UNIVERSITY OF CALIFORNIA, SAN DIEGO

Acquisition of Network Graph Structure

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

in

Psychology

by

Jason Jeffrey Jones

Committee in charge:

Professor Harold Pashler, Chair
Professor Charles Elkan
Professor David Huber
Professor Timothy Rickard
Professor Terrence Sejnowski

2011

The Dissertation of Jason Jeffrey Jones is approved, and it is acceptable in quality and form for publication on microfilm and electronically:

_____

_____

_____

_____

_____
Chair

University of California, San Diego

2011

DEDICATION

To my parents,

James and Patricia.

TABLE OF CONTENTS

Signature Page..... iii

Dedication..... iv

Table of Contents..... v

List of Figures..... vii

List of Tables…..... ix

Acknowledgements..... x

Vita..... xi

Abstract..... xii

Chapter 1: Introduction..... 1

1.1 Why Study Network Graph Acquisition? 1

1.2 What is a Network Graph? 2

1.3 Organization of this Dissertation..... 3

Chapter 2: The Anthropomorphic Acquisition Hypothesis and the Learning of Network Graph Structure..... 5

2.1 Introduction and Experiment 1..... 5

2.2 Experiment 2..... 21

Chapter 3: Scale-Free Superiority, Egocentric Bias and Network Centrality Heuristics in Social Graph Learning..... 28

Chapter 4: Interactions of Graph Structure and Training Regimen: How Best to Teach Differing Graphs..... 56

4.1 Introduction and Experiment 1..... 56

4.2 Experiment 2..... 72

Chapter 5: Modeling Learning Performance..... 92

Chapter 6: Concluding Remarks..... 98

References..... 100

Appendix..... 105

LIST OF FIGURES

Figure 1: An example social network graph..... 5

Figure 2: Random graph, scale-free graph and degree distributions..... 13

Figure 3: d' in each condition in Experiment 1..... 17

Figure 4: Error rates as a function of trial by graph type..... 18

Figure 5: Error rates as a function of trial by surface description..... 20

Figure 6: d' in each condition in Experiment 2..... 22

Figure 7: The percentage of subjects' responses that were YES given the implied genders of the name stimuli..... 27

Figure 8: Example random, caveman and scale-free graphs..... 30

Figure 9: d' in each combination of experimental conditions..... 42

Figure 10: d' for test trials that did or did not include the "You" node, separated by Task Description..... 43

Figure 11: Accuracy as a function of eigenvector centrality..... 44

Figure 12: Hit rate as a function of eigenvector centrality..... 45

Figure 13: False alarm rate as a function of eigenvector centrality..... 46

Figure 14: d' as a function of eigenvector centrality..... 46

Figure 15: c as a function of eigenvector centrality..... 47

Figure 16: The number of potential edges in a graph as a function of the number of nodes..... 56

Figure 17: A ring lattice graph of 8 nodes..... 58

Figure 18: An irregular graph formed by rewiring two edges of the ring lattice..... 59

Figure 19: d' in each condition in Experiment 1..... 67

Figure 20: A Watts-Strogatz graph illustrating potential symmetry errors..... 71

Figure 21: Entropy as a function of the probability of an event X..... 74

Figure 22: Li entropy as a function of graph density for graphs of 4, 6 and 8 nodes…. 75

Figure 23: Network diagrams of the Low, Medium and High Density graph stimuli… 79

Figure 24: A valid edge training trial..... 81

Figure 25: An invalid edge training trial..... 81

Figure 26: A test trial in which the subject chose NOT FRIENDS as a response when the correct answer was FRIENDS..... 83

Figure 27: d' in each combination of graph density and training regimen condition..... 84

Figure 28: Accuracy as a function of training for subjects and five versions of the Perfect Memory Model..... 89

Figure 29: Accuracy as a function of training for subjects and PM5..... 90

Figure 30: Accuracy as a function of test node centrality for subjects and PM5..... 91

Figure 31: Accuracy as a function of training for subjects and the Frequency Heuristic Model..... 93

Figure 32: Accuracy as a function of training for subjects and the Frequency Heuristic Model..... 94

Figure 33: Accuracy as a function of test node centrality for subjects and FHM..... 95

Figure 34: A multinomial model that subsumes the Perfect Memory and Frequency Heuristic Models..... 97

LIST OF TABLES

Table 1: Truth Tables for Experiments 1 and 2..... 25

ACKNOWLEDGEMENTS

I would like to acknowledge Professor Harold Pashler for his role as the chair of my committee and for generous financial support over the course of my graduate career.

I would also like to acknowledge everyone who has provided feedback on this research. My work is better than it would otherwise be due to the questioning, challenging and critiquing of Professor James Fowler, the members of the Human Nature Group and the attendees of the UCSD Cognitive Brownbag.

Finally, I wish to acknowledge the patience, understanding, love and support of my family and friends during the course of my graduate education. Thank you all. I couldn't have done it without you.

VITA

2000 Bachelor of Science, Washington University in St. Louis

2005 Master of Arts, University of California, San Diego

2011 Doctor of Philosophy, University of California, San Diego

PUBLICATIONS

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

Jones, J.J. and Pashler, H. (2007). Is the Mind Inherently Forward-Looking? Comparing Prediction and Retrodiction. Psychonomic Bulletin & Review, 14(2), 295-300.

ABSTRACT OF THE DISSERTATION

Acquisition of Network Graph Structure

by

Jason Jeffrey Jones

Doctor of Philosophy in Psychology

University of California, San Diego, 2011

Professor Harold Pashler, Chair

A network graph describes the web of connections between entities in a system. Network graphs are a flexible abstraction; they are equally useful in representing which neurons communicate in a flatworm's brain and which international terrorists collaborate. Understanding how humans learn the structure of network graphs will be useful both to maximize the efficiency of teaching natural networks and to minimize cognitive complexity when designing artificial networks.

To this end, I conducted five experiments in which subjects learned which objects in a set were connected. For example, some learned "who is friends with whom" in a social network. These experiments yielded several results. Strong support was found for

the hypothesis that the deep structure of a graph affects how quickly it will be learned. Scale-free graphs were acquired more readily than other graphs. Much less support was found for the hypothesis that the surface description given a graph affects learning. For example, learning a social network was no easier or harder than learning a transportation network. The manner in which the learning *task* was described also had no discernible effect on the rate of acquisition.

Learning followed two patterns within each graph. First, if a node of strong salience was present (e.g. a person node labeled "You" in a social network), then edges involving that node were learned quickly. Second, learners responded to network centrality. They responded most accurately when queried on an edge that involved two central nodes.

In the last experiment, I show that graphs of medium density are difficult to train, and care must be taken to match training technique to graph structure. I also demonstrate that visual depictions of graphs are generally better teaching material than verbal descriptions.

Finally, I propose a general model for network graph acquisition. In this model, the learner initially relies on a frequency heuristic to identify which nodes have many connections and which have few. This allows for educated guessing regarding which edges are valid. Slowly, explicit knowledge is accrued until the subject has perfect knowledge and need no longer rely on a heuristic.

