the current findings (the graph on top) and the results from Experiment 1 (bottom graph).

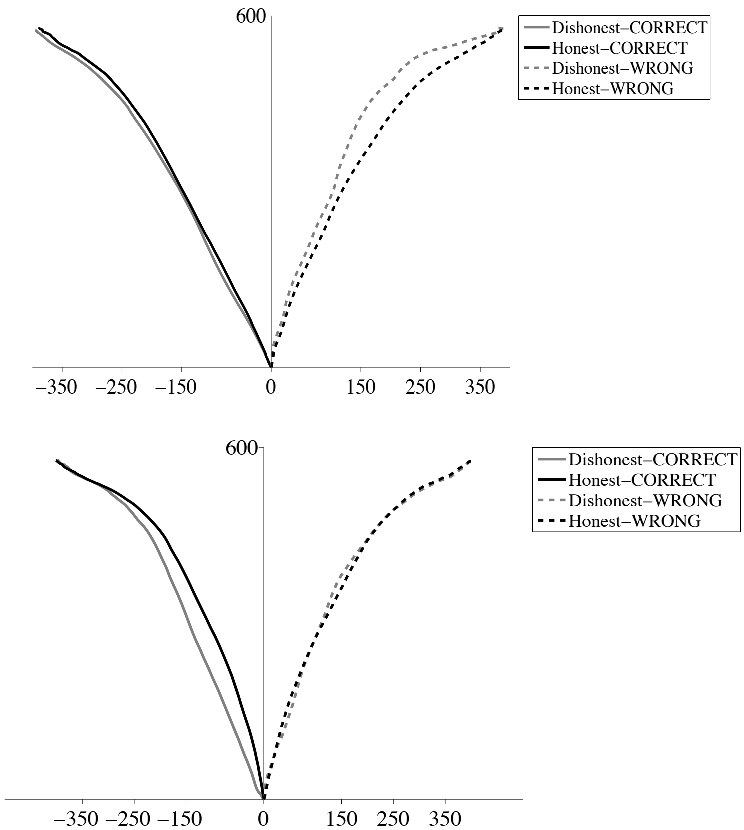

Figure 8. The figure on top shows the average trajectories of "Honest" (black) vs. "Dishonest" (gray) subjects while reporting correct (solid line) or wrong (dashed line) predictions in the current experiment. The bottom figure shows the corresponding trajectories from Experiment 1**.**

A possible explanation, given that there is approximately no difference between "Honest" and "Dishonest" in *correct* trajectories, is the influence of repetition priming. In correct trials participants only have to choose the option that is consistent with what they saw on the screen (heads or tails), thus they are much faster in choosing the matching option. On the other hand, when reporting a wrong prediction, they have to click on an option that is inconsistent with what they were shown on the screen. Thus they are expected to take more time to respond. However, there is a substantial difference between "Honest" and "Dishonest" trajectories in *wrong* trials with "Dishonest" participants exhibiting more complexity when being honest. Their average trajectory captures a desire to report the alternative deceitful option. Although, both groups are slower in reporting an inconsistent response, "Dishonest" people are slower with more deviated trajectories. Next, we compare these trajectory shapes by extracting measures from them and conducting further analysis on their quantified properties in terms of extent, complexity, etc.

Mouse-trajectory properties. The total time and motion time are temporal measures, whereas other variables mainly capture the dynamic changes along the mouse trajectory coordinates. Table 5 provides the mean values and standard deviations of the dependent variables for *correct* and *wrong* trajectories grouped by honesty. For each of the four dependent variables, a linear mixed effects model with a fully specified random effects structure was conducted. I used honesty (Honest vs. Dishonest), response type (Correct vs. Wrong), and the interaction term between them as fixed effects. As random effects, I had intercepts for subjects, as well as by-subject random slopes for the fixed effects.

A noticeable difference in the analysis of the current experiment is that in Experiment1 and 2 total accuracy was treated as both a continuous and discrete effect. However, for simplicity and clarity of presentation, in the current experiment I report the results by classifying the participants into "Dishonest" (total accuracy >= 70%) and "Honest" (total accuracy < 70%). Thus, in the current analysis I chose honesty as a discrete fixed effect. A summary of results is provided in Table 6.

The models for temporal measures showed that response type has significant effects on total time (β = -138.42, p < .001). All participants exhibited shorter reaction times (about 138 millisecond faster) when reporting a correct prediction; by clicking on the response that was consistent with the outcome of the coin flip. The results fit with the intuitive interpretation that emphasizes the effect of repetition priming.

Table 5
Means and standard deviations of the mouse-trajectory variables by honesty and response type

	Dishonest				Honest			
	Correct		Wrong		Correct		Wrong	
Variable	M	SD	M	SD	M	SD	M	SD
Total time (ms)	1021.89	420.70	1184.05	634.22	1025.47	465.72	1138.90	630.03
Distance (pixels)	853.88	271.50	963.79	383.46	874.51	376.02	896.76	357.27
x-range(pixels)	447.63	128.58	484.33	152.07	457.60	181.74	463.55	163.87
x-flips	0.90	1.01	1.46	1.45	0.95	1.07	1.17	1.26

Table 6
Coefficient estimates from mixed-effects models predicting variables

Variable	Dishonesty			Response type			Interaction		
	Coeff	SE	t	Coeff	SE	t	Coeff	SE	t
Total time (ms)	11.17	65.76	0.17	-138.42	33.65	-4.11**	-65.21	72.42	-0.90
Distance (pixels)	14.16	47.58	0.30	-70.45	18.04	-3.90**	-84.88	40.26	-2.10*
x range(pixels)	2.72	22.07	0.12	-25.20	7.32	-3.44**	-26.91	15.49	-1.73
x flips	0.10	0.11	0.90	-.31	0.06	-4.93**	-0.20	0.13	-1.49

* p<.05. ** p<.01.

The model for *x*-range also supports the effect of response type. Reporting a trial as correct is a significant predictor of *x*-range (β = -25.2, p < .001). Correspondingly, we found a marginally significant interaction between response type and dishonesty for *x*-range (β = -26.91, p = .08), suggesting that "Dishonest" participants had less curved mouse trajectories when answering correct versus wrong. However, "Honest" participants experienced a less dramatic change in their *x*-range when reporting a correct versus a wrong prediction. This offers some preliminary evidence that mouse trajectories of "Dishonest" participants were pulled towards the deceitful alternative when they were being honest, whereas "Honest" participants did not show any substantial difference in their trajectories regardless of the response type. Figure 9 illustrates the interaction between response type and honesty.

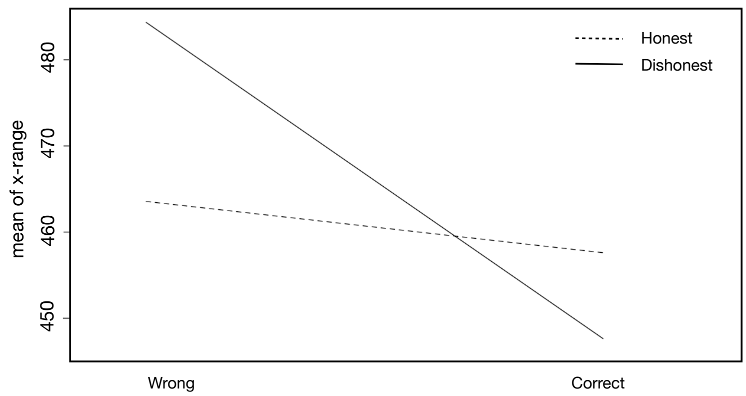

Figure 9. Average x-range grouped by honesty and response type.

The number of x-flips, which can be an indicator of hesitation, is significantly influenced by response type (β = -0.31, p < .001). When reporting a correct prediction, participants changed the direction of their mouse cursor about 30% less compared to trials reported as wrong. This suggests that all participants experienced less hesitation while choosing the consistent response either truthfully or as a deceptive answer.

The other variable that shows the cognitive process correspondent to the mouse trajectory is the absolute distance that the mouse cursor has traveled. Distance is significantly predicted by response type (β = -70.45, p < .001), as correct trials contain shorter and more direct trajectories. Moreover, there is a significant interaction between honesty and response type (β = -84.8, p = .03), suggesting that "Dishonest" participants had longer and more curved mouse trajectories while reporting a prediction as wrong.

2.3.4 Discussion

The difference between "Honest" and "Dishonest" participants, which was previously captured in *correct* trajectories, is now evident in *wrong* trajectories. Put another way, "Dishonest" participants merely exhibit a deviation from others when they are reporting their prediction as wrong (i.e. being honest). In such situations their trajectories demonstrate a pull towards the deceitful answer, indicating a desire for lying. More than facilitation in lying, the results revealed a signature of temptation to cheat in "Dishonest" participants, while being honest. Thus, it appears that changing the nature of the task directly influences the facilitated process. In the current task, it is not dishonesty that is facilitated; rather the more intuitive process that maintains response consistency is the dominant process.

The task setup dictates an advantage for responses that are consistent with the outcome of the coin flip. These responses happen to be the tempting option and therefore make it difficult to study the dynamics of dishonest responses. Although repetition priming makes correct responses easier for all participants, it helps distinguish "Honest" and "Dishonest" participants when they select an answer opposite to what they were shown. It is hard for everyone to report an inconsistent response but more so for "Dishonest" participants, as they attempt to develop a self-serving strategy as well. The results reveals that the task modification forces the difference between the two groups to appear in *wrong* trajectories rather than *correct* trajectories. "Dishonest" participants tend to be slower and more hesitant compared to "Honest" participants. They show longer and more curved trajectories, indicating their desire to lie when they chose the truthful option.

2.4 Discussion

In these series of experiments, I used a novel action-dynamics paradigm to test two competing theories concerning cognitive processes underlying dishonesty. Participants were put in a tempting situation, where they could cheat to earn money. They reported the accuracy of their prediction by moving the mouse cursor to the top right or top left of the screen (labeled as Correct or Wrong). The goal was to track their mouse movements while reaching for the option that can be chosen either truthfully or by cheating. As mouse trajectories are shown to be representative of underlying cognitive processes in decision-making, we expected honest and dishonest decisions to demonstrate significantly different shapes of mouse trajectories. Assuming that honesty is the grounded process, as predicted by the first theory (Gilbert, 1991; Duran, Dale, & McNamara, 2010; Duran & Dale, 2012; McKinstry et al., 2008; Spence et al., 2001), it should happen spontaneously with less hesitation and in shorter time, which results in shorter and more direct mouse trajectories. On the other hand, if lying to serve self-interest is the more facilitated process (Greene & Paxton, 2009) or the "automatic tendency" (Shalvi et al., 2012), as stated by the second theory, we expect dishonest decisions to happen faster with more confidence and less effort, causing shorter and less curved mouse trajectories.

Based on the current findings, people show less complexity in their mouse trajectories when they are being dishonest in incentivized tasks. "Dishonest" participants showed shorter and relatively direct mouse trajectories with no significant signs of hesitation. While being deceptive, their trajectories were more direct with less deviation towards the alternative response, illustrating less effort and time. "Honest" participants on the other hand, were demonstrating hesitation with longer trajectories and more attempts to change the direction of the mouse cursor. Consistent with the second theory (Shalvi et al., 2012; Greene & Paxton, 2009), our results suggest that cognitive processes can indeed be facilitated during dishonest decisions. However, motivation/self-interest is a key task variable that determines which theory will make the right relevant predictions.

Importantly, a feature that distinguishes the two theories is the type of task each implies. It appears that, the second theory is often invoked in tasks where the dishonest response is known in advance and serves the interests of the participant; tasks on which the first theory is based, tend to be prompted deception paradigms, in which self-serving dishonest decisions may not be present. The goal in the present study was to design a task in which we could extract the real-time cognitive processes of participants who are tempted to lie. By doing so, we were able to track the time course of processing, and test predictions from these two accounts.

Considering a more general and task-independent interpretation of the results, we could see the behavior observed from both "Honest" and "Dishonest" participants as a parallel-competition between different alternative choices that are fighting to dictate the final output process (McKinstry et al., 2008). Thus, there may not be an inevitable sequential order to honesty and dishonesty, where one process is always the preliminary tendency, being occasionally blocked by the other one. Rather, we suggest the possibility of a continuous parallel competition between the two options, which can favor any of the two, given the circumstance and the nature of the task.

Experiment 3 suggests that task variables indeed determine whether and how self-serving biases are reflected in cognitive dynamics. The present manipulation of the task seems to offer a more realistic picture of naturalistic deception by capturing the temptation to behave dishonestly. In a tempting situation, dishonesty is not always the facilitated tendency; rather it is one of the competing processes that may or may not come to govern responses. Importantly, it is not easy to disentangle all different task variables that are producing the effects. One cannot be sure if repetition priming is responsible for all of the observed effects. However, these findings start a promising trend for discussion.

Analyzing the dynamics of human hand movements in deceptive behavior reveals a good deal about the decision making processes underlying such behavior. Tracking the action dynamics helped capture a nearly real-time picture of deceptive behavior as it unfolds. Moreover, a signature of hesitation and temptation was found in recorded data as "Dishonest" and "Honest" participants made gradual movements towards an honest or deceitful option. Nevertheless, despite all effort, deception has been shown to be extremely difficult to study. As a rich tapestry of complex and multi-level cognitive processes, deception resists simplicity. It imposes social and moral costs on subjects which makes it extraordinarily difficult to control in an isolated experimental design. On the one hand it is nearly impossible to keep track of one's deceptive behavior but not giving them the fear of being caught. On the other hand, detecting deception post hoc is repeatedly shown to be a true challenge even for experts (Bond, 2008). Moreover, there are many different types of deception present in human behavior while there is no agreed upon definition for this complex phenomenon in the literature. Different types of deceptive behavior, cheating, and dishonesty are driven by a wide range of motivations, moral backgrounds and carried by a combination of cognitive processes. All this makes deception not a perfect area for studying fine grained real-time processes in a controlled and isolated setting. That is why this paper is going to change gears towards a relatively more narrowed cognitive phenomenon that provides an opportunity to get an accurate picture of the involved mental processes.

Chapter 3

3. Action dynamics of Prediction

In the current chapter, *prediction in sequence learning* is the scenario that I will be exploring to make a case for how action dynamics are weaved into the mental processes as they unfold. While deception tends to be an extraordinarily difficult phenomenon to capture, this chapter will take advantage of action dynamics to study a more narrowly scoped and tangible process instead: learning.

In the following sections, I will explain the theoretical background and experimental plan for studying the onset of prediction in environments with different levels of regularity and structure. While there is much research conducted on various aspects of sequence learning in humans, the role of prediction in these learning processes has been neglected. What follows is an attempt to shed light on the emergence of predictive mind and the role of prediction in implicit or explicit learning of sequences. The current chapter provides a literature review on prediction as the main function of the brain while presenting statistical learning as a framework for studying prediction. Further, this chapter lays out a precise experimental design and provides evidence for the role of prediction in statistical learning by offering data from a series of mouse tracking experiments.

3.1 Prediction

Imagine you are reading a book at a coffee shop and, while busy reading, reach for the cup that you think is full of coffee. Your hand picks up the cup with a certain amount of force and it suddenly goes too far high in the air. That is the moment you realize the cup was empty. Where did you get the estimation on how much force you should put in picking up that cup? In other words, where does the prospect of the experience of picking up a cup of coffee come from? The same goes when you go up the stairs and climb through an imaginary extra step. Immediately after your foot goes through the imaginary step your brain realizes something went wrong. In such experiences, everything points out to an *expectation* that the brain held but the input data did not meet.

"Brains, it has recently been argued, are essentially prediction machines" (Clark, 2013, p. 1). In his article on predictive brains, Andy Clark draws evidence on how brains are shown to support perception and action by a never-ending attempt to match the input data to their top-down expectations or as we will refer to it in this paper: *predictions*. Clark suggests that "brains as hierarchical prediction machines" is the approach that provides the strongest explanation, to-date, for how the brain works in unifying perception, cognition, action and even attention (2013). The brain does this through hierarchical generative models, the main responsibility of which is to minimize the error in prediction of input. It appears that the higher-level systems in the brain are constantly trying to compare the bottom-up input data to the models they hold about the cause or source of the input signal. Any error in these expectations would be corrected to help the brain adapt to the environment and increase accuracy in allocating the source of input.

The top-down expectations would also dictate the lower-level systems on how to capture the sensory input. In other words, a copy of the top-down expectations would be sent to the lower level systems to structure the bottom-up input in the most expected pattern. The brain will not form a model of the world by passively accumulating input data from the lower level systems; rather it tailors the input data to fit one of its best models of the current state of the world. A very efficient way for encoding input data in a hierarchical generative model is to code only the surprising aspects of the input. That means the only necessary coding is where the data are not consistent with the predictions of the model. Otherwise the model can easily fill in the gaps and make sense of the input. This kind of coding is known as predictive coding and is not an invention of the computational neuroscience field. Predictive coding was first used as a data compression technique (Said & Pearlman, 1993), where in image transmission the only data transferred from an encoded image are the values of the unpredictable pixels.

In an account of unifying perception and action by means of prediction it is important to note that they all work towards active minimization of prediction error. Learning is the process through which the internal models get updated to minimize the prediction error. Thus, prediction and learning are closely entangled. In the current section, I focus on a certain type of learning as it is the main theme of this chapter: sequence learning. What follows is a general review on sequential learning through the lens of prediction and how the notion of predictive behavior plays a role in this area. By that, I hope to lay out a platform for the experimental work presented in the future sections.

3.2 Sequential Learning

"*Thinking of going to the next pattern in a sequence causes a cascading prediction of what you should experience next. As the cascading prediction unfolds, it generates the motor commands necessary to fulfill the prediction. Thinking, predicting, and doing are all part of the same unfolding of sequences moving down the cortical hierarchy.*" (Hawkins & Blakeslee 2004, p. 158).

An ongoing debate in the field of sequential learning establishes a dichotomy between retrospective or memory-based learning versus predictive or implicit learning. Some researchers believe that there are two separate systems underlying implicit and explicit learning while others see these two as functions of the same system. The school of thought which believes that people can implicitly learn the abstract rules in a sequence of stimuli is referred to as the two-system view (Mewhort and Jamieson, 2009). On the other side of the debate are researchers who argue that this performance in learning sequences is not the result of a separate learning system a.k.a implicit learning. Rather, it is only due to recent memory of the local structures in the sequence. It is not necessarily a predictive learning system but a retrospective simple memory retrieval that gets better by more practice and more redundancy in the sequence (Jamieson & Mewhort, 2009; Cleeremans & McClelland, 1991).

3.3 Serial Reaction Time Paradigm

One of pioneering frameworks in studying sequential learning is the Serial Reaction Time task introduced by Nissen and Bullemer (1987). In their initial experiment, subjects saw lights appearing at four different locations on a computer screen and were instructed to press one of the four keys each of which were assigned to a particular light location. Subjects were assigned to two different conditions with the main difference being in the stimuli sequences; in the repeating sequence a same pattern of sequence repeated over and over while in the random condition the appearance of lights on the screen was chosen to be random and followed no specific pattern.

Through 8 blocks of 100 trials subjects in the random condition showed no meaningful difference in their reaction time, whereas in repeating condition subjects improved significantly showing 163ms decrease in their reaction time. Subjects' accuracy was high throughout the task but in the repeating condition they showed improvement by practice, especially in the initial blocks. Almost every subject in repeating condition reported that they had noticed a pattern in the sequence of stimuli and reproduced the pattern by pointing to the locations on the screen in the order they believed was the correct order.

In a follow-up experiment, Nissen and Willingham (1989) recruited patients with amnesia to compare their performance with control healthy subjects. The authors suggest the fact that patients with amnesia could learn the task despite severe impairments in consciously recalling verbal and nonverbal information provides support for the separateness of the procedural system.

To explore the retrospective theory of sequential learning, (Jamieson & Mewhort, 2009), designed experiments as well as simulations that used modifications of existing memory retrieval models to capture participants' performance in sequential learning. Their empirical results are consistent with previous results (Nissen & Bullemer, 1987; Nissen & Willingham, 1989), showing that more structured sequences cause faster reaction times, even though participants can not explicitly point out the rules of the grammar. However, their simulation results replicate performances very similar to SRT task by only relying on memory retrieval and no learning. They argue that the faster responses in repetition phases of SRT task denotes that people could be using the repeated local structures in a memory based learning.

In another critique of SRT task, Jackson and Jackson (1995) argue that studies of implicit learning including the famous Serial Reaction Time task may not be assessing what they are widely believed to be. The authors mention how the measures used in such tasks do not provide any evidence for having an entirely different subsystem for implicit learning, as the task does underestimate different sources of information available to subjects. They demonstrate that subjects' knowledge of serial-order could be independent of subjects' knowledge of the statistical structure of the sequence.

A recent experimental paradigm designed to explore the retrospective and predictive processes underlying sequence learning is presented by Duran and Dale (2010). This paradigm, inspired by Hunt and Alsin (2001) and Nissen and Bullemer (1987), requires subjects to predict the location of the next appearance of the target on the screen. Participants have to move the mouse-curser to where they think the next target will show up and click on the target as soon as it appears. The coordinates of the mouse-cursor are recorded while participants are moving towards the predicted location. Compared to the classical paradigms of serial learning, this task is simple, short, and only captures the start of the learning behavior. Nevertheless, these continuous movement patterns provide an advantage for this study as it is possible to observe and analyze not just the final decision of the participants but the process that gave rise to that decision and the corresponding reaction time. The dynamics of making that decision, as it unfolds in hand movements through time, reveal the initial mental processes underlying the learning behavior. This could offer further explanations for how retrospective and predictive models interact in sequence learning. The results suggest that most participants do, in fact, exhibit predictive behavior through different strategies. Although the participants may not be predicting a particular location, especially in a randomly generated sequence of target locations, they show anticipatory behavior of some kind that could lead to more specific predictions when there is more regularity in the sequence of stimuli presentation. The question arising is: when does the change in prediction pattern happen?

What is the golden regularity point at which most participants start going for more particular predictions rather than being in a general anticipatory mode? How long does it take for the participants to find assurance in structure of the environment and begin to predict more specifically and more confidently? These questions along with other concerns of the sort provide the theoretical motivation for the series of experiments presented here.

Each of the following experiments attempts to address different aspect of the questions that arise while exploring the onset of prediction in humans. The gradual modifications in the experimental design pursue the goal of capturing a rather realistic image of predictive behavior in environments with different levels of structure and uncertainty.

3.4 Experimental Framework

I shall take advantage of this section to establish the main experimental framework for all the studies proposed in the current document. This framework is inspired by Duran and Dale (2010); utilizing the same paradigm but different patterns of stimuli. The goal is to switch from presenting a random sequence of target locations to a more regular and predictable sequence of locations. In a between-subject design I wish to study the level of regularity in the environment which triggers predictive behavior in participants. Detailed modifications for each experiment would be discussed in the corresponding section.

3.4.1 The Stimuli Format

The interface is designed in Adobe Flash and consists of a fixed size 500-by-500-pixel white region in the center of the user's Internet browser window. The target is a 35-pixel diameter black circle that periodically shows up at one of the four corners of the white region. For each participant the circle appears 48 times with intervals of 750 milliseconds. The sequence in which order the target circle appears at corners of the white region is set differently in each of the three experimental conditions. The key constrain is that the target circle could not appear in the same location consequently. Moreover, the circle has to appear an equal number of times in each of the four positions (1,2,3,4) throughout the 48 trials. A view of the task interface is shown in Figure 10.

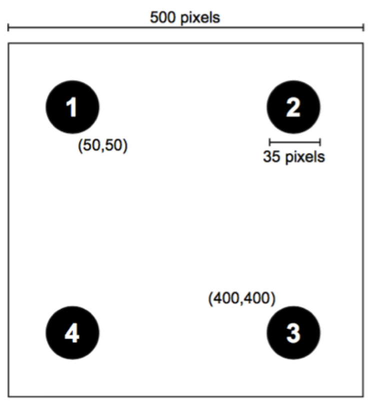

Figure 10. A screenshot of the task environment on a web browser (Dale et al. 2012)

3.4.2 The Stimuli Regularity

The major difference between experimental conditions lies in the structure of the 48 trial sequence. Changing the regularity in the pattern of presented stimuli provides the difference between conditions that participants are assigned to. There is more than one solution for producing sequences with a desired level of regularity. Here, I briefly describe two common approaches used in the literature for producing sequences that follow specific patterns of regularity.

A well-known approach in the field is to use a finite state language to produce grammatical sequences (Chomskey & Miller 1958, Reber 1967,Reber & Allen 1978). A set of letters (or numbers) is used as the vocabulary and combined based on a set of rules for constructing sentences that follow the grammar. A hypothetical state diagram for combining the numbers 1,2,3, and 4 based on a set of grammatical rules is shown in Figure 11. As shown in the figure, every sentence starts from $S_{init,}$ the initial state, and goes on to a state S_j. The end of the sentence is indicated by reaching S_{out} or the final state. The sentences could contain at least 3 letters. The rules of the grammar forbid any immediate repetition of the same letter, and therefore no recurrent loops exist at any state.

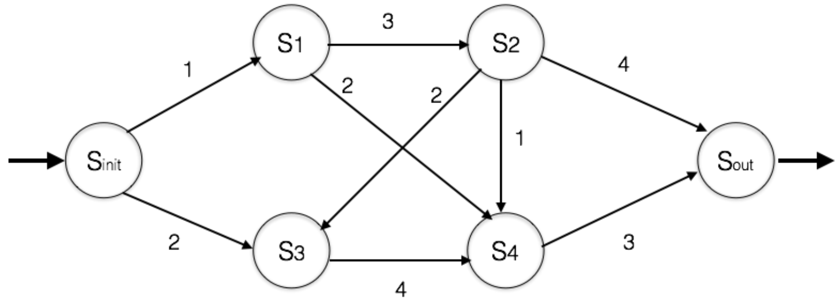

Figure 11. Hypothetical state diagram of the grammar used to generate the stimuli.

Another common approach used in the field of statistical learning is to define a predictability or redundancy measure for a sequence. Using this measure will make it possible to produce and compare sequences of letters or numbers with different levels f regularity. One such measure is introduced by Jamieson and Mewhort (2009). *G* is an indicator of redundancy in a sequence. The measure G quantifies the amount of structure in each grammar and is calculated based on Shannon's sequential uncertainty (Shannon & Weaver, 1949; See Jamieson & Mewhort, 2005, 2009 for details).

For the purpose of this dissertation, the redundancy measure G works efficiently. It is easy to compare the sequences and assign participants to conditions with different levels of regularity. Moreover, Duran and Dale (2010) have applied the same measure for the original experimental stimuli. Therefore, the stimuli for the current study were produced by setting G at certain levels in order to control the level of regularity in the sequence.

3.4.3 Measures

An important step in analyzing such rich dataset of nearly continuous action dynamics through time is to introduce proper metrics and variable to study. Here, I focused on calculating measures that help disentangle predictive and reactive performances in different trials. These measures are defined and explained below.

Distance to Next: This measure indicates the initial distance of the mouse cursor from the next target location. The Euclidian distance (in pixels) from the mouse cursor to the target position is calculated right at the moment the target appears at that location. If the participant has predicted correctly, this distance should be near 0 as the mouse cursor lands on top of the target circle exactly when it appears on the screen. On the other hand, if the mouse stayed near the last position of the target or in the center waiting for the next target to show up, they are most likely taking a reactive approach. Thus, *distance to next* is an indicator of predictive versus reactive behavior and has a negative correlation with prediction. Figure 12 (left) illustrates the *distance to next* in a snapshot of the task.

Distance from Previous: This measure is calculated to account for failed prediction attempts. A subject may perform a predictive behavior by moving the mouse towards where they believe is the next target position, but they might very well be incorrect in their prediction. This is still a predictive behavior and needs to be acknowledged even though was not successful. Initial *distance from pervious* is therefore calculated as the maximum distance from the previous target at the 750-ms mark. The smaller this value, the closer the mouse cursor was to the previous location of the target and the less effort from the subject to get engaged in predictive behavior. On the right panel of Figure 12 you can see *distance from previous* denoted on the task screen.

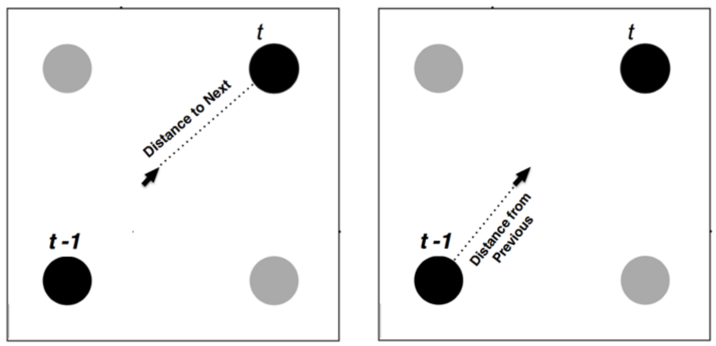

Figure 12. From left to right the figures illustrate the *distance to next* and *distance from previous* respectively.

3.5 Experiment 1: Random and structured

The goal of this experiment is to present a random sequence of target locations in the first half of the task and switch to a more regular and predictable sequence in the second half. I hope to show whether and how participants adjust their strategies to fit the changes in patterns of sequences.

3.5.1 Participants

The participants were recruited through Amazon Mechanical Turk, a crowd-sourcing website where researchers post desired tasks and online users will participate in the experiment. Monetary compensation was paid to the participants in exchange for their time. A total of 70 participants were recruited for this task, out of which only 66 completed the task providing reliable data. All participants received a monetary reward of $.75.

3.5.2 Stimuli

There were three main conditions designed for this experiment. In all conditions the first 24 appearances of the target circle were chosen to be random based on a sequence with 25% regularity. The second 24 appearances were distinguishing between the conditions. Subjects were randomly assigned to 55%, 89% and 100% regularity conditions which means after the first random sequence of target locations they got assigned to conditions with low, medium and high regularity respectively. Regularity of the sequences were calculated based on a measure introduced by Jamieson and Mewhort (2009). This measure, *G*, is the extent to which a 48-position sequence is redundant. Here, the sequence for each condition was constructed by concatenating a 24-position random regularity and 24-position more regular sequence borrowed from Dale et al. (2012). Table 7 shows the random section of the stimuli sequence as well as the regular sections.

Table 7
Blocks of Stimuli Sequence with different regularity levels

Regularity	Sequence
G = .25	4−2−3−2−1−2−1−2−4−1−3−2−1−4−1−3−1−3−2−4−1−3−1−3
G = .55	4−2−3−4−2−1−3−1−3−1−3−1−3−4−2−4−2−4−2−1−3−1−3−1
G =.89	2−3−4−1−2−3−4−1−2−3−2−3−4−1−2−3−4−1−2−3−4−1−2−3
G = 1.0	4−3−2−1−4−3−2−1−4−3−2−1−4−3−2−1−4−3−2−1−4−3−2−1

3.5.3 Procedure

When participants choose to join our experiment on Amazon Mechanical Turk, they are redirected to a webpage that hosts the Adobe Flash interface for the experiment. The program randomly assigns each of the participants to one of the three conditions: 55%, 89%, and 100% regularity. Participants were informed that they would see a black circle periodically appear at each of the four corners of the screen. They had to click on the target circle as soon as it popped up. All participants had to go through a 24-position random sequence of target locations regardless of their assigned condition. Following this phase, they had to perform on a more regular pattern of location sequences based on the conditions they were assigned to. This study has a between subject design. The circle appears once every 750milliseconds at a corner determined by the predefined sequence of locations. They could move or hold the mouse cursor anywhere on the screen while waiting for the target to show up at the next location. While the mouse cursor is travelling around the screen the coordinates of the cursor position are recorded every 25 milliseconds. Thus, the trajectory of the movements made by each participant throughout each trial was saved and reported at the end of the experiment. The mouse coordinates were being constantly recorded throughout all 48 trials. At the end of 48 trials, participants are asked if they found any specific pattern in the sequence of target locations.

3.5.4 Results

Data from 66 participants were imported to extract measures and examine changes in mouse trajectories. Example mouse-trajectories for two subjects are depicted in Figure 13. The subject on the right has spent a lot of time in the center whereas the figure on the left is showing a subject with more predictive trajectories, regardless of their accuracy.

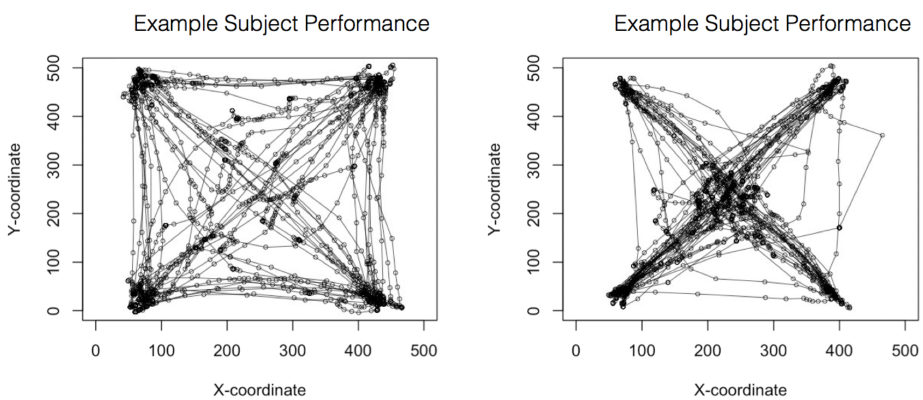

Figure 13. Example mouse-trajectories for two subjects through all 48 trials.

Distance to next as an important measure of predictive behavior was analyzed on a trial-by trial basis in each condition. I expected subjects in 100% regularity condition to show the lowest *distance to next* values and the 55% regularity condition to have the highest average *distance to next* values. In consistence with expectations, condition 55% showed the longest distances in pixels (M_{dist_next} = 273.62). Condition 89% was the second highest *distance to next* in pixels among all conditions (M_{dist_next} = 234.11), whereas condition 100% with the most regularity had the shortest *distance to next* average (M_{dist_next} = 227.28). Figure 14 shows trial-based boxplots for changes in *distance to next* values for all three conditions. Each panel illustrates the values for one of the conditions through 48 trials.