Chapter 1: Introduction

1.1 Why Study Network Graph Acquisition?

The success of Task Force 121's mission to find Saddam Hussein in post-war Iraq depended upon their understanding and use of social network graphs. Capturing and interrogating Baath government officials (i.e. the 54 men and 1 woman identified in the "Most Wanted Iraqis" deck of playing cards) yielded surprisingly little information as to Saddam's whereabouts. Saddam's personal social network graph – painstakingly pieced together by analysts using Hussein family photo albums and eyewitness accounts of who spent time with Saddam – ultimately identified the individuals most likely to know where he was. The central nodes in this graph turned out to be Saddam's closest bodyguards, and one of them – Mohammed Ibrahim Omar al-Musslit – pinpointed the location at which Saddam was found hiding.

This is but one example of a graph structure enabling important insight. Network graph models are a substrate for much current research. Network graphs have been used to describe both the series of chemical reactions that enable neuron-to-neuron communication (Bhalla & Iyengar, 1999) and the organization of the nervous system as a whole (Koch & Laurent, 1999). Graphs are in use in social science to explain the diffusion of behavior (McDermott, Fowler & Christakis, 2011) and ideas (Scherer & Cho, 2003) through networks of personal interactions.

As network theory has advanced recently, one important perspective has been missing. If more and more formerly unruly areas of research are to be tamed by the study of graphs, then it is important to understand how the human brain confronts the graph.

Graphs are expressive data structures and capable of boundless complexity. The brain is a very finite thing and limited in its ability to learn by the constraints of space, energy and especially time. How does the brain meet the challenge of representing these entities of combinatorial complexity? What is lost? How best to present graph structure to minimize that loss? These are important questions.

As a practical example, imagine yourself back in the year 2007. Someone hands you a list of the liabilities each bank in the United States has to other banks. Were you to look at that list, Bear Sterns and Lehman Brothers might look like just two more banks (as they did to regulators at the time). Transformed into a graph, you could make better use of this information. The central location of these nodes in the liability network would be alarming. Perhaps it would be useful for psychology to know how information about discrete relationships is aggregated into mental representations of graph structures.

1.2 What is a Network Graph?

A network graph is a formalism describing some number of entities and how they are connected. Those entities are represented by nodes in the graph. Nodes are connected by edges that represent the relationship of interest.

A canonical example is the social network graph. A social network contains people, and people are connected by a relationship – say friendship, for example. In the graph, nodes represent people and edges represent friendships.

In this work, I will use the words network and graph interchangeably. Both refer to the entire system – the collection of nodes and edges. The term “network graph” is used as well (mostly in titles) to differentiate graphs the formalism describing a network

from graphs that are graphical depictions of data (e.g. a bar graph). I also use the words "node" and "vertex" interchangeable to mean the same thing. Likewise, "edge," "link," and "connection" are synonymous.

An important distinction regarding edges is that between valid and invalid. A valid edge is one that exists in the graph. For example, the edge between Alice and Bob is valid in the social network if Alice and Bob are friends. An invalid edge is one that possibly could exist, but does not in the graph in question. The edge between Bob and Carol is invalid if Bob and Carol are both in the same network, but are not friends.

1.3 Organization of this Dissertation

Chapters 2 through 4 present experimental results. Each was written first as an independent manuscript. There may thus be some redundancy between these chapters, for which I apologize. Hopefully, the reader will agree that some of the points made bear repeating.

In Chapter 2, I introduce a hypothesis that will drive much of the research discussed in this dissertation. The effects of a graph's deep structure and its surface description on acquisition are investigated. In Chapter 3, the effects of additional graph structures and the properties of individual nodes on acquisition are explored. In Chapter 4, I address the practical question of how best to present graph structure information. Interactions with graph properties reveal this is a question with no simple answer. Chapter 5 contains two additional analytic approaches to the empirical data. Using the framework of signal detection theory, I confirm and further interpret previous conclusions. I also develop two computational models of learning for the task of graph

acquisition and compare their performance to human performance. Chapter 6 brings the dissertation to an end with concluding remarks.

Chapter 2: The Anthropomorphic Acquisition Hypothesis and the Learning of Network Graph Structure

2.1 Introduction and Experiment 1

For more than five decades, researchers have been studying how humans acquire and retain information about social networks. In this tradition, social networks are represented by graphs, and the graphs consist of nodes (representing people) and edges (representing the relationships between people). Figure 1 is an illustration of a social network represented as a graph. By manipulating such graphs – specifically by controlling which edges are present and which are absent – researchers create social network stimuli that conform to or diverge from people's natural expectations (De Soto, 1960).

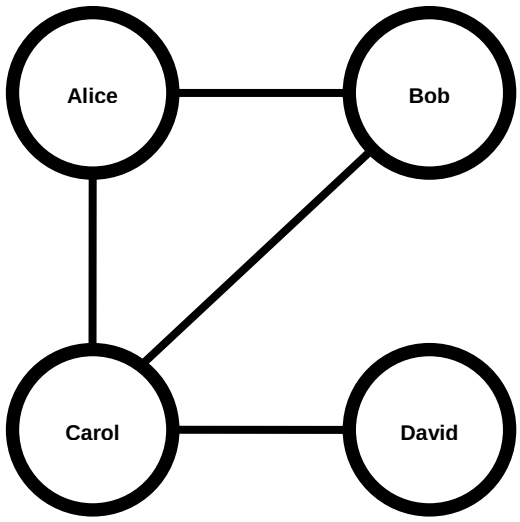

Figure 1. An example social network graph. Nodes represent people, and edges represent the relationships that exist between them.

Many reliable effects have been discovered and replicated numerous times. For instance, it is clear there is an expectation in people's minds that friendships are reciprocal (Cottrell, 1975; De Soto, 1960; De Soto, Henley & London, 1968; De Soto & Kuethe, 1959; Janicik & Larrick, 2005; Krackhardt & Kilduff, 1999; Walker, 1976). In other words, if Alice is friends with Bob, we assume Bob is friends with Alice.

Another consistent finding is an expectation of transitivity (De Soto, 1960; De Soto, Henley & London, 1968; De Soto & Kuethe, 1959; Janicik & Larrick, 2005; Krackhardt & Kilduff, 1999; Poitou, 1970; Tsujimoto & Robertson, 1978). Subjects assume that two people with a common friend are likely to be friends. For instance, if Alice is friends with Bob, and Bob is friends with Carol, then it is assumed Alice is also friends with Carol.

In the experimental work to date, researchers have severely limited the size of the social networks subjects learned. For the set of experimental papers cited in this article, the median and modal number of nodes in the stimulus graph structure is 4. It is easy to understand the reasons experimenters use small graphs. Small graphs allow for very tight control of experimental variables. Small graphs can be learned more quickly than large graphs, and many experimenters use perfect knowledge of the graph as the learning criterion. Large graphs quickly become very complex structures.

An unfortunate consequence of only studying the learning of small graphs is that little is known about how subjects acquire information about larger social networks. This article seeks to extend our knowledge of social network acquisition to larger graphs (about the size of a small class) while also proposing and testing two novel hypotheses.

The two hypotheses concern the effects of deep structure and surface description

on the rate at which connections within a network can be learned. The specific hypotheses (discussed momentarily) are driven by the more general *anthropomorphic acquisition hypothesis*. The anthropomorphic acquisition hypothesis states that humans use the social network as a metaphor when learning any new network. Humans have existing knowledge and beliefs about social networks that they apply to new networks. It follows that the more a novel network stimulus resembles a true human social network, the less adjustment to belief necessary during learning, and the more quickly it should be acquired. The experiments to follow will investigate whether the anthropomorphic acquisition hypothesis is borne out both on a deep level of abstract structural similarity and at a surface level of semantic descriptive similarity.