I conducted a linear mixed effects model with subject as a random factor (Baayen, Davidson, & Bates, 2008), and condition (55, 89, 100) as fixed factor. The results confirmed that subjects in condition 55 showed significantly longer *distance to next* throughout the experiment (β = 39.51, p = 0.02). Condition 100, although having shorter *distance to next* values on average, was not significantly different from Condition 89 (β = -6.83, p = 0.6). This is while participants in every condition show high *distance to next* values when they are in the first half of the experiment with random sequence of target locations. That means when there was less randomness in the sequence, participants in condition 100 (M_{dist_next} = 197.14) and condition 89 (M_{dist_next} = 205.51) were significantly better in making predictions compared to participants in condition 55 (M_{dist_next} = 267.28). The pattern of results stays the same, but more significantly so, when the model is run on only the second half of the trials for each condition. The model showed reliably significant decrease in predictive behavior for condition 55 ($\beta = 61.77$, $p = .006$).

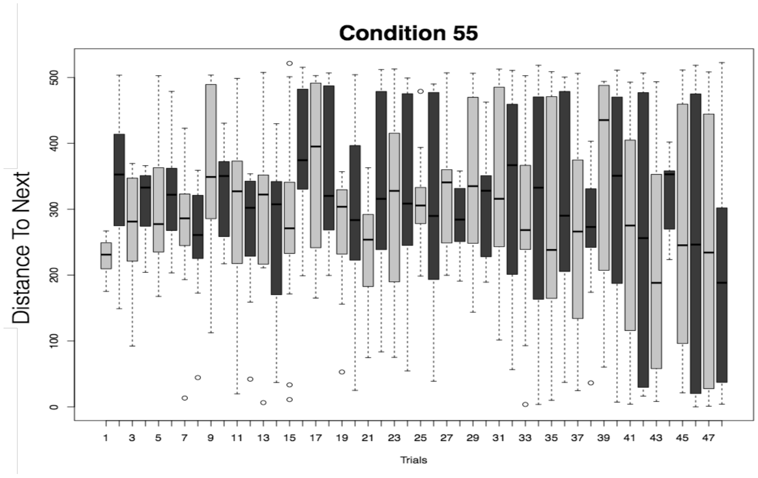

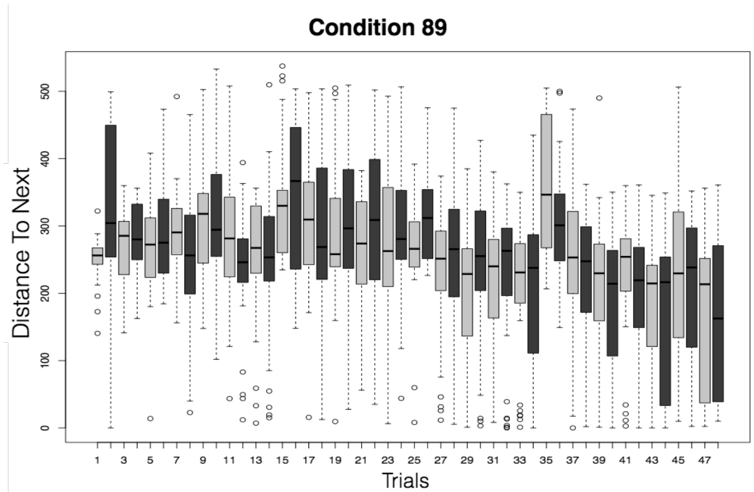

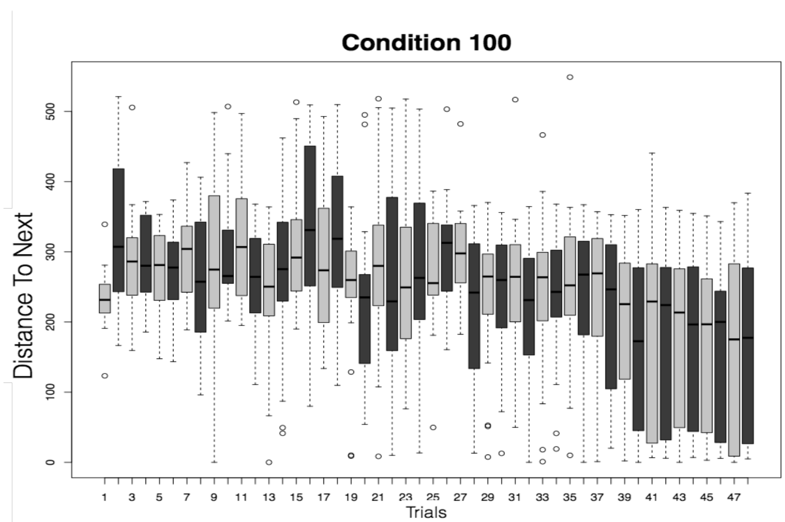

Figure 14. Trial-based boxplots for changes in *distance to next* values. Each panel illustrates the values for one of the conditions through 48 trials.

Figure 15 shows a plot of average *distance to next* values through 48 trials for each condition. Each curve illustrates the average values for one of the conditions through 48 trials. As it is apparent in the curve for condition 89% (red line in Figure 15) there are some trials, which do not obey the general pattern of *distance to next* values. It seems like in these trials something extraordinary happens that causes a sudden decrease in *distance to next* values, suggesting that participants felt more willing to make predictions in these specific trials. The nature of this phenomenon, however, is still ambiguous. This ambiguity provides the motivation for conducting follow-up experiments where there are more than three levels of regularity as conditions. Testing more sequences with variant levels of structure would capture a more accurate picture of the phenomenon and how prediction changes as the environment becomes more and more structured. Next chapter is an attempt to include more experimental conditions in order to provide fine-grained analyses of predictive behavior in statistical learning.

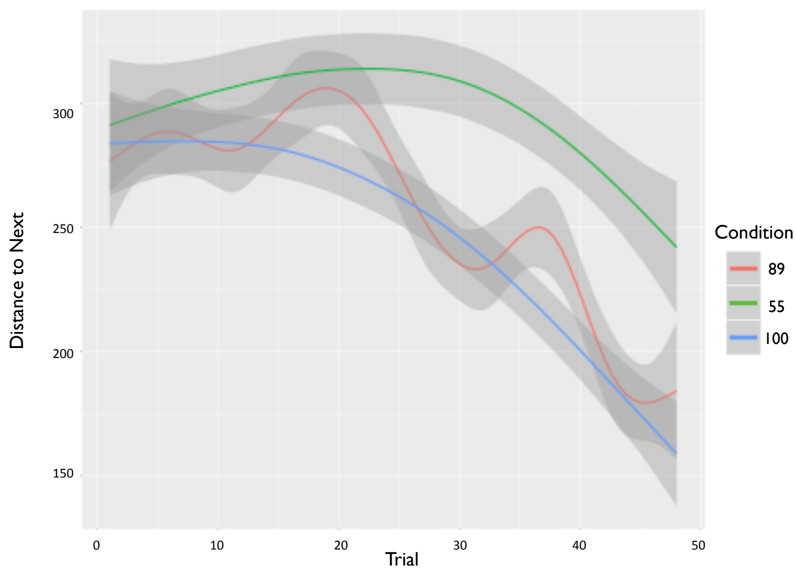

Figure 15. Average *distance to next* values for all there conditions.

3.6 Experiment 2: More Sequences

This experiment is in fact a more advanced version of the preliminary experiment presented in the previous section. In this version, more conditions (i.e. levels of regularity) are added in order to test the role of prediction in sequential learning under more controlled levels of structure in the environment. As described in the previous experiment, participants are presented with a sequence of target locations. This sequence is random in the first half and becomes more structured in the second half. I hope to show whether and how participants adjust their strategies to fit the changes in patterns of sequences. Having more conditions provides a better view of how participants interact with the stimuli and what drives their potential strategies.

3.6.1 Participants

Participants were recruited through Amazon Mechanical Turk. There were more participants acquired in this version of the experiment as I intended to get a more precise view of the phenomenon with more power in detecting any potential effect. A total of 170 subjects were recruited for this task; 162 completed the task, providing reliable data. Participants were compensated $.75 for their time.

3.6.2 Stimuli

There were 7 conditions designed for this experiment. In all conditions there were 48 trials, where each trial was defined as a single appearance of a circle in one of the corner locations on the screen. The first 24 appearances of the target circle were chosen to be completely random (no G assigned). The second 24 appearances were the ones distinguishing between the conditions. Participants were randomly assigned to sequences with the following regularity levels: 44%, 55%, 68%, 76%, 79% 89%, and 100%. In each condition after the first random part of the sequence, participant was presented with the rest of the sequence with one of the listed regularity levels. The levels were chosen to cover a spectrum of low (44%) to high regularity level (100%). From here on, the conditions are named after their regularity level, thus they will be referred to as condition **44**, **55**, **68**, **76**, **79**, **89**, and **100**. AS mentioned earlier, regularity of the sequences was calculated based on G, a measure introduced by Jamieson and Mewhort (2009). The sequence for each condition was constructed by concatenating a 24-position random regularity and 24-position more regular sequence. Table 8 shows the stimuli sequences in detail.

Table 8
Stimuli Sequences For All Conditions

Condition	Sequence
44	1–4–3–2–4–1–3–2–4–2–3–1–3–4–2–1–2–4–1–3–2–3–4–1–2–1–4–3–4–3–2–4–1–2–3–2–1–2–1–4–3–4–3–2–3–2–3–2
55	1–4–3–2–4–1–3–2–4–2–3–1–3–4–2–1–2–4–1–3–2–3–4–1–4–2–3–4–2–1–3–1–3–1–3–1–3–4–2–4–2–4–2–1–3–1–3–1
68	1–4–3–2–4–1–3–2–4–2–3–1–3–4–2–1–2–4–1–3–2–3–4–1–4–3–4–3–4–3–2–1–4–3–2–1–4–3–2–1–4–3–2–1–4–3–2–1
76	1–4–3–2–4–1–3–2–4–2–3–1–3–4–2–1–2–4–1–3–2–3–4–1–3–2–3–2–4–1–3–2–4–1–4–1–3–2–3–2–3–2–4–1–3–2–4–1
79	1–4–3–2–4–1–3–2–4–2–3–1–3–4–2–1–2–4–1–3–2–3–4–1–2–3–1–4–1–4–1–4–2–3–1–4–2–3–1–4–2–3–1–4–2–3–1–4
89	1–4–3–2–4–1–3–2–4–2–3–1–3–4–2–1–2–4–1–3–2–3–4–1–2–3–4–1–2–3–4–1–2–3–2–3–4–1–2–3–4–1–2–3–4–1–2–3
100	1–4–3–2–4–1–3–2–4–2–3–1–3–4–2–1–2–4–1–3–2–3–4–1–4–3–2–1–4–3–2–1–4–3–2–1–4–3–2–1–4–3–2–1–4–3–2–1

3.6.3 Procedure

The program randomly assigned each of the participants to one of the conditions: 44, 55, 68, 76, 79, 89, and 100. Identical to the procedure in Experiment 1, participants were informed that they would see a black circle periodically appear at each of the four corners of the screen. They had to click on the target circle as soon as it appeared. While the mouse cursor was travelling around the 500-by-500-pixel white region on the screen, the coordinates of the cursor position were recorded every 25 milliseconds. Similar to the previous experiment, participants had to go through a 24-position random sequence of target locations followed by a more regular pattern of sequence based on the condition they were assigned to.

3.6.4 Results

Eight participants were discarded from the analysis, as they did not complete the task. Out of the 162 remaining participants there were an average of approximately 23 subjects assigned to each condition. The participants in 100% regularity condition were expected to demonstrate the most predictive behavior and therefore the shortest *distance to next* average. As seen in Experiment 1, the hope here was to replicate the ascending trend of *distance to next* values as the conditions get less and less regular and predictable. Consistent with expectations, condition 100 showed the shortest average *distance to next* (M_{dist_next} = 242.19, sd= 127.40) while condition 44 showed the longest average distance in pixels (M_{dist_next} = 309.5, sd= 111.49). This is while participants in every condition show high *distance to next* values when they are in the first half of the experiment with random sequence of target locations. Figure 16 shows the decreasing pattern in *distance to next* as the sequence gets more structured. Moreover, Figure 17 shows the difference in average distance to next between the random half and the structured half of each condition. Detail values of variables are reported in Table 9.

Table 9
Means and Standard Deviations of the Mouse Trajectory Variables by Condition in Experiment 1 (New Sequences).

	Time and Trajectory Variables					
	Distance to Next (pxl)		Response Time (ms)		Distance from Prior (pxl)	
Condition	M	SD	M	SD	M	SD
44% (N= 24)	309.5	111.5	704.7	387.39	144.43	148.44
55% (N= 16)	285.36	137.3	602.66	341.94	224.45	170.02
68% (N= 22)	269.3	127.97	638.54	347.44	200.93	151.36
76% (N= 18)	281.55	106.73	550.98	176.18	192.69	127.15
79% (N= 29)	284.64	125.84	602.8	229.1	201.29	148.08
89% (N= 32)	250.23	131.67	645.62	716.93	209.11	141.46
100%(N= 21)	242.19	127.40	571.65	233..02	207.56	134.49

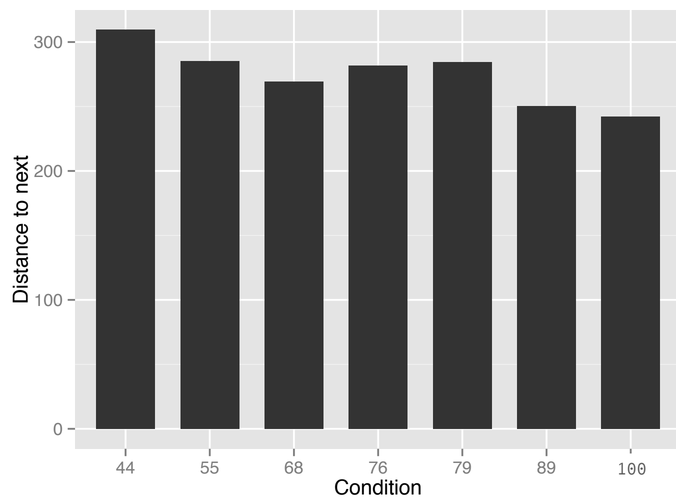

Figure 16. Average distance to next for all conditions in Experiment 2.

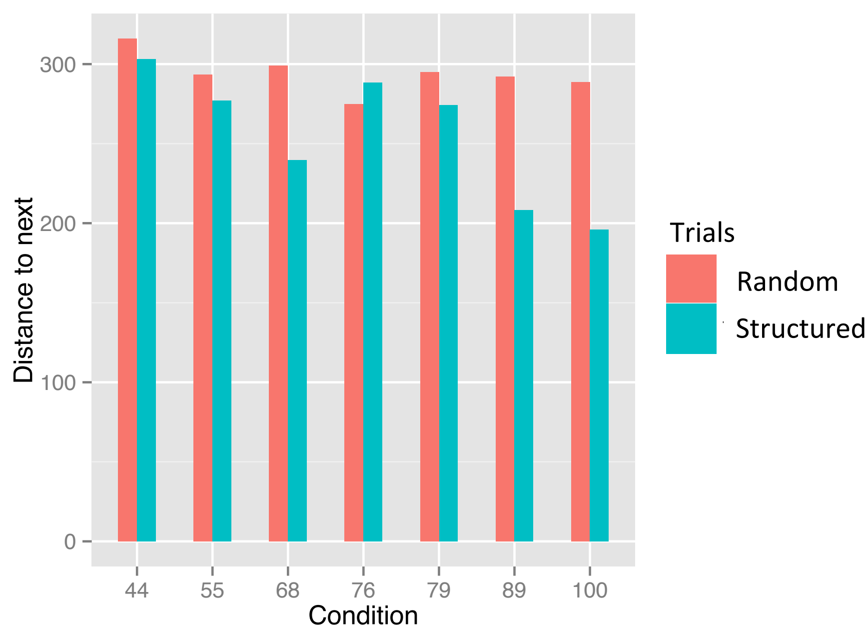

Figure 17. Comparison of average *distance to next* in random and structured parts of the sequence for each condition.

I conducted a linear mixed effects model with Subject as a random factor (Baayen, Davidson, & Bates, 2008), and condition (44,55,68,76,79,89,100) as a numerical fixed factor. The model showed reliably significant decrease in predictive behavior for conditions with lower regularity. The results confirmed that subjects in conditions with less regularity showed significantly longer *distance to next* throughout the experiment (β = -5.38, $p < 0.001$). The pattern of results stays the same, but more significantly so, when we run the model on only the second half of the trials for each condition. This suggests that as expected, people tend to apply more predictive strategies as the environment gets more structured. *Distance to next* as an indicator of predictive behavior shows reliable decrease as we introduce more structure.

Further, I conducted a linear mixed effects model using condition as a categorical factor. All conditions were compared against condition 100. I expected the performance in lower regularity conditions to show significant difference compared to condition 100. This is while the more regular conditions are expected to exhibit a similar performance to what we see at condition 100. In this experiment 68% regularity seems to be the level that puts participants in predictive mode as their performance starts to be significantly different from the levels below it.

The results approximately follow the decreasing pattern that was expected (see Figure 16). However, the conditions 76 and 79 appear to break the pattern by a surprising divergence from the general decreasing trend. Despite being a highly structured sequence, the average *distance to next* at these two conditions show a significant difference from what we see in a perfectly regular sequence (100% regularity). This indicates that, overall, participants had a difficult time predicting the location of next stimuli in these conditions even though the general level of regularity was relatively high in the sequence. This discrepancy raises important questions about the structure of the sequence and what other factors besides the over-all regularity level in a sequence can affect the predictive behavior in participants. One graph that could shed light on the detailed transitions throughout each condition is the average curve of distance to next on a trial-by-trial basis. Figure 18 provides a graph with such curve for all conditions including 76 and 79. The graph shows how the value of *distance to next* (averaged over all participants) changes as a new target in the sequence is revealed.

A closer look into the curves associated with conditions 76 and 79 shows a wave-like pattern in the sequence. It appears that there might be some unintentional local structures present in parts of the sequence. Condition 76 is particularly oscillating quite often with raises in the *distance to next* followed by falls. In order to get an even closer look into these oscillations, Figure 19 (top panel) provides a boxplot of the average performance of participants in condition 76 on a trial-by-trial basis. You can see the exact order of the sequence and the changes in average *distance to next* as the trials unfold. The figure makes it clear that in some cases the participants were led to believe that a pattern is happening and as soon as they got into predictive mode the pattern suddenly breaks and leaves them with an error in prediction. This error would effect a measure like *distance to next* as they have moved their mouse away from the target even though they were engaged in predictive behavior. This is where looking at another measures such as *distance from prior* could provide more insight into the subjects' behavior. Bear in mind that the shorter this distance the more reactive the participant has performed, as this is a metric that shows how far the mouse cursor was from the previous target when the new one appeared. If the participant has started departing from the previous location in an attempt to make a prediction this measure still captures their predictive behavior even though they may end at the wrong location.

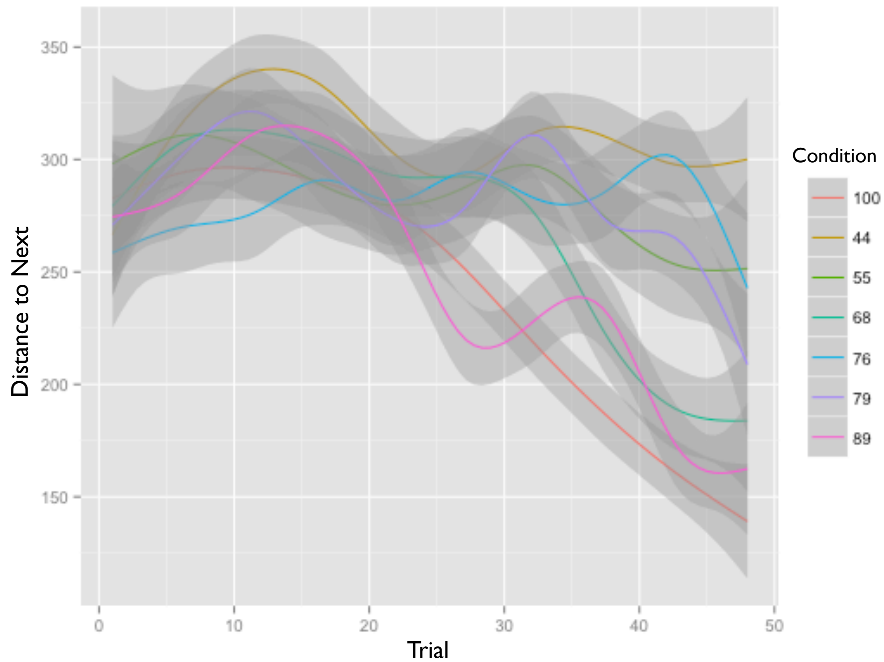

Figure 18. Average distance to next for all conditions throughout 48 trials

Figure 19 (bottom panel) depicts the average *distance from prior* on a trial-by-trial basis. I suspect that the local structure, even though observed only over a few trials, could have a noticeable effect on the predictive behavior in participants. This suggests that a sequence may be globally well structured (76% or 79% regularity) but the local structure could cause prediction errors and therefore not support predictive behavior. Thus, the general learning of the sequence will not look as good as expected for the level of regularity at that sequence.

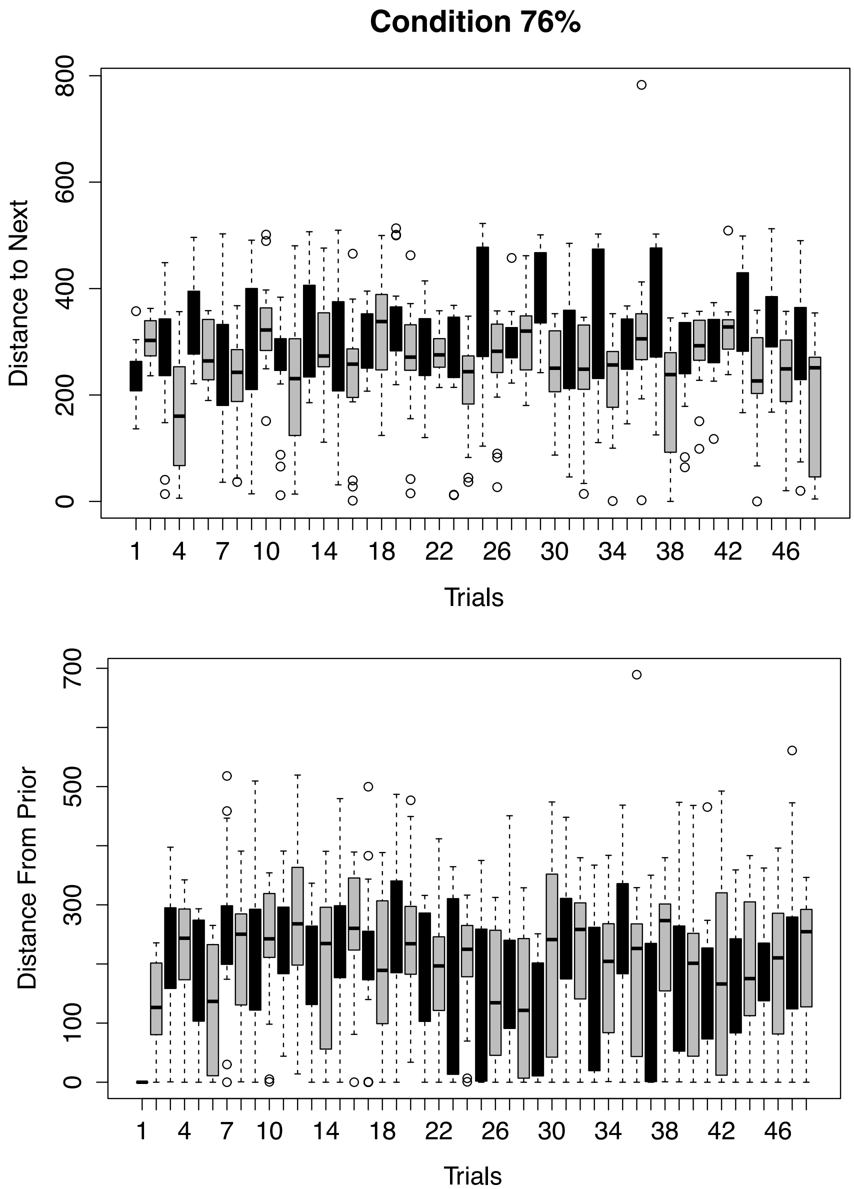

Figure 19. Average *distance to next (top) and distance from prior (bottom)* on a trial by trial basis in condition 76.

To further study this matter I explored the potential interaction between condition and trial. While trial (i.e. how far in the sequence they are) and condition (i.e. how structured the sequence is) are both significant predictors of *distance to next*, there is a strong interaction between the two of them (β=-4.75, p < .0001). This suggests that as the participants go further in the sequence, in more regular sequences their behavior becomes more and more predictive. The interaction shows that local effects have a big impact on the general pattern. Next chapter will discuss and examine local structures and their role in sequence learning.

3.7 Experiment 3: Local Structures

In Experiment 1, we found the participants' behavior to be influenced by local structures even in random sequences. Local structures are short subsets of the sequence that happen to be more patterned than the whole sequence. Some participants showed a change in strategy and performance when encountering structured subsets of the sequence even if the length of this subsequence was as short as 4 to 5 trials.

In this experiment the goal is to intentionally create a structured subsequence in the heart of an otherwise random sequence. I expect the participants to realize the predictability of the subsequence and switch to predictive mode, which will result in shorter reaction times and shorter *distance to next*. Another goal that this follow-up experiment suffices is exploring the onset of prediction. How long should the structured sequence be for the participants to trust the pattern and start predicting? How long does it take for people to switch from a reactive mode in a random setting to a predictive mode in a setting perceived as structured and predictable? This experiment is designed to explore the role of local structures and how or if they trigger the predictive behavior in participants.

3.7.1 Stimuli

In this experiment there will be only one condition for all participants. The sequence consists of two random segments and one very patterned and repetitive subsequent. The hope is that the participants stop the reactive mode and start exhibiting predictive behavior as they notice the presented stimuli to become repetitive and more patterned. The location of the stimuli in first 20 trials of the sequence will be chosen randomly, followed by eight trials that resonates between location 1 and 2, in order to impose a local structure to the sequence. After the patterned subsequence of 8 trials there will be another 20 random trials. We predict that the participants will start performing more predictive as they find out about the pattern in the stimuli. The main goal of this experiment is to explore the onset of prediction. How structured should the sequence of stimuli be for the participants to trigger predictive mode.

3.7.2 Participants

This experiment was uploaded on Amazon Mechanical Turk. Data from 75 participants were collected and analyzed.

3.7.3 Procedure

The procedure in this experiment is similar to Experiment 1. The only difference lies in the number of conditions. In this experiment there is only one condition (as described in the stimuli section). All participants receive the same sequence consisting of three blocks of random, structured, and again random order respectively.

3.7.4 Measures

The measures used for analyzing the mouse trajectories and behavior follows the format described in the main stimuli section (*distance to next* as a measure of predictive behavior). Another important measure used for analyzing the data form this experiment, is focused on participants' behavior throughout the local structured block. The reaction time difference between trial 20 (the start of the structured block) and trial 22 will be compared to the reaction time difference between trial 26 and 28 (indicating the end of structured block).

3.7.5 Results

The first step in conducting analysis for this experiment is to show if the structured block did in fact make participants engage in more predictive behavior. Consistent with the expectations, the average *response time* and *distance to next* both dropped during the structured block and raised again as the sequence went back to being random. It is apparent that trial type (random v.s. structured) is a very significant predictor of *distance to next* ($\beta = -33.63$, $p << .001$). On average, the structured block shows 33 points decrease in *distance to next* compared to random blocks. This is also the case with *response time* ($\beta = -33.08$, $p = 0.003$) as the subjects are significantly faster in the structured subsequence. Figure 20 (left) demonstrates the rise and fall of average *distance to next* over 48 trials. Fig 20 (right) zooms into the same graph in order to depict the slope of change throughout the 8 structured trials. As it is shown in the figure, there is a sharp decrease in the distance to next as the structure block unfolds. A linear mixed effects model was conducted to test the significance of this decrease. The model confirms that there is a highly significant change in measures in the structures subsequence. Trial is a strong predictor of *response time* (β = -17.59, $p < .0001$) and *distance to next* (β = -11.06, $p < .001$) throughout the local structure while it does not provide any information in the random parts of the sequence (β = 0.58, p = 0.4).

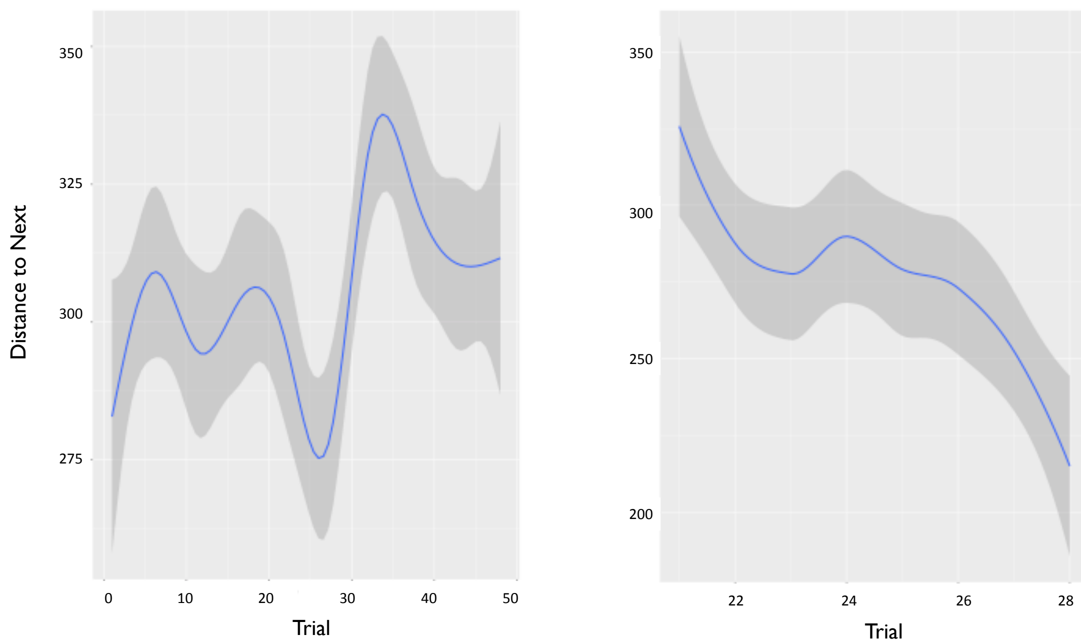

Figure 20. The figure on the left shows the average *distance to next* throughout all 48 trials. The figure on the right zooms into the graph through the structured block (trials 20-28).

More detailed analyses were conducted to compare the difference in participants' performance as the structured subsequence unfolds. As mentioned above, I calculated the difference in reaction time between the first two trials of the structured block as well as the difference between the last two trials of the structured block. The idea is that at the beginning of the structured block participants will still be in reactive mode under the impression that the sequence is still entirely random. However, by the time they get to the end of this structured block (trial 26-28) they have picked up on the pattern and start being more predictive which will result in shorter reaction times and shorter *distance to next*. It is important to note that in analysis I wont be comparing two absolute values, rather I will be comparing two differences. The bigger the difference the faster the participants were towards the end of the structured block. These faster reactions could be interpreted as more predictive strategies and better learning of the underlying structure of the stimuli.

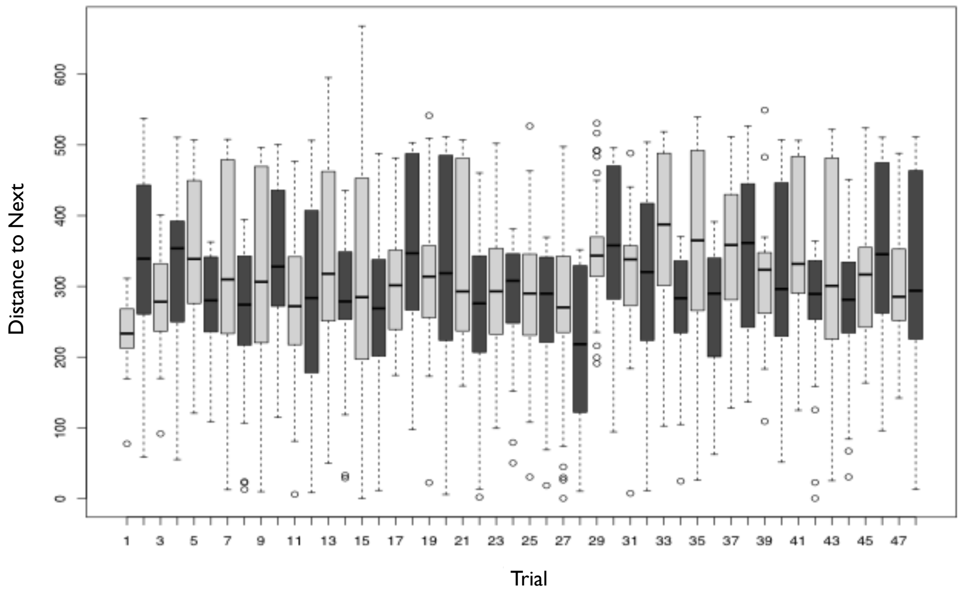

Figure 21. Trial by trial changes in average distance to next. It is apparent how the average distance to next decreases through trials 22-28.

A two sample t-test was conducted to compare the difference between trial 20 to 22 (M= 134.40, SD= 103.98) and the difference between trial 26 to 28 (M= 68.51 , SD=84.2). Results from the two tailed t-test show that there was a significant difference between the changes of *distance to next* from trial 20 to 22 compared to the change in trial 26 to 28 ($t_{(74)}$ = 3.1147, p = 0.002). The change in predictive behavior spikes as we get closer to the end of the structured block, suggesting that it only takes maximum of 6 trials for most of the participants to trust the structure of the environment and start predicting. They see patterns in less than 6 occurrences and act upon what they learn in such short exposure. Prediction appears to be a fundamental indictor of this learning process.