Scale-Free versus Random Graph Structure

The structural hypothesis begins with the assumption that human subjects will expect novel networks to resemble social networks they have encountered in the past. It has frequently been observed that well-documented human social networks reveal a formal structure best described as a scale-free graph (Csányi & Szendrői, 2004; Ebel, Mielsch & Bornholdt, 2002; Liljeros, Edling, Amaral, Stanley, Åberg, 2001; Wang, Moreno & Sun, 2006). In a scale-free social network, few people have many friends and many people have few friends.

(The name *scale-free* refers to the degree distribution. Scale-free graphs have many nodes that participate in few edges and a few nodes that participate in many edges. When plotted as a histogram, the degree distribution decays according to a power-law. The scale-invariance of power laws gives rise to the term scale-free.)

A natural graph structure to contrast with scale-free is the random graph. In a random graph, edges are randomly placed between nodes. In the context of a social network, a random graph would appear if in a group of people, dyads chosen at random became friends. A scale-free graph would arise if having more friends in the first place made one more likely to acquire new friends.

One can easily imagine how social networks come to be scale-free rather than random. It is probably not the case that in any dyad there is a constant chance of a friendship randomly occurring. Some individuals are easier to become friends with than others. Further, existing friendships increase the opportunities to meet new people and form new friendships. A large number of existing friendships is a signal that an individual would make a good friend, and might even serve as an incentive itself – if one is looking to maximize the number of people in one's extended network, adding a popular person as a friend is more profitable than befriending a loner. All of these forces act to concentrate friendships in a central core rather than spread friendships evenly among groups of people.

If past experience with human social networks drives expectations about novel networks, and past experience is likely to be with scale-free networks, the hypothesis that follows is that subjects will most quickly learn networks with a scale-free graph structure. In the experiments described in this paper, the hypothesis that scale-free network graphs are acquired more quickly than random network graphs will be tested directly.

Is Social Network Acquisition Exceptional?

Humans find human social networks exceptionally interesting. The online social network site facebook.com is the second-most visited website in the world, behind only

google.com (Alexa Internet, Inc., 2011). The movie *The Social Network* grossed nearly 100 million dollars in theaters (The Internet Movie Database, 2011). Before online social networks, people played The Kevin Bacon Game, the goal of which was to connect an actor to Kevin Bacon with as few co-starring links as possible. Academia is not immune to the fascination with social networks. Over the last ten years, the frequency of the term "social network" in the titles and abstracts of academic journal articles has increased at an average annual rate of 22% (Thomson Reuters, 2011).

Perhaps there is something exceptional about social networks that drives this fascination. There is at least one demonstration that people categorize human social relations differently than the relations between non-human stimuli. Cottrel (1975) had subjects classify stimuli based on three binary features. The stimuli consisted of two individuals' feelings toward an issue (each could be positive or negative) and their feelings toward each other (they could like or dislike each other). Consider this example: Doris is in favor of draft protests. Jenny is against draft protests. Doris and Jenny like each other. (The feelings toward each other were always reciprocal.)

Subjects in the study performed well when required to classify the stimuli based on balance. Balance, as defined by Heider (1958), occurs when two people's feelings toward each other are predicted by their feelings toward an object. If Doris and Jenny's opinions about draft protests are in agreement, they should like each other. If they disagree, they should dislike each other. If the state is of any other form, as it is in the example above, the situation is said to be unbalanced.

Subjects find balance the most difficult classification rule to learn when the stimuli are geometric forms rather than social relations. Shepard, Hovland and Jenkins

(1961) had subjects learn to categorize sequences of three shapes into two categories. Each shape could be large or small, so each stimulus consisted of three binary features. Balance dictates that the size of the third shape be dictated by the first two shapes. For example, the balance rule could require that if the first two shapes are unequal in size, then the third shape must be small. Otherwise the stimulus should be classified as unbalanced. (Another form of the balance rule could require that two unequal size shapes be followed by a large shape. The important distinction is that the required size of the third shape is dictated by the agreement or disagreement in size of the first two shapes.)

When subjects must learn the balance rule to categorize shape stimuli, they make more errors than they do for other rules and take more trials to reach steady perfect performance. This is not the case when the stimuli are social in nature. Even though the abstract form of the problem is the same, it is easier for people to learn that two individuals' feelings about an issue predict their feelings toward each other than it is to learn that the congruence of the sizes of two shapes should predict the size of a third shape.

In addition, some have made the argument that knowing how social networks are structured enhances one's own "social capital" and can lead to positive social and financial outcomes (Burt & Ronchi, 2007; Flynn, Reagans, Amanatullah & Ames, 2006). Thus, there may be advantages to being especially skillful at acquiring social network graphs that do not exist for other types of graphs.

On the other hand, there may be nothing exceptional about the acquisition of social networks. In favor of this hypothesis are the results of an experiment in which subjects' learning of two different graph structures was disrupted. It was found that

learning was slowed equally for social stimuli (people with military ranks) and non-social stimuli (nonsense syllables followed by numbers) (Poitou, 1970).

In the following experiments, learning of social, transportation and computer networks will be contrasted. If acquisition of social networks is exceptional, it will be apparent in the results.

EXPERIMENT 1

Method

Participants

102 UCSD undergraduates participated in the study.

Design

The design was a 2 (Graph Structure, within) x 3 (Surface Description, between) mixed design. Each subject was trained and tested on two Random graph structures and two Scale-Free graph structures. Upon beginning the experiment, each subject was randomly assigned to one of three surface descriptions: Social, Transport or Computer Network. All of the instructions and stimuli corresponded to this condition throughout the experiment.

Stimuli

All graphs consisted of 20 nodes. There were 190 potential edges (20 choose 2). Self-edge loops were not allowed. The number of potential graphs was very large: 2^{190}.

Graph Structure. Random graphs were generated by the process specified in Erdős & Rényi (1960). To form edges, two nodes were selected at random, and an edge placed between them. This process was repeated 20 times. If the two nodes drawn were

already connected, another draw was performed. As a result, each Random Graph contained 20 edges.

Scale-Free graphs were generated by the process specified in Barabási & Albert (1999). The graph began as two connected nodes. New nodes were added one at a time. Each new node formed an edge with each existing node with the probability $E_i/2N$, where E_i was the number of edges node *i* participated in and N was the total number of edges in the graph. To illustrate, the first node added to the seed set of two connected nodes had a 1/2 chance to form an edge with the first node and an independent 1/2 chance to form an edge with the second node.

This process leads to preferential attachment. The more edges a node already has, the more likely it will be to participate in new edges. This preferential attachment results in the power-law degree distribution characteristic of a scale-free graph. The resulting graphs contained a mean of 19.1 edges.

A deliberate choice was made to generate new graphs for each block for each subject. Using only one "Random" graph and one "Scale-Free" graph would greatly reduce the generalizability of any results found in that manner. Any results in that case could easily and probably appropriately be attributed to the particular graphs employed. By generating new graphs for each subject and each training block based on two different processes, much broader inferences can be made about the learning of the *family* of random graphs versus the *family* of scale-free graphs rather than a very narrow inference about one graph versus one other graph.

Figure 2 presents the combined degree distributions for all graphs used in this experiment and one random graph and one scale-free graph to illustrate the difference between graph structures.