3.8 Experiment 4: Forced Centering

Data from experiment 1 and 2 show that some participants choose to go back to the center of the screen after each trial. They take their mouse cursor to the center to be equidistant from all four positions on the screen. This strategy, called "*centering",* makes it more challenging to compare all participants and difficult to distinguish between predictive and reactive behavior. In Experiment 4, I change the procedure so that all participants are forced to go back to the center after each trial. They will have to click on the center for the next stimuli to appear. This modification will normalize the measures from different strategies while participants can still exhibit predictive behavior. The main difference will be reflected into their performance after clicking on the center and during the 750ms before the next stimuli shows up. The goal is to replicate the results from experiment 1 with a more standard baseline for comparison between subjects, regardless of their strategy. If the effects hold under the new circumstances, we can have a more accurate picture of predictive behavior while at the same time replicating the results from the original experiment. Moreover, this additional step in the procedure refines the definition of prediction as it was always a debate whether going back to the center and waiting in an equidistance position from all four corners could count as a type of predictive behavior. On the one hand, participants who apply "centering" strategy are relatively more engaged in a predictive behavior compared to those who will let the mouse cursor sit where the previous stimulus was and passively wait for the next target to appear. On the other hand, it is also true that centering is not a fully predictive behavior as it still involves waiting for the target to show up before any directed move.

3.8.1 Stimuli

The stimuli in this experiment are exactly the same as Experiment 2. Each condition contains a sequence of 48 trials half of which are completely random and the other half tends to become more regular. The conditions are distinguished based on the level of regularity in the second half of the sequence. There will be seven conditions with the regularity levels 44%, 55%, 68%, 76%, 79%, 89%, and 100%.

3.8.2 Participants

This experiment was also uploaded on Amazon Mechanical Turk. 100 participants were recruited from the United Sates. All participants were offered a monetary compensation of $0.75for their time.

3.8.3 Procedure

Similar to Experiment 1, participants were instructed to click on the location of the next stimuli as fast as they can. However, the main difference between this experiment and Experiment 1 lies in the procedure. After reading the instructions they were directed to a page where they saw a red square placed at the center of the screen. A text box on the screen instructed participants to click on the red square in order for the next stimuli (black circle) to show up. 750ms after they click on the center square the first stimulus appears. Each participant was randomly assigned to one of the conditions described in Experiment 1.

3.8.4 Results

As the goal for this experiment is to replicate the results from Experiment 2 with more homogenous baseline for all participants, the analyses were chosen to be similar to the ones from Experiment 2. However, the analyses for this version of the experiment are more organized with an opportunity to compare all types of predictive/reactive strategies with a common baseline. Measures used for these analyses are also the same as the ones used in the previous experiment, namely *distance to next* and response time.

Out of 100 participants, recruited through Amazon Mechanical Turk, one participant failed to complete the task and was therefore excluded from the analysis (N=99). The bar-plot in Figure 22 shows that the pattern in the data stays very similar to Experiment 2. While the over-all pattern is the same, the difference between random and structured sections of the sequence tend to be more fluctuant (See Figure 23).

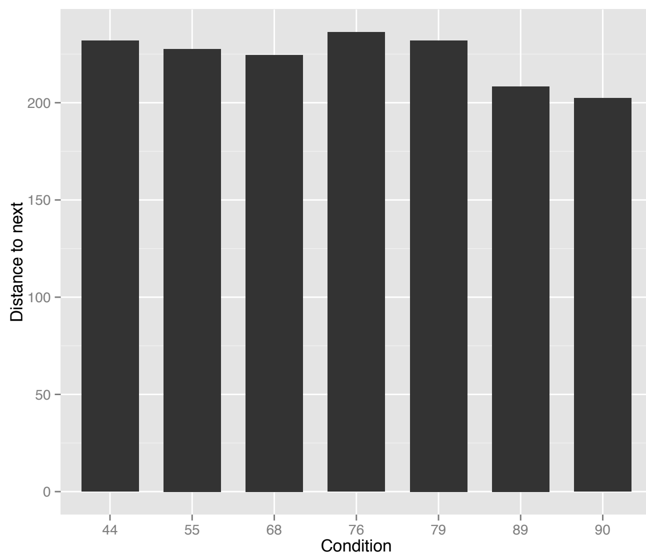

Figure 22. Average distance to next for all conditions in Experiment 4.

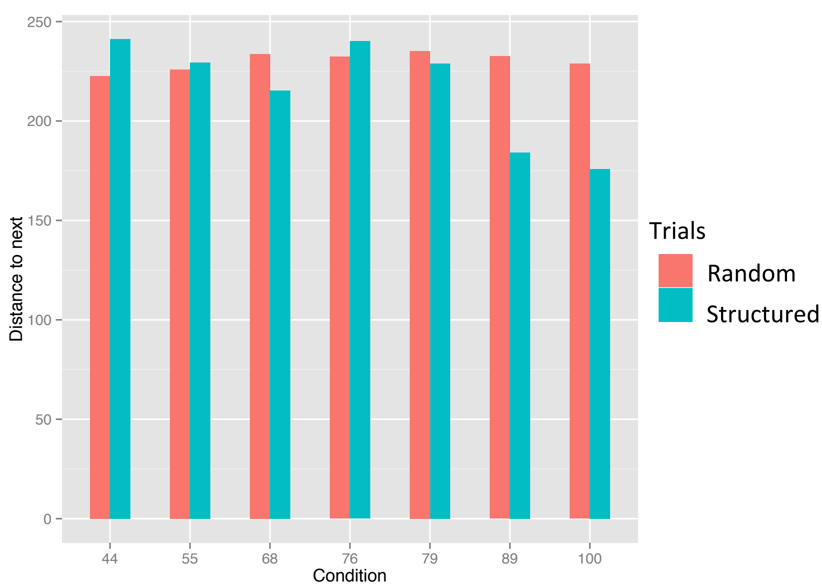

Figure 23. Comparison of average *distance to next* in random and structured parts of the sequence for each condition.

Consistent with the results in Experiment 2, condition 100 showed the shortest average *distance to next* (M_{dist_next} = 242.19) while condition 44 showed the longest average distance in pixels (M_{dist_next} = 309.5). This is when participants in every condition show high *distance to next* values when they are in the first half (i.e. the random half) of the sequence of target locations (See Table 10 for more details).

Table 10
Means and Standard Deviations of the Mouse Trajectory Variables by Condition in Experiment 3 (Forced Centering).

	Time and Trajectory Variables					
	Distance to Next (pxl)		Response Time (ms)		Distance from Prior (pxl)	
Condition	M	SD	M	SD	M	SD
44% (N= 14)	231.95	125.46	547.77	511.13	297.03	98.23
55% (N= 13)	227.66	104.67	522.83	225.21	288.93	97.06
68% (N= 13)	224.57	108.62	531.20	203.66	277.83	77.48
76% (N= 17)	236.18	106.77	562.44	199.17	282.2	91.20
79% (N= 23)	232.05	90.70	562.27	267.65	268.53	80.33
89% (N= 7)	208.39	119.90	599.85	274.75	283.09	86.17
100%(N= 12)	202.32	120.42	598.26	497.49	286.96	90.59

A linear mixed effects model with subject as a random factor, and condition (44,55,68,76,79,89,100) as numerical fixed factor was conducted to assess the effects of structure on *distance to next* throughout 48 trials. The model showed that as we expected condition (β = -.61, p = 0.02) is a reliable predictors of *distance to next*. This indicates that despite being forced to go back to the center after each trial, participants still engage in predictive behavior as the environment gets more structured. This predictive behavior becomes more significant as they get further in the experiment and pick up on the potential patterns in the sequence. Therefore, there is also a significant interaction between trial and condition (β = -0.04, $p << 0.001$).

That being said, forcing all participants to go back to the center appears to undermine the difference between predictive and reactive behavior. Even though there is still the same pattern of changes in the average *distance to next* as the sequences get more and more structured the results are not as significant. It appears that forcing "centering" strategy upon all participants makes the gap between predicting and reacting participants to shrink. There are at least two explanations for this observation. On the one hand, forced "centering" could demotivate predictive participants from actively predicting the location of next stimuli as it introduces the additional task of relocating the mouse cursor to the center. This additional movement could be an extra cognitive load and a barrier on the smooth process of transition into predictive mode. On the other hand, this slight decrease in the difference between performance could be caused by forcing reactive participants to apply a semi predictive strategy such as "centering." In other words, in the current experiment participants cannot passively wait for the next stimuli to show up. They have to take their mouse cursor to the center and click to proceed. This boosts their predictive measures such as *distance to next* and puts them in a better position compared to the reactive participants in Experiment 2. Figure 24 shows the average distance to next for all conditions in Experiment 2 and the current experiment on the same graph.

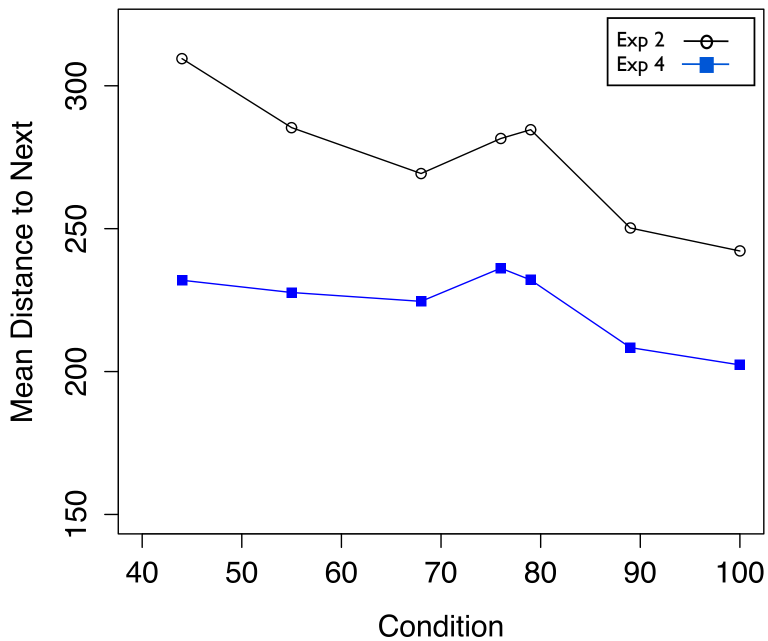

Figure 24. Average distance to next for all conditions in Experiment 2 and the current experiment.

An across experiment analysis was conducted to quantify the effect. The mixed effects model across Experiment 2 and the current experiment shows an extremely significant decrease in average *distance to next (ß= -47.9, p << .0001)*. The effect is very noticeable indicating that there is 47 units drop in average distance to next when participants are forced to go back to the center. Thus, the results confirm this boost in predictive measures in Experiment 4. The average distance to next tends to be lower for every single condition in Experiment 4. This suggests that forced "centering" did in fact induce some predictive behavior in participants causing the gap between the two types of participants to fade.

3.9 Experiment 5: Priming the Prediction

Experiments 2 and 4 (i.e. new sequence and forced centering) offer platforms to explore the predictive or reactive behavior in environments with different levels of structure in presence or absence of forced centering. However, in both those experiments participants were only instructed to react as quickly as they can. The word prediction was never explicitly used in the instructions. A simple yet crucial manipulation of the task could be a change in wording of the instruction. Participants will be instructed to predict the location of the next stimulus.

The objective of this modification is to prime the participants to adopt a predictive approach. This simple change in instructions will change the participants' mindset before they start the task. The hope is for them to actively predict and be aware of their errors. The action dynamics thus generate an error signal that improves learning. The results from this experiment could be compared to those of experiment 2 and 4.

3.9.1 Participants

The experiment was uploaded on Amazon Mechanical Turk and 108 participants were recruited. All participants were paid \$.75 for their time.

3.9.2 Stimuli and Procedure

The stimuli are identical to Experiment 2. The only difference in this experiment is driven by the instructions in procedure. Prior to the offset of the experiment, participants were explicitly instructed to predict the location of each target before it appears on the screen. Following is the exact wording of the instructions for this experiment:

> *"We want you to make PREDICTIONs! Can you predict the location of the next circle and click on it as soon as it appears? Click this box to begin, and first circle will appear"*

3.9.3 Results

Out of the original 108 participants, one person was excluded from the analysis due to their extremely long response times. Moreover, trials with response times longer than 20seconds were discarded from analysis as outliers.

A linear mixed effects model was conducted with condition and trial as fixed effects and subject as a random factor. The results show that the same pattern observed in the original experiment is present in this modification of the experiment as well. Trial and condition are still significant predictors of *distance to next* as the more regular sequences cause reliably shorter *distance to next* and therefore more predictive behavior.

The main goal of this experiment, however, is to assess whether participants tend to be actively more predictive because they have been primed to actively predict. Or rather, there is no certain change in their behavior. In order to show this effect, I ran analysis across two experiments. The measures from the current experiment were compared to the results from Experiment 2. Figure 25 depicts the average distance to next for all conditions in both experiments. The graph shows that participants in the priming condition are in fact more engaged in predictive behavior. *Distance to next* is significantly shorter () and response time also shows decrease in the current experiment compared to Experiment 2.

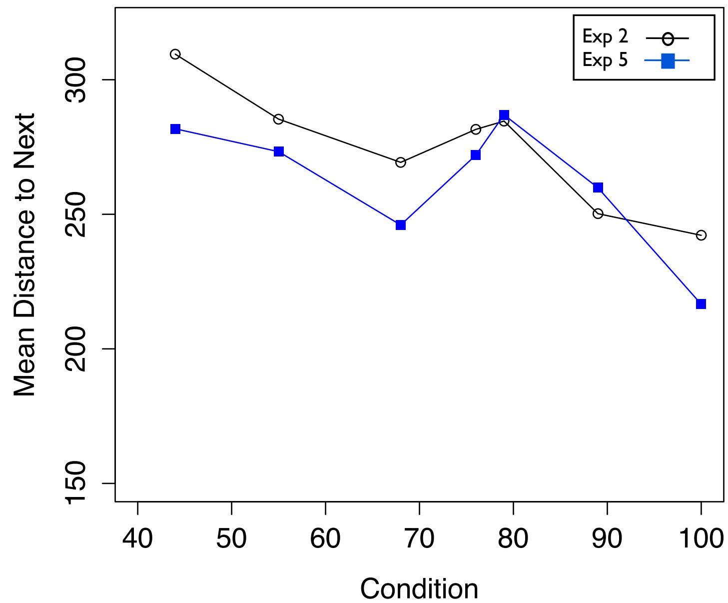

Figure 25. Average distance to next for all conditions in Experiment 2 and the current experiment

A linear mixed effects model was conducted across Experiment 5 and Experiment 2 to explore the effect of priming on predictive behavior. I looked at experiment type as a predictor of *distance to next*. The model showed a non-significant decrease in *distance to next* under priming instructions (5.9 units drop in average *distance to next,* p=.4). Further, I restricted the analysis to only include the regular half of the sequence, which also did not show any significance (ß = .5, p = .9). However, surprisingly, the more significant change happened in comparison between random halves of the sequence (ß = -12.9, p < .05). It appears that the effect is more salient when there is less regularity in the sequence. That is, participants will start being predictive in structured environments anyways, but being primed to make predictions enhances their predictive behavior in a random environment as well. To explore this further, a model was designed to study the interaction between condition and experiment type across the two experiments. The results reemphasize the very significant role of regularity condition (ß = -1.1, p << .01) as well as a significant interaction between condition and experiment (ß = 0.8, p < 0.1). This suggests that participants in the priming version showed significantly more predictive behavior in more regular conditions.

The results support the hypothesis that the onset of prediction could be moved to the very beginning of the sequence by priming the participants to actively predict. Participants transit into predictive mode more easily in this condition. However, the main question may be whether the decision to make predictions can in fact influence the learning process in sequential learning. In other words, could an explicit effort for making predictions influence the implicit learning of patterns in the sequence? A good follow up study to explore this issue would be presenting participants with a secondary test to see if they are able to pick the sequence resembling the target sequence with a probability higher than chance level.

3.10 . Discussion

A recent approach in cognitive science argues that prediction is a core concept underlying perception, action, and cognition, to the extent that brains could be referred to as "prediction machines" (Clark, 2013, p.1). This framework is a valuable reference for unifying views of perception, action, cognition and attention. Extending experimental paradigms to explore prediction in different aspects of human behavior could facilitate tests of this claim. To serve that end, the current document offered an experimental paradigm to study the onset of prediction in statistical learning.

The results in the literature have suggested that participants tend to rapidly switch into a predictive mode almost as a discrete strategy. In this project I designed tasks, inspired by previous work, to explicitly explore the onset of predictive behaviors. Through a mouse-tracking task, participants get to learn the statistical structure in a sequence of flashing dots while their mouse movements are being recorded. Pilot results suggested that participants engage in more predictive responses as the regularity of the statistical sequence increases. These preliminary findings revealed the significant effect of the statistical structure on triggering the rapid onset of predictive behavior. The follow up experiments lay out an incremental design for manipulating important aspects of the study to tap into the mental processes underlying predictive behavior.

Despite all the limitations in the experimental design and measures, what is novel in this research is how I provided a tool to capture the time course of the onset of prediction. These continuous measures also present an opportunity for observing the predictive and reactive strategies applied by participants throughout the experiment as well as the shifts in strategy as the regularity of the stimuli changes.

Chapter 4

4. Simulation

Time is an inevitable character of prediction. Predicting future events requires a chronological structure that underlies the sequence of events. Simulating a prediction task, therefor, would obviously impose a need for simulation of the temporal order of events. There have been various attempts to implement time in the literature (Jordan, 1986; Jordan & Rosenbaum, 1988; Fowler, 1980). Many of these approaches look at time as an explicit entity or separate dimension of the input, whereas more recent strategies have looked at time as an implicit feature which could be treated by its effect on the process (Jordan, 1986; Elman 1990). In connectionist models this implicit account of time could be represented as a memory of events. Whereas, probabilistic models capture the concept of time by updating the posterior probability using pervious events.

Here, I propose two main simulation approaches to model the predictive behavior observed in humans. One group of connectionist models which provide infrastructure for simulating memory are Simple Recurrent Neural Networks or SRNs. Section 1 of this chapter offers a description on how I simulated the process of sequence learning by the aid of recurrent neural networks. The other approach of modeling taken in this document is Bayesian modeling as it is well suited for simulating prediction based on the likelihood of events happening in the environment. Section 2 is an attempt to answer the question of how and to what extent Bayesian models can capture the predictive strategies involved in statistical learning.

4.1 Simple Recurrent Neural Network

Implicit account of time in connectionist models has been implemented via different strategies. Jordan (1986) used recurrent connections to implement a dynamic memory for the first time. In his model, a copy of the previous state of the system was fed to the input level in order to have an influence on the training process and the next round of output production. He implemented this by adding extra units, called state units, to the input level. These nodes would be the same number as output units and will copy the previous values of the output units to be involved in future calculation of the weights and output values. There will be direct connections between each node on the output level and each state node. Figure 26 (left) shows a rough sketch of Jordan (1986) recurrent network where not all connections are depicted. Inspired by Jordan (1986), Elman (1990) introduces a recurrent neural network with internal representation of time. In this network structure, the hidden unit patterns are fed back to the system. The same number of nodes as hidden units are added to the input level, called “context” nodes. These nodes receive connections from the hidden layer; the connections will always have a fixed weight of 1.0 and will copy the previous state of the hidden units into context units. This approach allows the system to hold a memory of the former state of the network and make predictions based on that memory. Panel B in Figure 26 shows a recurrent neural network structure proposed by Elman (1990).

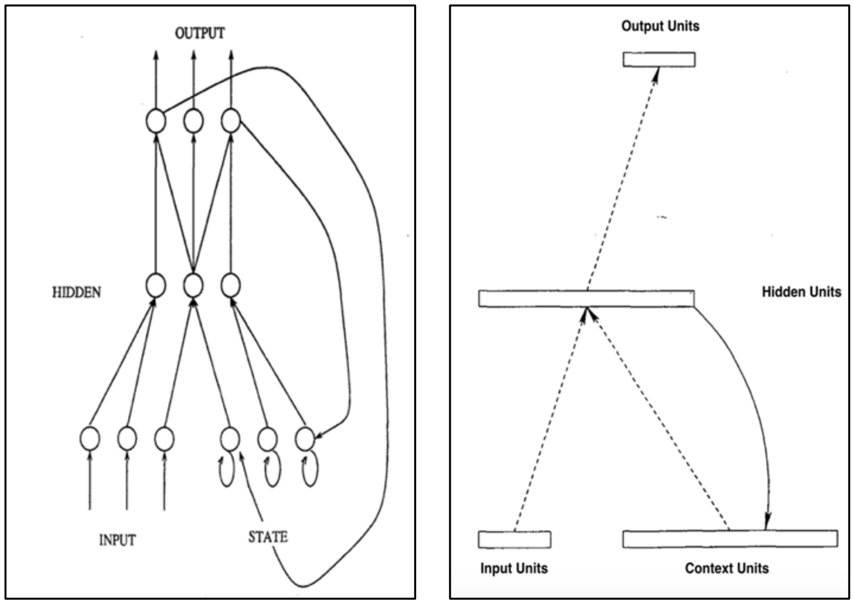

Figure 26. On the left is the structure of the recurrent network proposed in Jordan (1986). The image on the left shows the structure of the network designed by Elman (1990) where the hidden units have backward connections to the context units.

Elman showcases this recurrent network by solving the XOR problem. Using recurrent connections to produce memory offers an efficient solution for XOR problem. In this problem the network has to get the input [0 0, 0 1, 1 0, 1 1] and produce the output [0, 1, 1, 0] respectively. Elman concatenates each input string and the output for that to make a long sequence of numbers. This sequence of input will then be fed to the network and after training the network learns to predict every third element of the sequence as it is predicted by the previous two.

I propose a structure where a copy of the activation of hidden units is sent back to the input layer to be treated as memory of the previous step. Adding memory to the structure of the network helps us capture the concept of time as it is the fundamental parameter in a sequential learning task.

4.1.1 Implementation

In the current paper we have a prediction task similar to the XOR problem with more complexity and less regularity in the sequence. Here we will have sequences made of integers 1,2,3, and 4. Each sentence consist of 48 elements and the order of elements in the sequence is determined by sampling at different levels of regularity. A very predictable sequence would be of 100% regularity or G=1.0 (for more information about G please see Jamieson & Mewhort 2009). An example of a 100% predictable sequence would be:

[4 3 2 1 4 3 2 1 4 3 2 1 4 3 2 1 4 3 2 1 4 3 2 1 4 3 2 1 4 3 2 1 4 3 2 1 4 3 2 1 4 3 2 1 4 3 2 1 4 3 2 1]

Each number in the sequence was replaced with a unique four-digit binary code to be fed to 4 input units:

1 → 0 0 0 1 2→ 0 0 1 0 3 → 0 1 0 0 4 → 1 0 0 0

The network consists of 4 input units, 10 hidden units, 10 context units and 4 output units. Each input was a 4-bit binary vector fed to the input level. Input units and context units are all connected to the hidden layer. The hidden units feed forward to the output units while also having backward connections to the context units. A simplified structure of the network is demonstrated in Figure.27.

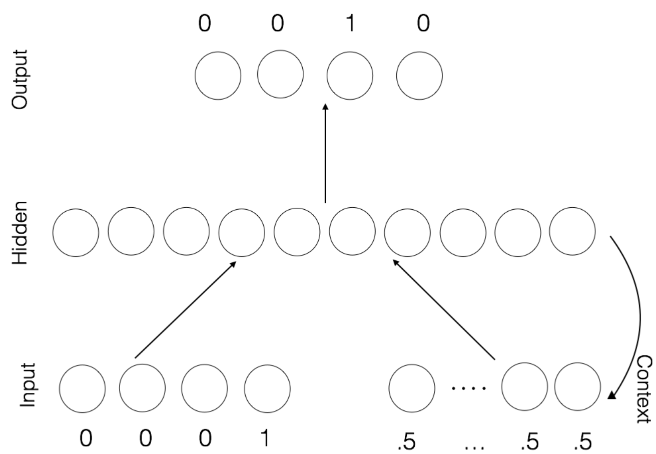

Figure 27. Structure of the recurrent network proposed in this paper

The network was trained by presenting 4-bit input vectors one at a time, when it was expected to predict the next input. At time step ***t*,** the first input vector of four elements is fed to the input units. The context units are all initially set to 0.5. Both input units and context units then activate the hidden layer. Hidden units will activate the output units with a feed forward connection. A copy of hidden units' activations will be fed to the context units through a backward connection from the hidden layer; the weights on this set of connections will always be set at 1.0 with no change during the training. At time step ***t+1*** the input units will receive the second input vector in the sequence while the context units will be fed the values from the hidden layer at time ***t***. This new set of input will activate the hidden layer and consequently the output layer.

The training consisted of simple back propagation approach (Rumelhart, Hinton, & Williams, 1986) where the output pattern is compared to the target output and the error is used to adjust weights on the input to hidden connections as well as hidden to output connections. The network was trained over 10000 iterations of a sequence with 24 elements (i.e. 24 vectors of 4-bit length) of a random sequence. The weights from that training stage was then used to initiate the training for the sequence with 24 elements of a more regulated sequence. This training was also iterated 10000 times. Testing was done on a sequence of the same length with a different order but the same level of predictability.

The network was separately trained and tested on each condition. Figures 28 through 10 illustrate the mean-squared-error plots for all three conditions. Note that each plot is made by concatenating two mean-squared-error plots: the training on 24-element random sequence followed by the training on 24-element regulated sequences.

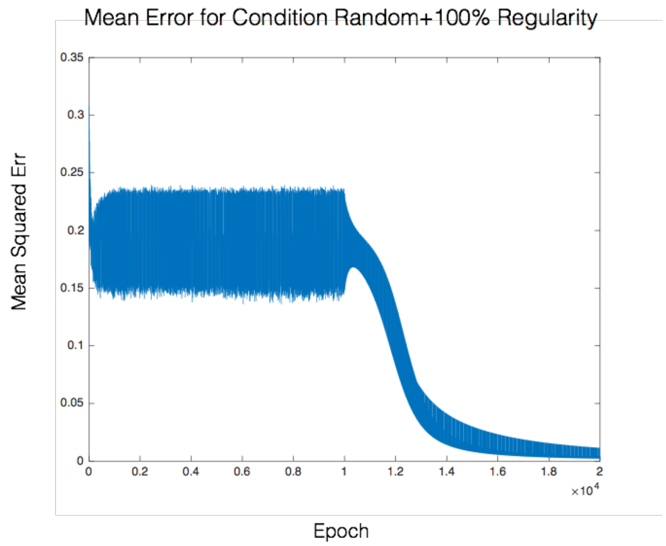

Figure 28. Mean-squared-error plot for training of the condition with random sequences followed by 100% regularity sequences

As shown in Figure 28, no learning happens when the network is being trained on fully random sequence of inputs (the first half if the error plot). Network behavior stays random as is the input. However, in the second half of the plot when the 100% regular sequence of input is presented to the network it starts to predict the next input pretty quickly. The mean-squared-error plot depicts rapid decrease in prediction error which results in accurate predictions towards the end of 10000 iterations of the network.

Figure 29 (left) shows the error plot for condition 89. The first half of this plot is devoted to random sequences which does not lead into any learning. However, the network behavior changes in the second half of the error plot. The network starts predicting with a hope of learning.

A closer look into the second half of the graph (Figure 29 right) shows a clear zoomed-in view of the prediction-errors made by the network. As it appears in the zoomed-in graph on the right, the network is learning the predictable subsets of the sequence (89% of it) and based on the pattern it has picked up confidently makes incorrect predictions about the irregular parts of the sequence. Therefor, it is evident that although some learning is happening the network is incapable of predicting the unpredictable anomalies.

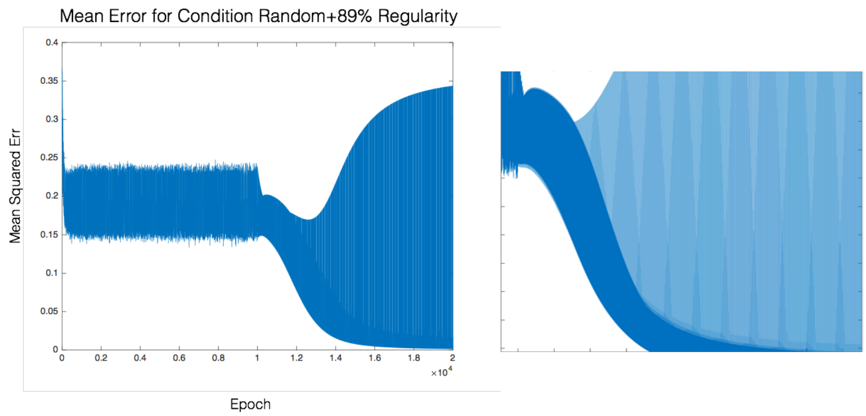

Figure 29. Mean-squared-error plot for training of the condition with random sequences followed by 89% regularity sequences.

We expect to see the same pattern of errors to repeat in the 55% condition. In fact, the mean-squared-error plot for 55% condition shows the same pattern with even more severity in making confidently wrong predictions. A zoom-in capture of the second half of the plot (Figure 30 left) gives us a better view on how the network makes incorrect predictions 45% of the time, while learning the correct predictions rest of the time.

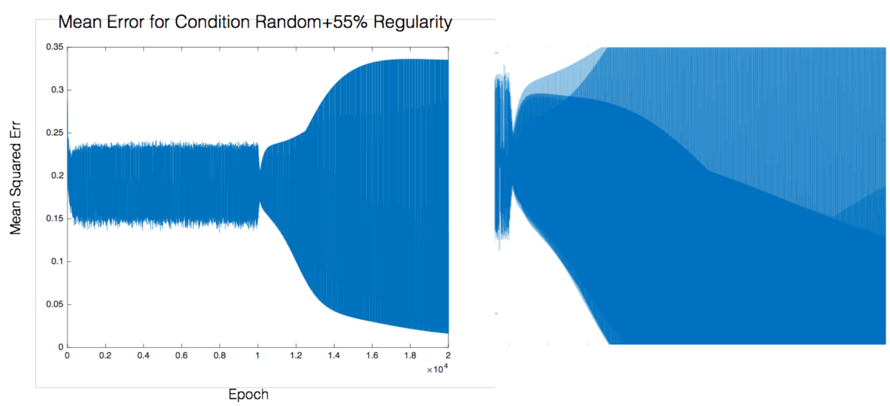

Figure 30. Mean-squared-error plot for training of the condition with random sequences followed by 55% regularity sequences.

4.1.2 Discussion

Here we simulated the process of sequence learning by the aid of recurrent neural networks. In general, these networks hold a memory of the previous activation of the output. We proposed a structure where a copy of the activation of hidden units is sent back to the input layer to be treated as memory of the previous step. Adding memory to the structure of the network helps us capture the concept of time as it is the fundamental parameter in a sequential learning task.

The goal for this network was to be exposed to the same sequence of stimuli as human subjects were in the experimental design presented in the previous section. In the case of 100% regularity sequences, the network quickly learns to predict the next element of the sequence with very low error. In a less regular sequence, such as 55% condition, the network shows some learning compared to pure randomness but is still far from ideal prediction of the whole sequence. Considering the nature of the sequence this tends to be a perfectly expected behavior.

It appears that the network succeeds in mimicking human performance in such tasks. During the random trials the network is having a hard time predicting future locations of the target. As soon as some regularity is offered the behavior of the network changes and it starts making confident predictions, whether correct or wrong. Even though Recurrent Neural Network model presented here fits the human behavior data very well, this model does not shed any light on how humans make such decisions and apply certain strategies. All this model does is to mimic the iterative learning process and match the end results with those of humans. Using models that can provide a better insight into how humans do so could be a valuable contribution to predictive learning. Thus, the next two sections are dedicated to discussing more transparent models of this phenomenon.

4.2 Entropy Model of Memory

In the previous section, I utilized Recurrent Neural Networks to account for time as a key component of sequence learning. In this section, I take another route to simulate the effect of time, or more specifically, memory in statistical learning. In a more human-like approach, here I simulated a sliding time window over each given sequence to act as an active memory of the previous steps in the sequence. Since different conditions hold different levels of structure in the stimuli sequence, I was able to calculate an entropy score (See Shannon, 1948) within the scope of the sliding window for each of the sequences. This was done by designing a contingency table containing the probability of moving from each position to any of the other possible positions. For instance, in a 100% structured sequence with a repeating pattern of (1,2,3,4) the chance of going from position 1 to 2 is 100% whereas the probability of going from position 1 to 3 is absolutely zero. This gives a measure of entropy in the sequence indicating how much new information is provided by revealing the next position in the sequence, or in other words, how surprising the sequence is.

One core question in the current study is how much the entropy and structure in the sequence correlate with implicit learning, or rather, predictive learning. Here, by simulating the role of memory in this statistical learning process I made an attempt to visualize the relation between memory of the sequence and committing predictive behavior.

In each iteration of the simulation, the length of the sliding window was manipulated, ranging from 4 trials ago to 23 trials ago. The goal was to simulate the effect of memory capacity in remembering certain number of steps before the current step. For each memory size the entropy was calculated at each slide of the window across all conditions of the experiment (44%, 55%, 68%, 76%, 79%, 89%, 100%). The calculated entropies were then compared to the *distance to next* results for each condition across all participants looking for possible correlations. The results show an interesting correlation between the length of the sliding window (indicating the memory capacity) and *distance to next* (indicating predictive behavior). The curve in figure 31 demonstrates how the correlation between entropy and predictive behavior changes as a function of memory capacity. The correlation is low when the length of the sliding window is short (4 to 8 steps back). However, the correlation tends to peak around middle length window of memory (10 to 14 steps back) and decreases or flattens out in longer memory windows. These results suggest that memory of the previous steps do play an important role in implicit learning and predictive behavior. Figure 31 shows the changes across all conditions while Figure 32 demonstrates the same phenomenon for each level of structure to capture the correlation pattern of each regularity condition.