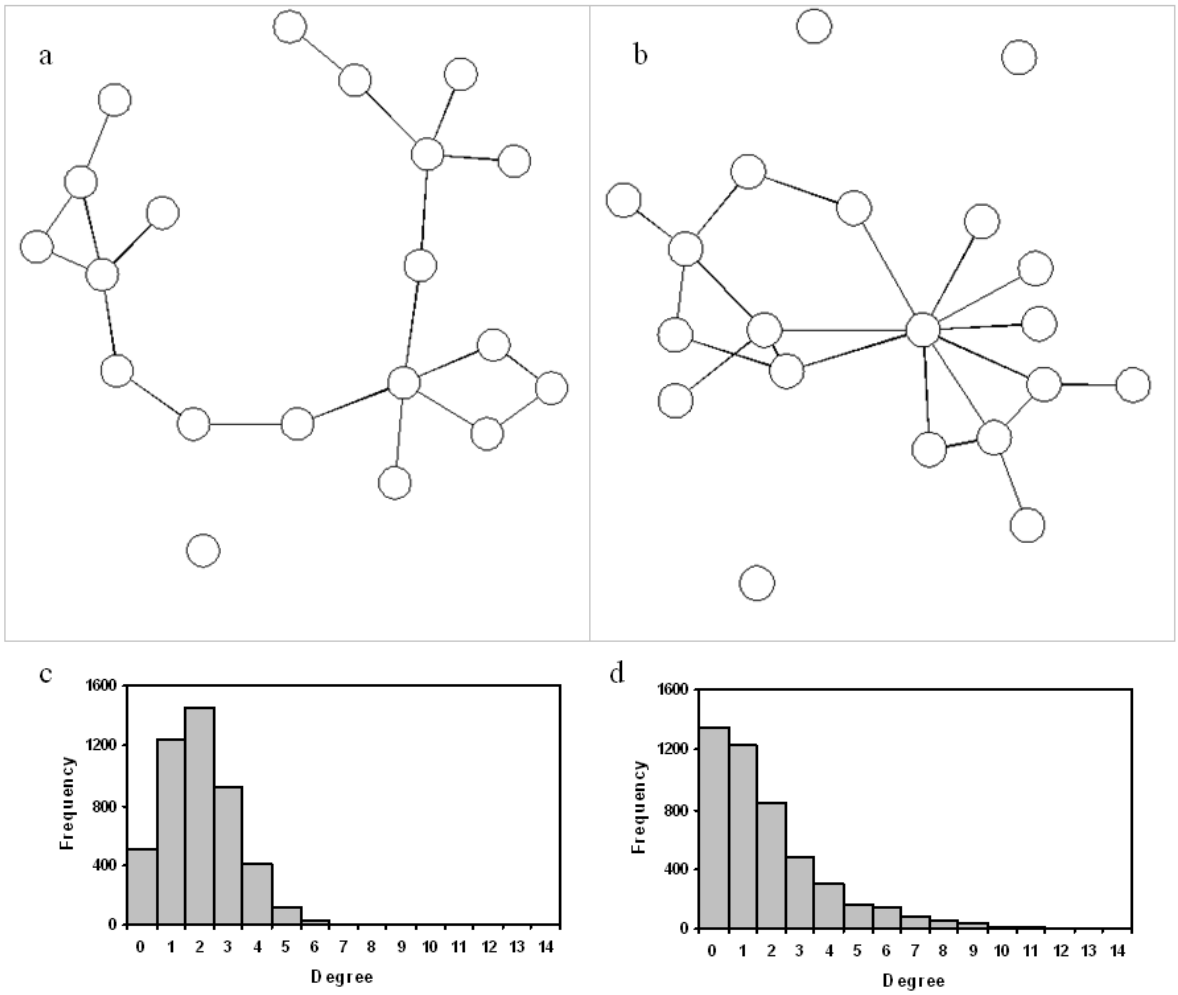

Figure 2. (a) An example of a random graph. (b) An example of a scale-free graph. (c) A histogram of the degree distribution for all nodes in all random graph stimuli. (d) A histogram of the degree distribution for all nodes in all scale-free graph stimuli.

Surface Description. The instructions and stimuli corresponded to one of three surface descriptions: Social, Transport or Computer Network. Subjects in the Social Network condition were instructed to learn "Who is friends with whom?" in a class of 20 students. Lists containing 10 male names and 10 female names were constructed by

choosing names from the Social Security Administration's list of popular baby names in the 1990's (Social Security Administration, 2011). The initial letter of each name was unique within each list.

Subjects in the Transport Network condition were instructed to learn which cities were connected by train tracks in a set of 20 cities. City names were drawn from online lists of small cities in various countries. The initial letter of each city name was unique within each list.

Subjects in the Computer Network condition were instructed to learn which computer servers were connected by fiber optic wires in a set of 20 computer servers. The server names were constructed by concatenating a letter, three digits and a ".net" or ".com" suffix. The initial letter of each server name was unique within each list.

Lists of all stimuli are included as Appendix A. At the beginning of each block, the stimuli names were randomly mapped to nodes in the abstract graph structure.

Procedure

The experiment was administered online through a PHP/mySQL/Flash Web application. Subjects followed a link to the experiment website from the subject pool administration software.

Subjects first read through several short paragraphs of instructions. They were informed that they would be learning the connections between 20 students/cities/computers. In the case of students, they were told to learn who was friends with whom. In the case of cities, they were told to learn which cities were connected by train tracks. In the case of computers, they were told to learn which computers were connected by fiber optic wires.

Subjects were informed that all links were reciprocal – if A was connected to B, then B was also connected to A. Subjects were also instructed to expect test trials in which two stimuli would be presented and they were to answer YES the items were connected or NO the items were not. Subjects were instructed to give their best guess if they did not know the answer.

In the experiment, training trials alternated with test trials. There were 4 blocks of 80 training and test trials each. In each block, the subjects were trained and tested on a new graph. Half of the subjects were given the pattern Random, Scale-Free, Random, Scale-Free, and the remaining subjects were given the complementary pattern.

Training Trials. In a training trial, the name of one node was presented at the top of the screen. Directly below, in a smaller font, were the words "is friends with" for social networks or "is connected to" for transport and computer networks. In two columns below these words appeared the names of all the nodes the trained node formed an edge with. If the trained node participated in no edges, the token --NONE-- appeared instead of the list of connected nodes.

Nodes were sampled for training without replacement from the list of 20 nodes until all had been shown. This process then repeated. Thus, over the course of 80 trials, each node was the focus of a training trial exactly four times.

The training trial information remained onscreen for 3 seconds. A blank, white screen followed for 1 second before the onset of a test trial.

Test Trials. In a test trial, two names appeared next to each other. Below the names was the prompt "Friends?" in the Social condition or "Connected?" in the

Transport and Computer conditions. Below the prompt were two equally-sized light-grey buttons labeled YES and NO.

Feedback was provided after the subject made a response. The button for the correct answer was highlighted, and a 300 millisecond sound clip of a bell was played for a correct response or a buzzer for an incorrect response.

For each network, a matching number of valid (existing) and invalid (non-existent) test edges were generated. All valid edges (approximately 20, depending on structure) and an equal number of invalid edges (chosen at random from the many possible invalid edges) comprised the set of test edges. These edges were randomly sampled without replacement until all had been shown. At that point, all of the constructed test edges became available again for re-testing. This procedure was followed so that responses should be distributed evenly between YES and NO, even though most edges in the graph were invalid.

Results

No significant effects of block were found. The remaining analyses collapse across blocks. d' within a block was the primary dependent variable.

A mixed-design ANOVA with Greenhouse-Geisser correction was performed to test for effects of graph structure and surface description. d' for each combination of conditions is graphed in Figure 3. Subjects' d' in Scale-Free blocks was .74 and was reliably higher than the .46 d' scored in Random blocks $F(1, 99) = 50.16$, $MSE = .08$, $p < .001$. The main effect of the between-subjects manipulation Surface Description was significant $F(2, 99) = 7.13$, $MSE = .43$, $p = .001$. Additionally, the interaction of Graph

Structure and Surface Description reached marginal significance $F(2, 99) = 2.74$, $MSE = .08$, $p = .07$.

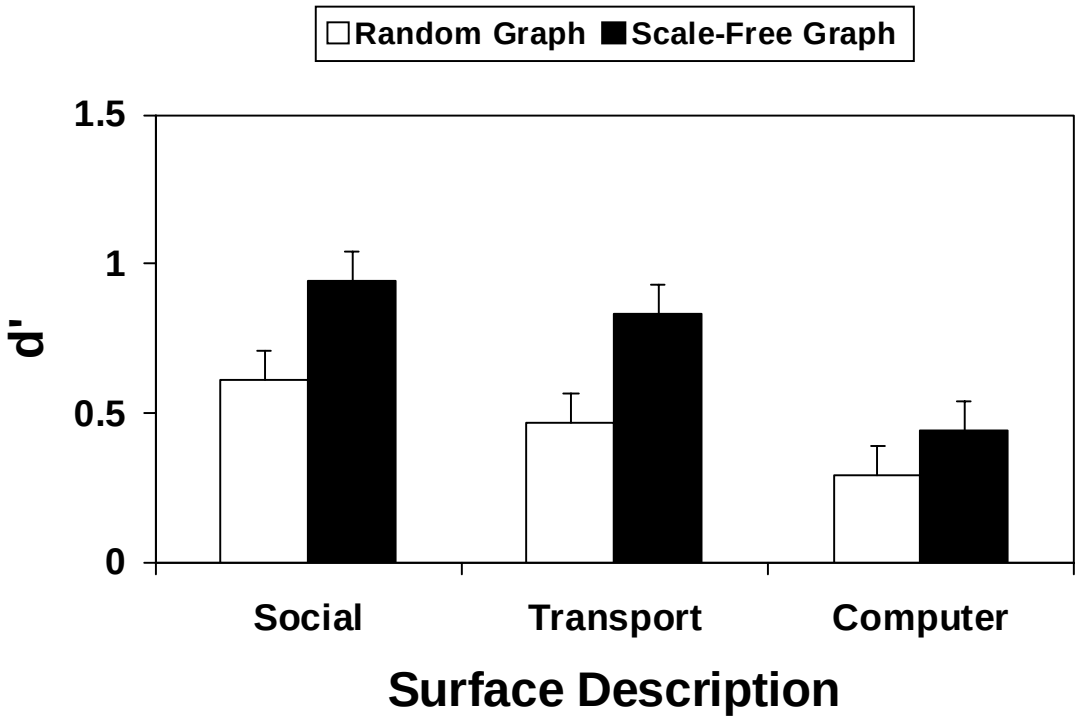

Figure 3. d' in each condition in Experiment 1. All values are reliably greater than chance performance of d' = 0.

Post-hoc comparisons imply the main effect of Surface Description was driven by the inferior performance of subjects in the Computer Network condition. d' was higher for both the Social Network $t(66) = 3.95$, $p < .001$ and Transport Network $t(66) = 2.56$, $p = .01$ conditions.

Discussion

There is strong support for the hypothesis that scale-free graphs are easier to learn than random graphs. Figure 4 depicts the learning curves under these two conditions.

While there is significant learning in both, the rate of learning is clearly faster for scale-free graphs.

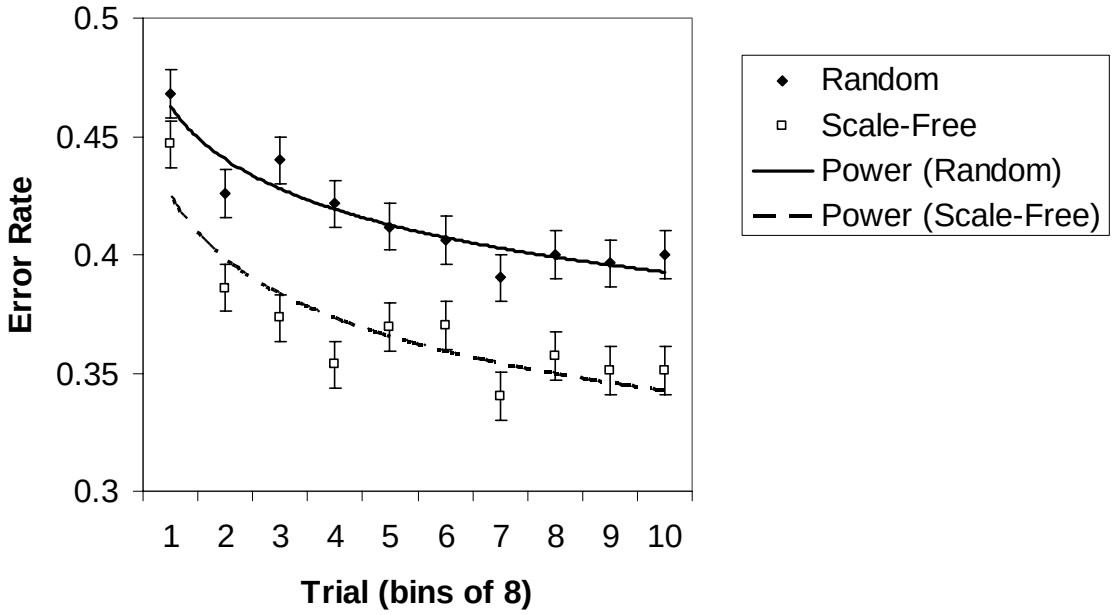

Figure 4. Error rates as a function of trial by graph type. Subjects acquired scale-free graph structure faster than random graph structure.

The Anthropomorphic Acquisition Hypothesis predicts this outcome because the scale-free network structure is congruent with the learner's expectations of network structure, because these expectations are based on social networks. However, the same effect could be due to the statistics of the learning procedure and the strategies subjects adopt to make educated guesses early in learning. For example, determining whether a node is highly-connected or isolated is a good initial strategy. It allows one to make a judgment as to whether an edge is valid based on knowledge of whether the nodes involved generally have many connections or few. Further, this is a better strategy for

learning a scale-free graph than a random graph. This is true because scale-free graphs have greater degree variability than random graphs. (See the distributions in Figure 2.)

Learners can quickly gain information about how many edges a node participates in based on two sources of information in the experiment. First, the name of a node that participates in many edges will appear many times during training. This is because it must appear in the list of connected nodes for every node it is connected to. Second, a node with many connections will appear with a long list when it is the trained node. Isolated nodes have short lists. As we will see when subjects' learning is simulated in Chapter 5, this type of learning that focuses first on *how many* (rather than which) connections a node has may be the best explanation for subject performance in early training.