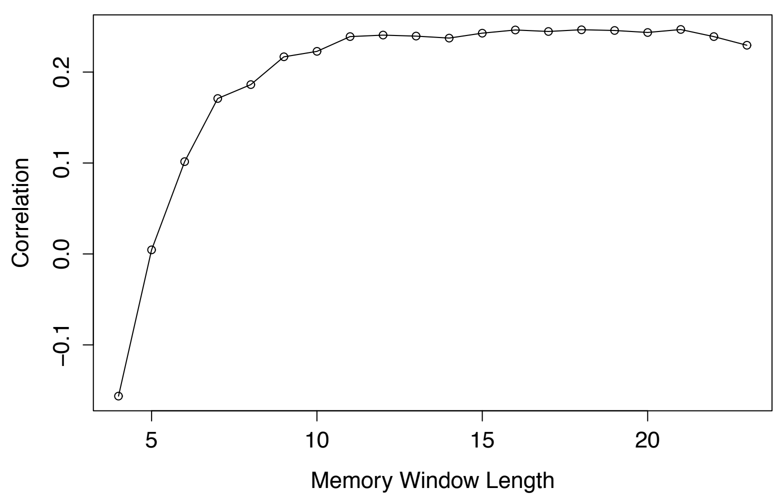

Figure 31. The correlation between entropy and predictive behavior as a function of memory capacity across all conditions

The model shows that having a memory of 12 +or- 2 items ago gives people enough confidence to engage in predictive behavior by making assumptions about the underlying structure of the environment. However, very short or very long memory spans do not necessarily show any facilitating effects on predictive behavior or implicit learning.

The number 12+-2 may remind the reader of the magic number 7+ or- 2 and the famous Miller paper on working memory capacity (Miller, 1956). The number depicted by the graphs in the current study tend to be bigger than 7+or-2, as this is not about active working memory capacity. It is possible that the system involved in keeping track of an implicit pattern in previous steps is not solely dependent on working memory. The distinction needs to be made between statistical learning and cognitive tasks requiring pure involvement of working memory.

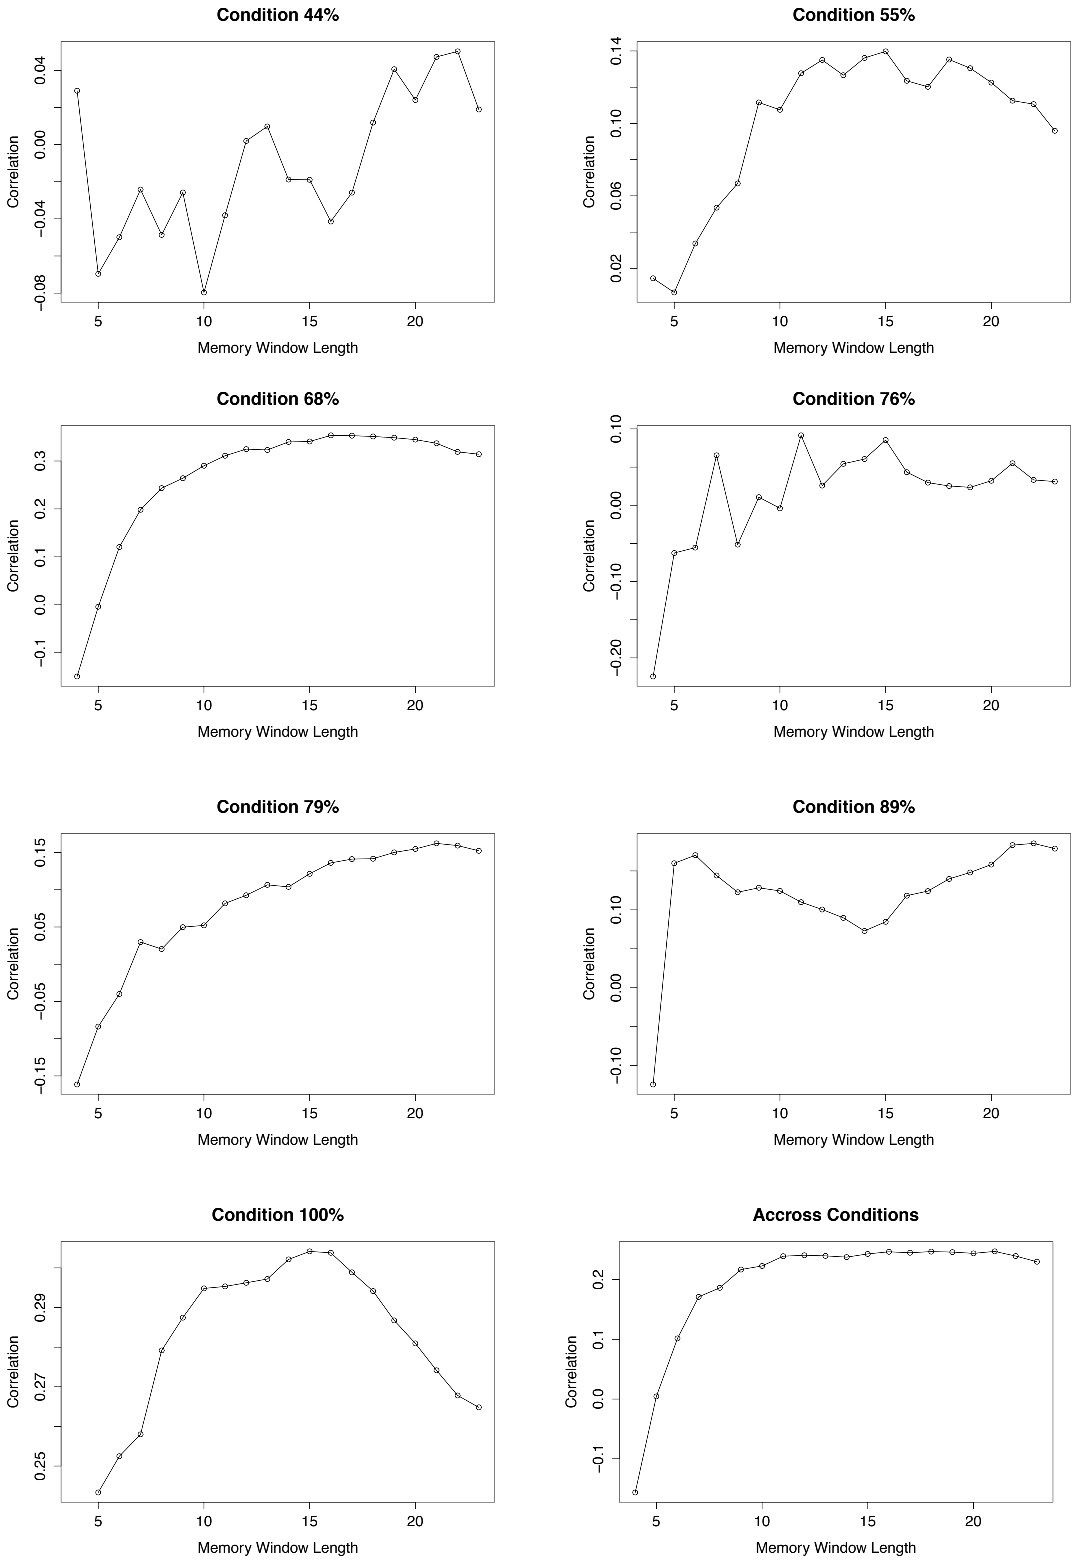

Figure 32. The correlation between entropy and predictive behavior as a function of memory capacity for each condition and across all conditions.

4.3 Bayesian Models and Future Direction

The use of Bayesian methods in cognitive science has considerably increased in the recent years. The current paper does not attempt to provide a detailed implementation of a Bayesian model for statistical learning. Nevertheless, it is only fair to mention the potential Bayesian models proper for this paradigm. In what follows, I provide a theoretical background on various Bayesian methods appropriate for modeling this phenomenon. The goal here is to show whether the problem of sequential learning, as described in this paper, is amenable to Bayesian modeling, and if so, which Bayesian approach provides an acceptable simulation of human behavior in this task. The actual implementation of these frameworks, however, would be the topic of a future work.

Bayesian models provide a method for how rational agents should infer from the environment to update their beliefs. Agents use a combination of new data, assumptions about the environment (likelihood), and general knowledge they have previously acquired (prior), to make an inference about the current state of the environment (posterior). Probabilistic models have been applied to different areas of cognitive research such as language processing, vision, deductive inference and decision-making. Another important area that many of these models are trying to address includes learning causal structure from patterns of statistical correlation (Griffiths, Kemp, Tenenbaum, 2008). The probabilistic nature of Bayesian models and their approach in making predictions based on a combination of new data and prior knowledge marks them as appropriate candidates for modeling predictive learning. Despite connectionist models that fit the data with a "black box" approach, Bayesian methods offer a more transparent view of how humans decide to predict in a semi-structured environment. The variety of Bayesian approaches in the field raises the question of which approach provides an acceptable simulation of human behavior in statistical learning.

One domain that has close resemblance to statistical learning might be the field of artificial grammar learning. Thus, Bayesian models offered in this field could very well be suitable for our purpose. Perfors, Tenenbaum, and Wonnacott (2010) present a domain-general model for learning artificial language using Hierarchical Bayesian models(HBM). They built up on a domain-general Bayesian framework for learning higher-level knowledge in an inductive form. Their hierarchical Bayesian model replicates experiment results by making inferences about the distributional statistics of verbs. Perfor et al. suggest that the model could capture other aspects of general language development while demonstrating that generalization in verb acquisition can emerge from more generic assumptions about the nature of learning.

Henderson and Tivot (2010) use a natural language parsing task to benchmark a class of Bayesian networks they have proposed for modeling structured prediction problems. The authors suggest that Incremental Sigmoid Belief Networks (ISBNs) proposed in their 2010 paper are especially appropriate for addressing complex structured prediction problems such as natural language processing. These models tend to make parsing the language possible due to their incrementally specified model structure. Thus ISBNs could be another potential rout for modeling our predictive learning task as it offers a memory of the previous events for making inferences about the upcoming stimulus.

The most relevant class of models proposed in the field are Dynamic Belief Models. Yu and Cohen (2009) use DBMs to model reaction time data in studies where humans observe non-existent patterns in random statistical sequences. In these types of experiments subjects see patterns of repetitions or alternations in a random sequence of stimuli and their response time and accuracy improve if the upcoming stimulus happens to follow the pattern they have imagined. This is very similar to the effect of local structures that we observed in our primary analysis on the pilot data from Experiment 1. The authors suggest that humans are used to find trends of statistical patterns in the stimuli as it is very rare to have purely random occasions in the real world environments where laws of physics and biology apply. They have modeled predicting stimuli based on previous observations using Bayes theory, and compared the models to participants' behavior. Dynamical Belief models seem to be the most appropriate approach to take for our experiment, as we are also interested in modeling human behavior while making predictions based on previously observed data.

Chapter 5

5. Discussion

The very promise of this dissertation was to offer a set of tools that goes beyond the here and now when studying the complexity of human cognition. The goal was to overcome the shortcomings introduced by traditional discrete measures of cognitive behavior such as Reaction Time. Is it possible to provide behavioral methods that go beyond the final outcome of a behavior and offer a peek into the process that gave rise to it? If so, could these methods improve what we know about complex cognitive processes? Could they shed light on aspects of emerging mind that stay out of reach with traditional discrete methods? And finally, could these new methods equip us with a better understanding of cognitive processes as complex and high-level as learning and decision making? The work offered in this dissertation shows that the answer to all of the above questions is a promising affirmative.

In the recent years, it is becoming clear that a rich and semi-continuous set of measures is extractable from relatively implicit behaviors such as hand or eye movements. It has been shown that spatial and temporal dynamics of motor movements (i.e. *action dynamics*) can shed light on the progression of high-level cognitive tasks. This rich body of data could provide a nearly real-time translation of unfolding cognitive processes.

In the current work, I chose two areas in which the action dynamics have not been explored nearly enough: deception and predictive learning. On the one hand this work takes advantage of action dynamics to get under the hood as these complex cognitive processes unfold in the brain. On the other hand, it utilizes these cognitive phenomena to demonstrate the power of action dynamics in studying higher level complex cognitive processes.

In this document, mouse tracking was utilized as a paradigm to conduct a series of studies in deception and predictive learning. This paradigm has only recently been evolved as means of studying action dynamics of cognition. The current work expands the use of this paradigm into studying deception in a naturally tempting setting as well as predictive learning in statistical structures. In what follows, I briefly touch upon the findings and limitations of each of the two studies. I will also revisit the findings of chapter 4, the computational modeling of predictive learning.

5.1 Action dynamics of deception

The control theory in deception states that humans have to actively suppress their truth bias in order to lie. While the control theory of deception seems to be elegantly supported by neuroimaging and experimental data, more recent threads of research provide evidence that challenges the confidence of this theory. The new evidence suggests that deception is not a secondary process resulting from active inhibition of the truth. Rather, it could be the default process that "automatically" happens in the presence of temptation in self-serving situations. Defying the notion of deception as a controlled inhibition of the truth, this view has also been supported by behavioral (Shalvi et al., 2012; Gino et al., 2011) and neuroimaging (Greene & Paxton, 2009) evidence.

In the present work, I utilized action dynamics and more specifically mouse tracking to contribute to this debate. The results showed that deceptive behavior could be facilitated when there is a self-benefit in being deceptive. People show less complexity in their mouse trajectories when they are being dishonest in incentivized tasks. "Dishonest" participants showed shorter and relatively direct mouse trajectories with no significant signs of hesitation. While being deceptive, their trajectories were more direct with less deviation towards the alternative response, illustrating less effort and time. "Honest" participants on the other hand, were demonstrating hesitation with longer trajectories and more attempts to change the direction of the mouse cursor.

There are many complications and limitations in studying deception. No matter how precise of a picture the action dynamics depict, they still have to deal with these limitations. One very common obstacle is that there is no unified concept of deception or a specific type of behavior referred to as deception. Rather, deceptive behavior contains a wide range of actions that could be categorized as deceptive while also falling under more general types of cognitive behavior. This speaks into how researchers may be tapping into conceptually different processes involving various brain networks while presumably studying deception.

Another difficulty with deception research is the social and moral aspect, which prevents participants from exhibiting natural behavior. The shame and fear of being caught influences the participants' behavior and adds noise to the real picture. One advantage of the study presented here is that it offers an anonymous setting, reassuring participants that there is no chance of being caught. On the other hand, most experimental designs in the field are inspired by deception detection frameworks as that was the historical goal of the field. These frameworks were later adopted to also study the neural/mental processes involved in deception, but are still not entirely suitable for such goal. Participants are usually instructed to be dishonest or honest which makes many of the studies in the field to be far from a real-world implementation of deception and temptation. In the current work, using mouse tracking gave me an opportunity to implement a genuinely tempting situation where the individual participates in the task with no explicit instruction on deception. Rather, they genuinely decide whether to act honestly or dishonestly from the comfort of their own home while not being worried about getting caught or socially shamed. All in all, mouse tracking provides a framework to study deception in more ecologically valid setting while shedding light on the underlying processes involved in making a decision to lie.

5.2 Action Dynamics of Prediction

Many researchers in different domains of cognitive science have identified prediction as central to perception, cognition, and action. Yet the methods used in studying these fields do not necessarily capture the forward thinking and predictive behavior involved. Phenomena such as the onset of prediction and strategy switch are not easily (if at all) captured by discrete and ambiguous measures like Reaction Time.

The paradigm offered in this paper makes it possible to study the unfolding of prediction as the core of learning and decision making. It allows for testing various hypotheses. For one, I focused on studying the level of structure in the environment and how that correlates with the onset of prediction. It appears that predictive behavior is a wager that should only be worth it for the individual to get engaged in. The findings show that not every participant switches to a predictive mode. Among those who do switch, the majority only do so if the underlying structure of the environment looks promising. Individuals stay reactive and only in anticipation when they are confused about the pattern or structure of the sequence. The results however, show that observing a local pattern of structure with the length as short as 4-6 trials is enough to trigger the wagers. Moreover, explicit priming of the participants results in noticeably more predictions. When people think they are expected to be predictive they take bets even on purely random sequences. They engage in predictive behavior even though it may not help with accuracy.

There is an emphasis on the role of prediction as the basis for learning in this task. That brings us back to the debate in the field over retrospective or memory based learning versus predictive or implicit learning. The findings from the studies conducted here, tend to side with predictive learning more so than the alternative. Participants may stay passive and not actively predict in all settings, but it does not take too much for them to exhibit their internal expectation of the sequence. It appears that there is no such thing as pure reactive approach. Whether they act upon it or not, participants still make a model or expectation of the next step and judge the structure based on that. The entropy model in chapter 4 shows how the memory of the previous steps could play a role in how eager the person is to take a predictive approach. Having a memory of approximately last 10 steps could provide the participants with confidence in their expected model of the sequence. At that point, whether they decide to be predictive or stay reactive, is still a notion of forward thinking on their end. Hence, it is worth noting that the findings suggest that both memory and prediction are in the heart of this learning process no matter we settle on explicit or implicit learning.