The data concerning surface descriptions are mixed, however. Figure 5 depicts the learning curves under these three conditions. If one only compared the Social Network and Transport Network conditions, the conclusion one would reach is that performance is similar across surface descriptions. If one contrasted either with the Computer Network condition, however, one's conclusion might be that the structure of a computer network is uniquely difficult to acquire. Furthermore, the computer network condition was unique in the fact that subjects did not show a reliable learning advantage for scale-free over random graphs.

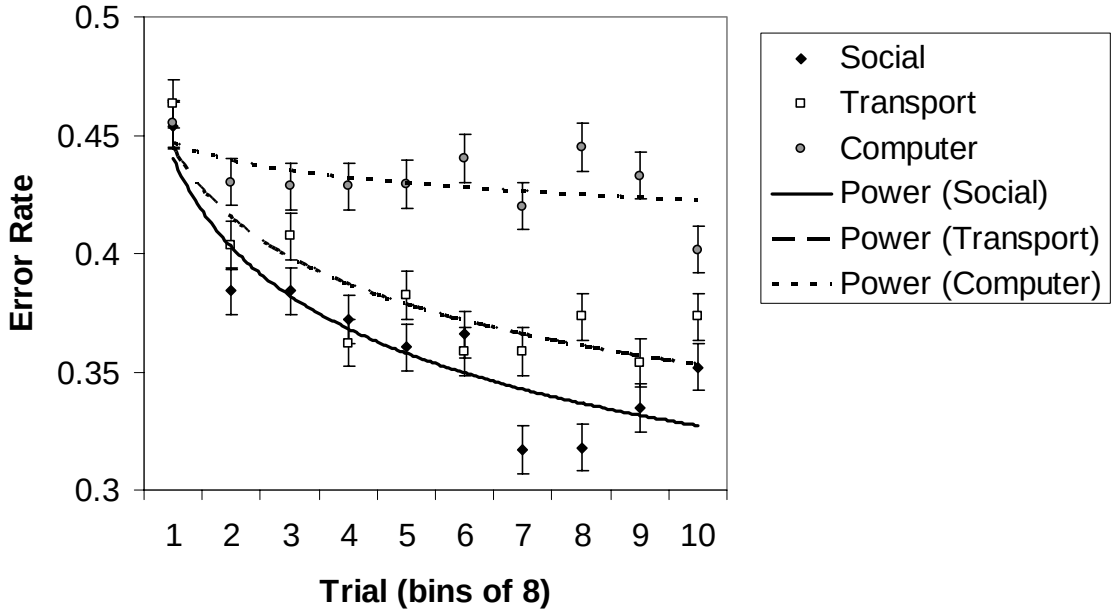

Figure 5. Error rates as a function of trial by surface description. Subjects' acquisition in the Computer condition was slower than in the other two conditions.

These results may have an analog in a previous study of associative recognition (Clark, 1992). In the associative recognition paradigm, subjects first study pairs of words (AB, CD, EF). Then, they must distinguish between intact pairs (AB) and rearranged pairs (AF) in a recognition test. In Clark's experiments, when the studied words were high-frequency words, subjects did a better job of discriminating intact pairs from rearranged pairs. Clark and others (Kelley & Wixted, 2001) have argued this effect is due to subjects' ability to augment pair recognition by sometimes recalling intact pairs. High-frequency words are more easily recalled than low-frequency words, and therefore associative recognition may be facilitated more often by recall. It could be argued that a similar effect accounts for the current results if we consider the pronounceable, word-like proper name stimuli of the Social and Transport conditions to be more like high-

frequency words and the unpronounceable, random string stimuli of the Computer condition to like low-frequency words.

It seems likely that the poor performance subjects displayed in the Computer Network condition was due to the node labels, rather than any consideration of the nature of the network. Therefore, a second experiment was undertaken to perform a more equitable test of the power of surface descriptions to affect the acquisition of graph structures. Specifically, the names of the nodes were equated across the surface description conditions.

2.2 Experiment 2

This experiment duplicates the design and procedure of Experiment 1 except for differences mentioned below.

Method

Participants

177 UCSD undergraduates participated in the study.

Stimuli

In this experiment, the four lists of names were the same for "students," "cities," and "computers." The instructions included the sentence "All of the names of the students/cities/computers will be foreign, so PAY CLOSE ATTENTION." The stimuli are listed in Appendix A.

Results

Again, no significant effects of block were found. The remaining analyses collapse across blocks. d' for each combination of conditions is graphed in Figure 6.

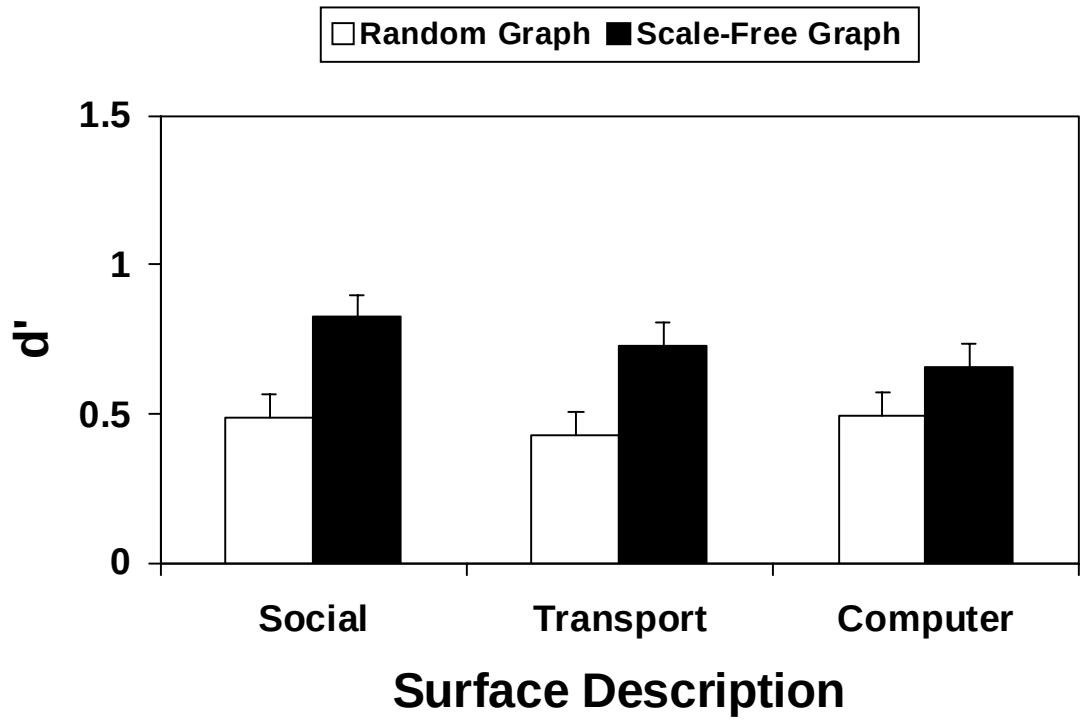

Figure 6. d' in each condition in Experiment 2. All values are reliably greater than chance performance of d' = 0.