The experimental design here, was considerably simpler than the actual serial reaction times studies with hundred of blocks of different sequences. This was partially intentional as I am interested in scenarios more resembling the real world situations. Nonetheless, it would be very beneficial to develop this framework for tasks precisely comparable to the classic studies in the field of statistical learning. Another future approach that certainly needs to be explored is the level of learning and awareness in the task. The work here was focused on prediction more than the learning itself. However, the common end goal of the sequential learning studies is gauging the participants' implicit / explicit learning of the underlying structure in the sequence. Admittedly, the current work does not put enough emphasis on this aspect of the phenomenon. As demonstrated by Duran and Dale (2012) a proper follow-up study would be to add a recognition and or awareness phase to the end of the task. This would give us measures to assess the correlation between predictive behavior and learning.

5.3 Simulation

In the final section of this document I presented computational models in an attempt to simulate the results from experiment 1&2 in chapter 3, where participants tend to show predictive behavior as the sequences become more and more structured. While the connectionist model offers an exceptionally good fit for the results, the hope is that the Entropy Correlation model provides a more transparent picture of how participants produce those results. The main challenge in modeling this type of behavior would be to simulate strategy switches in participants' behavior. That will impose limitations and complications in the implementation of more transparent models.

5.4 Conclusion

Through this journey I made an attempt to show how it is only fair to start decoupling from traditional discrete measures of cognitive behaviors. Despite being reliable and informative tools, these measure are relatively blind to the unfolding of a cognitive process. I introduced action dynamics as new and more powerful alternatives. They offer what traditional measures did and much more. Using two series of studies, I demonstrated how these tools are empowering cognitive scientist by equipping them with simple, reliable and effective tools to get a real-time view of the mental processes in progress. I chose two scenarios that specifically lend themselves to an action dynamics study, as they are resulted from inherently dynamic processes orchestrated in the brain. What shown here, however, is generalizable to many other areas in cognitive science. In the recent years, researchers have shown that this thread has only begun to offer its power in capturing a more accurate picture of cognition.

Action dynamics approach, similar to any other, has its own limitations. Although it offers considerably more information compared to traditional measures in the field, there is still ambiguity in what give rise to certain behaviors. Motivations, brain sections involved in performing that cognitive behavior, and causal relations behind it would still be out of reach in these approaches. Mouse-tracking in particular shows mainly manual tendencies, whereas, combining it with eye-tracking and oculomotor measures could extend the range of reactive and predictive strategies captured. Coupling action dynamics with neuro-imaging and non-invasive brain manipulation studies that could establish causal relations could be a valuable tool in extending our knowledge of complex cognitive processes.

6. References

Abe N, Greene JD (2014). Response to anticipated reward in the nucleus accumbens predicts behavior in an independent test of honesty. J Neurosci 34:10564–10572

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

Barnes CM, Schaubroeck J, Huth M, Ghumman S (2011). Lack of sleep and unethical conduct. Organ Behav Hum Decis Process 115:169–180

Bond, G. D. (2008). Deception detection expertise. *Law and Human Behavior*, *32*(4), 339-351.

Chomsky, N., & Miller, G. A. (1958). Finite state languages. *Information and control*, *1*(2), 91-112.

Clark, A. (2013). Whatever next? Predictive brains, situated agents, and the future of cognitive science. *Behavioral and Brain Sciences*, *36*(03), 181-204.

Cleeremans, A., & McClelland, J. L. (1991). Learning the structure of event sequences. *Journal of Experimental Psychology: General*, *120*(3), 235.

Coco, M. I., & Duran, N. D. (2016). When expectancies collide: Action dynamics reveal the interaction between stimulus plausibility and congruency. *Psychonomic bulletin & review*, *23*(6), 1920-1931.

Dale, R., Duran, N. D., & Morehead, J. R. (2012). Prediction during statistical learning, and implications for the implicit/explicit divide. *Advances in Cognitive Psychology*, *8*(2), 196.

Dale R, Duran ND (2013) Dealing with complexity differently: from interaction-dominant dynamics to theoretical plurality. Ecol Psychol 25(3):248–255.

DePaulo BM, Kashy DA (1998) Everyday lies in close and casual relationships. J Pers Soc Psychol 74(1):63.

DePaulo BM, Kashy DA, Kirkendol SE, Wyer MM, Epstein JA (1996) Lying in everyday life. J Pers Soc Psychol 70(5):979.

Duran, N. D., & Dale, R. (2009). Predictive Arm Placement in the Statistical Learning of Position Sequences. In *Proc. of the 31st annual conference of the cognitive science society*.

Duran ND, Dale R (2012). Increased vigilance in monitoring others' mental states during deception. In: Miyake N, Peebles D, Cooper RP (eds) Proceedings of the 34th annual conference of the cognitive science society. Cognitive Science Society, Austin, pp 1518–1523

Duran ND, Dale R, McNamara DS (2010) The action dynamics of overcoming the truth. Psychon Bull Rev 17:486–491

Elman, J. L. (1990). Finding structure in time. *Cognitive science*, *14*(2), 179-211.

Fischbacher U, Fo¨llmi-Heusi F (2013) Lies in disguise—an experimental study on cheating. J Eur Econ Assoc 11(3):525–547

Fowler, C. (1980). Coarticulation and theories of extrinsic timing control. *Journal of Phonetics,*
8, 113-133.

Freeman JB, Dale R, Farmer TA (2011) Hand in motion reveals mind in motion. Front Psychol 2:59

Freeman, J. B., & Ambady, N. (2011). When two become one: Temporally dynamic integration of the face and voice. *Journal of Experimental Social Psychology*, *47*(1), 259-263.

Freeman, J. B., & Dale, R. (2013). Assessing bimodality to detect the presence of a dual cognitive process. *Behavior research methods*, *45*(1), 83-97.

Freeman, J. B., Penner, A. M., Saperstein, A., Scheutz, M., & Ambady, N. (2011). Looking the part: Social status cues shape race perception. *PloS one*, *6*(9), e25107.

Freeman, J. B., Pauker, K., & Sanchez, D. T. (2016). A perceptual pathway to bias: Interracial exposure reduces abrupt shifts in real-time race perception that predict mixed-race bias. *Psychological Science*, *27*(4), 502-517.

Freeman, J. B., Pauker, K., Apfelbaum, E. P., & Ambady, N. (2010). Continuous dynamics in the real-time perception of race. *Journal of Experimental Social Psychology*, *46*(1), 179-185.

Freeman, J. B., & Ambady, N. (2009). Motions of the hand expose the partial and parallel activation of stereotypes. *Psychological science*, *20*(10), 1183-1188.

Freeman, J. B., & Ambady, N. (2010). MouseTracker: Software for studying real-time mental processing using a computer mouse-tracking method. *Behavior Research Methods*, *42*(1), 226-241.

Gallivan, J. P., & Chapman, C. S. (2014). Three-dimensional reach trajectories as a probe of real-time decision-making between multiple competing targets. *Frontiers in neuroscience*, *8*.

Ghez, C., Gordon, J., Ghilardi, M. F., Christakos, C. N., & Cooper, S. E. (1990, January). Roles of proprioceptive input in the programming of arm trajectories. In *Cold spring harbor symposia on quantitative biology* (Vol. 55, pp. 837-847). Cold Spring Harbor Laboratory Press.

Ghez, A. M., McCarthy, D. W., Patience, J. L., & Beck, T. L. (1997). The multiplicity of pre-main-sequence stars in southern star-forming regions. *The Astrophysical Journal*, *481*(1), 378.

Gilbert DT (1991) How mental systems believe. Am Psychol 46:107–119

Gino F, Schweitzer ME, Mead NL, Ariely D (2011) Unable to resist temptation: how self-control depletion promotes unethical behavior. Organ Behav Hum Decis Process 115:191–203

Gold, J. I., & Shadlen, M. N. (2007). The neural basis of decision making. *Annu. Rev. Neurosci.*, *30*, 535-574.

Greene JD, Paxton JM (2009) Patterns of neural activity associated with honest and dishonest moral decisions. Proc Natl Acad Sci USA 106:12506–12511

Griffiths, T. L., Kemp, C., & Tenenbaum, J. B. (2008). Bayesian models of cognition.

Gunia BC, Wang L, Huang L, Wang J, Murnighan JK (2012) Contemplation and conversation: subtle Influences on moral decision making. Acad Manag J 55:13–33

Hawkins, J. & Blakeslee, S. (2004) On intelligence. Owl Books/Times Books.

Henderson, J., & Titov, I. (2010). Incremental sigmoid belief networks for grammar learning. *Journal of Machine Learning Research*, *11*(Dec), 3541-3570.

Hehman, E., Stolier, R. M., & Freeman, J. B. (2015). Advanced mouse-tracking analytic techniques for enhancing psychological science. *Group Processes & Intergroup Relations*, *18*(3), 384-401.

Hibbeln, M., Jenkins, J. L., Schneider, C., Valacich, J. S., & Weinmann, M. (2017). How is your user feeling? Inferring emotion through human–computer interaction devices. *MIS Quarterly*, *41*(1).

Horwitz, R., Kreuter, F., & Conrad, F. (2017). Using mouse movements to predict Web survey response difficulty. *Social Science Computer Review*, *35*(3), 388-405.

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

Ipeirotis PG (2010) Demographics of Mechanical Turk (Tech. Rep. No. CeDER-10 01).

Jackson, G. M., & Jackson, S. R. (1995). Do measures of explicit learning actually measure what is being learnt in the serial reaction time task? A critique of current methods. *Psyche*, *2*, 20.

Jamieson, R. K., & Mewhort, D. J. K. (2005). The influence of grammatical, local, and organizational redundancy on implicit learning: An analysis using information theory. Journal of Experimental Psychology: Learning, Memory, and Cognition, 31, 9 – 23.

Jamieson, R. K., & Mewhort, D. J. K. (2009). Applying an exemplar model to the serial reaction-time task: Anticipating from experience. *The Quarterly Journal of Experimental Psychology*, *62*(9), 1757-1783.

Johnson, K. L., Freeman, J. B., & Pauker, K. (2012). Race is gendered: how covarying phenotypes and stereotypes bias sex categorization. *Journal of personality and social psychology*, *102*(1), 116.

Jordan, M. I. (1986). *Serial order: a parallel distributed processing approach. Technical report, June 1985-March 1986* (No. AD-A-173989/5/XAB; ICS-8604). California Univ., San Diego, La Jolla (USA). Inst. for Cognitive Science.

Jordan, M.I., & Rosenbaum, D.A. (1988). *Action* (Tech. Rep. *No.* 88-26). Amherst: University of Massachusetts, Department of Computer Science.

Kahneman D (2011) Thinking, fast and slow. Macmillan.

Koop, G. J., & Criss, A. H. (2016). The response dynamics of recognition memory: Sensitivity and bias. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *42*(5), 671.

Lagzi, F., & Rotter, S. (2015). Dynamics of competition between subnetworks of spiking neuronal networks in the balanced state. *PloS one*, *10*(9), e0138947.

Lazerus, T., Ingbretsen, Z. A., Stolier, R. M., Freeman, J. B., & Cikara, M. (2016). Positivity bias in judging ingroup members' emotional expressions. *Emotion*, *16*(8), 1117.

Lepora, N. F., & Pezzulo, G. (2015). Embodied choice: how action influences perceptual decision making. *PLoS computational biology*, *11*(4), e1004110.

Magnuson JS (2005) Moving hand reveals dynamics of thought. ProcNatl Acad Sci USA 102(29):9995–9996.

Mattek, A. M., Whalen, P. J., Berkowitz, J. L., & Freeman, J. B. (2016). Differential effects of cognitive load on subjective versus motor responses to ambiguously valenced facial expressions. *Emotion*, *16*(6), 929.

McKinstry C, Dale R, Spivey MJ (2008) Action dynamics reveal parallel competition in decision making. Psychol Sci 19:22–24

Mason W, Suri S (2012) Conducting behavioral research on Amazon's Mechanical Turk. Behav Res Methods 44(1):1–23

Mason WA, Watts DJ (2009) Financial incentives and the performance of crowds. In: Proceedings of the ACM SIGKDD workshop on human computation, pp 77–85.

Mead N, Baumeister RF, Gino F, Schweitzer M, Ariely D (2009) Too tired to tell the truth: self-control resource depletion and dishonesty. J Exp Soc Psychol 45:594–597

Murakami, M., & Mainen, Z. F. (2015). Preparing and selecting actions with neural populations: toward cortical circuit mechanisms. *Current opinion in neurobiology*, *33*, 40-46.

Nissen, M. J., & Bullemer, P. (1987). Attentional requirements of learning: Evidence from performance measures. *Cognitive psychology*, *19*(1), 1-32.

Nissen, M. J., Willingham, D., & Hartman, M. (1989). Explicit and implicit remembering: When is learning preserved in amnesia?. *Neuropsychologia*,*27*(3), 341-352.

Papesh, M. H. (2015). Just out of reach: On the reliability of the action-sentence compatibility effect. *Journal of Experimental Psychology: General*, *144*(6), e116.

Perfors, A., Tenenbaum, J. B., & Wonnacott, E. (2010). Variability, negative evidence, and the acquisition of verb argument constructions. *Journal of child language*, *37*(3), 607-642.

Quétard, B., Quinton, J. C., Mermillod, M., Barca, L., Pezzulo, G., Colomb, M., & Izaute, M. (2016). Differential effects of visual uncertainty and contextual guidance on perceptual decisions: Evidence from eye and mouse tracking in visual search. *Journal of vision*, *16*(11), 28-28.

Reber, A. S. (1967). Implicit learning of artificial grammars. *Journal of verbal learning and verbal behavior*, *6*(6), 855-863.

Reber, A. S., & Allen, R. (1978). Analogic and abstraction strategies in synthetic grammar learning: A functionalist interpretation. *Cognition*, *6*(3), 189-221.

Roitman, J. D., & Shadlen, M. N. (2002). Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *Journal of neuroscience*, *22*(21), 9475-9489.

Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1988). Learning representations by back-propagating errors. *Cognitive modeling*, *5*(3), 1.

Said, A., & Pearlman, W. A. (1993, November). Reversible image compression via multiresolution representation and predictive coding. In *SPIE Symposium on Visual Communications and Image Processing*.

Shannon, C. E. (1948). A mathematical theory of communication, Part I, Part II. *Bell Syst. Tech. J.*, *27*, 623-656.

Shannon, C., & Weaver, W. (1949). The mathematical theory of communication. Urbana, IL: University of Illinois Press.

Seymour TL (2001) A epic model of the 'guilty knowledge effect': strategic and automatic processes in recognition. Diss Abstr Int Sect B Sci Eng 61(10-B):5591

Seymour TL, Schumacher EH (2009) Electromyographic evidence for response conflict in the exclude recognition task. Cogn Affect Behav Neurosci 9(1):71–82

Shalvi S, Eldar O, Bereby-Meyer Y (2012) Honesty requires time (and lack of justifications). Psychol Sci 23:1264–1270

Song JH, Nakayama K (2008) Target selection in visual search as revealed by movement trajectories. Vision Res 48(7):853–861

Spence SA, Farrow TF, Herford AE, Wilkinson ID, Zheng Y, Woodruff PW (2001) Behavioural and functional anatomical correlates of deception in humans. NeuroReport 12:2849–2853

Spivey MJ, Dale R (2006) Continuous dynamics in real-time cognition. Curr Dir

Psychol Sci 15:207–211

Spivey MJ, Grosjean M, Knoblich G (2005) Continuous attraction toward phonological competitors. Proc Natl Acad Sci 102:10393–10398

Suri S, Watts DJ (2011) Cooperation and contagion in web-based, networked public goods experiments. PLoS One 6(3):e16836.

Smeding, A., Quinton, J. C., Lauer, K., Barca, L., & Pezzulo, G. (2016). Tracking and simulating dynamics of implicit stereotypes: A situated social cognition perspective.

Tabatabaeian, M., Dale, R., & Duran, N. D. (2014). Facilitation in Dishonesty is Subject to Task Constraints. *Proceedings of the 36th Annual Conference of the Cognitive Science Society*, 1562–1567.

Tabatabaeian, M., Dale, R., & Duran, N. D. (2015). Self-serving dishonest decisions can show facilitated cognitive dynamics. *Cognitive processing*, *16*(3), 291-300.

Verschuere B, De Houwer J (2011) Detecting concealed information in less than a second: response-latency based measures. In: Verschuere B, Ben-Shakhar G, Meijer E (eds) Memory detec- tion. Cambridge University Press, London

Vrij A, Mann SA, Fisher RP, Leal S, Milne R, Bull R (2008) Increasing cognitive load to facilitate lie detection: the benefit of recalling an event in reverse order. Law Hum Behav 32(3):253.

Walczyk JJ, Roper KS, Seemann E, Humphrey AM (2003) Cognitive mechanisms underlying lying to questions: response time as a cue to deception. Appl Cogn Psychol 17:755–774

Xiao, K., & Yamauchi, T. (2017). The role of attention in subliminal semantic processing: A mouse tracking study. *PloS one*, *12*(6), e0178740.

Yu, A. J., & Cohen, J. D. (2009). Sequential effects: superstition or rational behavior?. In *Advances in neural information processing systems*(pp. 1873-1880).