The main effect of graph structure was replicated. Subjects' d' in Scale-Free blocks was .74 and was reliably higher than the .47 d' scored in Random blocks $F(1, 174) = 66.45$, $MSE = .09$, $p < .001$. The main effect of Surface Description was not significant $F(2, 174) = .51$, $MSE = .46$, $p = .60$. The interaction of Graph Structure and Surface Description again reached marginal significance $F(2, 174) = 2.50$, $MSE = .09$, $p = .09$. Unlike in the last experiment, d' was reliably higher in the Scale-Free + Computer condition than in the Random + Computer condition $t(58) = 2.86$, $p = .006$.

General Discussion

It is clear that the structure of a graph determines how quickly it may be learned. Knowledge of the structure of scale-free graphs was acquired more quickly and more completely than random graphs of similar density.

One interpretation of these results is a validation of the anthropomorphic acquisition hypothesis. Subjects used their knowledge of social network structure to shape their expectations concerning all novel networks. True social networks are of scale-free structure. Therefore, new networks of scale-free structure are learned most efficiently regardless of the supposed context given by the surface description.

An alternative explanation is possible by way of information theory. Intuitively, it would seem that it requires more information to describe a random graph than a scale-free graph. Take, for instance, the graphs in this experiment. Between 20 nodes, there exist 190 potential undirected edges (self-edges excluded). The maximum number of bits necessary to describe one graph in particular is 190 – one bit for each edge to denote if it is valid or invalid.

To learn the structure of a random graph requires all 190 bits to be provided. A partial description of the graph still leaves maximum uncertainty about the structure of the unrevealed portions of the graph. The known edges provide no information about which unknown edges exist or do not.

In contrast, a partial description of a scale-free graph limits the possible configurations of the rest of the graph. If it is assumed (or induced) that some nodes attract edges more readily than others, then the known edges place probabilities on the existence of the unknown edges. (Note that it must be assumed or induced that the graph is scale-free for this to be true, however. If one fully expects a graph to have random structure, even the revelation of 19 edges of a scale-free graph does not inform one of where the 20th edge “should” be.)

The evidence is not clear as to whether surface description alone has an effect on graph acquisition. In Experiment 1, it appeared that describing a graph as the set of fiber optic connections between computer servers made the structure of those graphs especially difficult to learn. However, that deficit was successfully erased by giving the servers more pronounceable, less confusable names. One could argue that computer network engineers have already discovered this effect. 74.125.224.114 is referred to as google.com rather than the numeric IP address. Computer servers in data centers are often named after hobbits, Star Wars characters or other sets of named entities (Libes, 1989). No doubt these names make learning the connections between computers easier.

In no case was the learning of networks described as "social" exceptional. In this experiment, it did not matter whether subjects believed they were learning who is friends with whom, which cities are connected by train or which computers are wired together. In this experiment, it appears that the learning mechanism functioned in a manner agnostic to context and learned graphs based on their deep structure with no regard to the surface description.

Yes Bias

Table 1 contains truth tables for both experiments collapsed across conditions. Subjects display a strong YES bias. (The bias was similar across conditions, so separate tables per condition are not shown.) This bias is especially surprising considering how few edges actually exist in the graphs to be learned. Most edges in the graphs were invalid, and yet subjects were more likely to respond YES than NO when asked if two nodes were connected.

Table 1. Truth Tables for Experiments 1 and 2.

Experiment 1				
		True Answer		
		No	Yes	Marginal
Subject Response	No	26 %	16 %	42 %
	Yes	24 %	34 %	58 %
	Marginal	50 %	50 %	

Experiment 2				
		True Answer		
		No	Yes	Marginal
Subject Response	No	28 %	17 %	45 %
	Yes	22 %	33 %	55 %
	Marginal	50 %	50 %	

A similar effect was reported in one of the first experimental studies of the perception of social relations. De Soto and Kuethe (1959) had subjects judge the probability of one relation given another. For example, subjects were asked, "Jack likes Dave. Does Russ like Dave?" In what must have been a perplexing moment for subjects, some trials did not include a given relationship. Instead there was only a question (e.g. "Does Les like Al?"). For these types of questions, subjects assigned a subjective probability of 59% to a yes response. This subjective probability was reliably greater than 50% and higher than that for other relations such as "is happier than," "dominates," and "dislikes."

Both the past and the current finding suggest people have a tendency to assume two people are friends until evidence is provided otherwise. However, the yes bias in these experiments was not specific to the social network conditions. In fact, it appears in every condition. It may be that a yes bias is characteristic of graph learning in general.

It should also be noted that the yes bias decreased across learning trials within a block and across blocks. As subjects learned the structure of a particular graph and became familiar with the task in general, they became better calibrated.

Intra versus Intergender Friendships

In the Social condition of Experiment 1, the names used as stimuli strongly imply gender. Indeed, the names were chosen from the Social Security Administrations lists of most common boy and girl baby names.

Therefore, in each test trial, the friendships in question could be either intragender (MM or FF) or intergender (MF or FM). Interestingly, this factor had no discernable effect on accuracy. By the end of the experiment, subjects had learned to associate MF and FM pairs just as well as MM and FF pairs.

Subjects had to adjust their prior beliefs to reach this point, however. Figure 7 presents the percentage of responses that were YES categorized by the genders of the two test stimuli names. The general tendency toward a YES bias is evident in that all values are above 50%. The YES bias is even greater for intragender pairs. A paired-samples t-test comparing frequency of YES responses for intra versus intergender pairs within subjects revealed a significant difference $t(33) = 3.71$, $p = .001$.

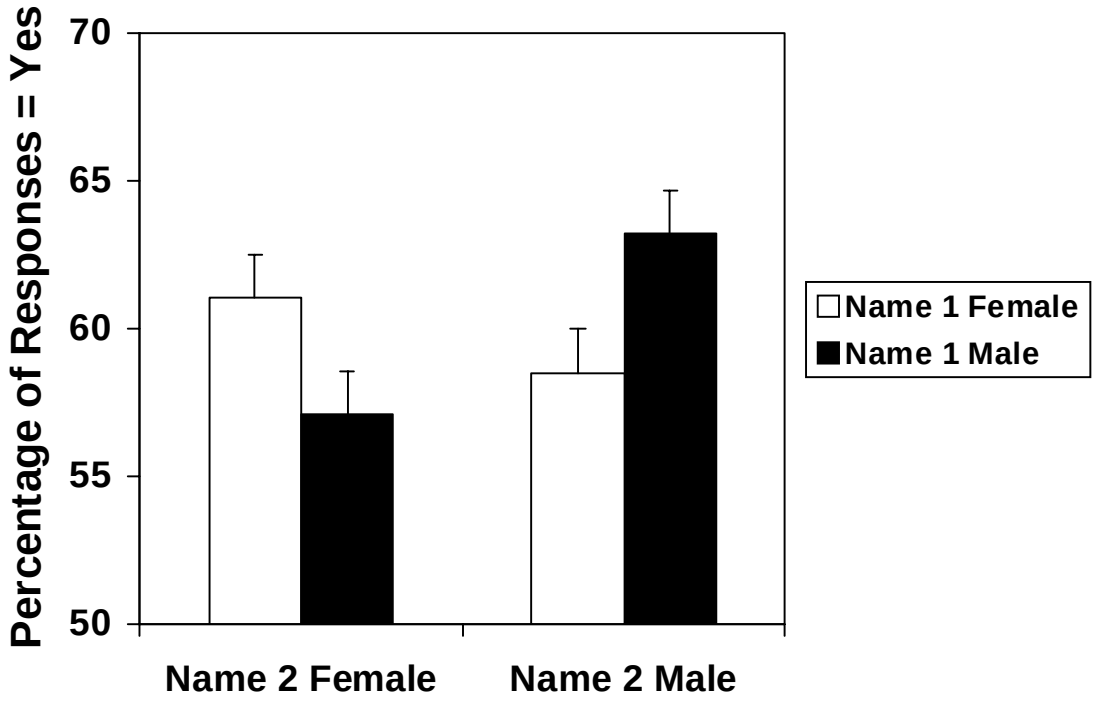

Figure 7. The percentage of subjects' responses that were YES given the implied genders of the name stimuli.

Summary

This investigation is the first to experimentally test the acquisition of social networks larger than just a handful of nodes. Two novel hypotheses were tested. Support was found for the hypothesis that humans learn scale-free graphs more efficiently than random graphs. Evidence was mixed as to whether the surface description of a network affects its acquisition. In this experiment, it did not appear to matter whether a graph was described as the set of connections between people, cities or computers. As long as the labels assigned to nodes were remained constant across domain, all graphs were learned equally well.

Chapter 3: Scale-Free Superiority, Egocentric Bias and Network Centrality Heuristics in Social Graph Learning

The intelligence teams that tracked down both Saddam Hussein and Osama Bin Laden made use of the tools and terminology of social network analysis. In the case of Saddam Hussein, success came when the central nodes in his informal social network were identified. A pair of brothers who had been childhood friends with and were the closest bodyguards of Saddam were the key captures that brought the U.S. military's hunt to Saddam's spider hole. In the case of Bin Laden, the important break came by finding a courier that linked the network of Al Qaeda in Iraq and the leadership of Al Qaeda proper. It was this broker between networks who held the critical knowledge necessary to track the people close to Bin Laden, and ultimately the terrorist leader himself.

Not all social network graphs involve high-value intelligence targets, obviously. Each of us has our own graph, which we must come to know and constantly update in order to make new friends, find new opportunities and avoid uncomfortable social missteps. With this paper, I hope to begin a discussion concerning how social network structure is acquired and how it is represented in the mind.

Beyond the immediate practical benefits that may be delivered – such as suggestions for training intelligence analysts more efficiently – the study of social network acquisition will supply data pertaining to broader theoretical questions. Understanding how people come to know social network graph structures will provide insight as to how people come to know graph structures generally. Given that formal graphs are a data structure well-suited to represent a wide range of complex systems

(websites connected by hyperlinks, proteins connected by coparticipation in metabolic reactions, neurons connected by synapses, banks connected by liabilities or any other phenomena in which entities form ties) it is necessary that we study how human cognition comprehends network structure.

Previous work has identified a number of reliable effects that emerge when people are tasked with learning a novel social network. People expect friendship relations to be reciprocal (Cottrell, 1975; De Soto, 1960; De Soto, Henley & London, 1968; De Soto & Kuethe, 1959; Janicik & Larrick, 2005; Krackhardt & Kilduff, 1999; Walker, 1976). That is, if Alice is friends with Bob, we assume Bob is friends with Alice.

People expect friendship relations to be transitive (De Soto, 1960; De Soto, Henley & London, 1968; De Soto & Kuethe, 1959; Janicik & Larrick, 2005; Krackhardt & Kilduff, 1999; Poitou, 1970; Tsujimoto & Robertson, 1978). Network learners assume that two people with a common friend are likely to be friends. For instance, if Alice is friends with Bob, and Bob is friends with Carol, then it is assumed Alice is also friends with Carol.

These studies predominately used small graphs as stimuli. (Most studies use 4-person social networks.) Additionally, the effects under examination involved dyads or triads rather than the graph structure as a whole. As a necessary evil, the current experiment eschews the simplicity and clean design afforded by small graphs to extend our knowledge of graph acquisition to larger structures. Instead of manipulating dyad or triad structure, the current study examines the effect of gross structural differences.

Three Graph Structures: Random, Caveman and Scale-Free

To assess how the deep structure of a social network graph affects how quickly it can be acquired, this experiment makes use of three categorically different graph structures. One example of each graph type is pictured in Figure 8. The general hypothesis under examination is that some structures will be more quickly acquired than others. More desirable would be a hypothesis that makes specific predictions about the ordering of acquisition difficulty. To fill that role, I propose the anthropomorphic acquisition hypothesis: the more a graph stimulus resembles true human social networks, the faster it will be acquired. In other words, the more properties a novel social graph to be learned shares in common with the general properties of real-world human social networks, the faster it will be learned. To understand how the anthropomorphic acquisition hypothesis applies to the current study, one must first know how much the graph structures to be used as stimuli have in common with true human social networks. To that end, these three structures are placed in an (admittedly somewhat subjective) order from least "humanlike" to most humanlike.

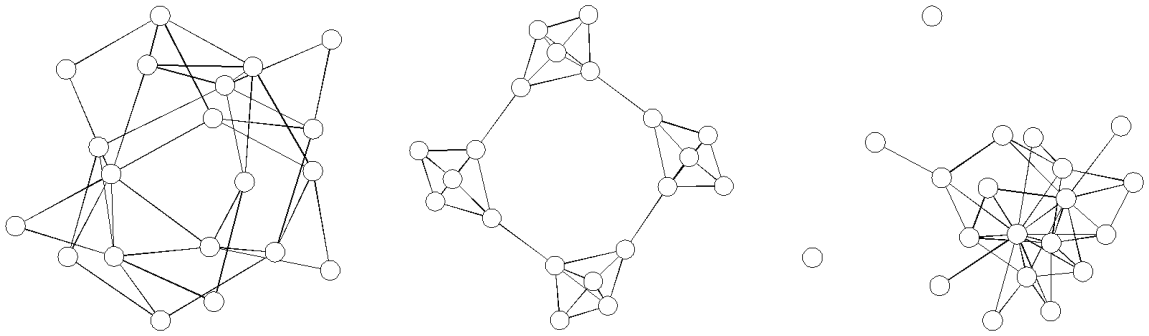

Figure 8. Left: an example random graph. Center: an example caveman graph. Right: an example scale-free graph.

Random graphs (Erdős & Rényi, 1960) are least like true human social networks. In a random graph, edges are placed randomly between nodes. In a social network, this would be the equivalent of each pair of people in a group having the same uniform probability of becoming friends. This process of placing edges leads to networks with low transitivity. This is unlike human social networks which exhibit higher transitivity, because in the real world, having a common friend makes it more likely that two people are friends.

Random graphs also exhibit a Poisson degree distribution, unlike human social networks. In random graphs, each node (or person) is likely to have a number of edges (friendships) that is close to the mean. The degree distribution of human social networks is better described by a power-law curve than a Poisson. In human social networks, a great many people have few friends and few people have a great many friends.

In these two ways, random graphs are unlike true human social networks.

Caveman graphs (Watts, 1999) are more like true human social networks. In a caveman graph, nodes are first organized into several disjoint, fully intraconnected cliques (or caves). Then, in each cave, one *intra*cave edge is removed and an *inter*cave edge created in its place. All caves become connected through a central ring of individuals. A caveman graph has high transitivity due to the cliquishness of the caves. In this way it imitates human social networks.

However, the degree distribution of a caveman graph is unlike that of a human social network. In fact, almost all the nodes in a caveman graph participate in the same number of edges. This fact makes the caveman graph a bad analog of human social

networks, because it clashes with the expectation that a few people should have many friends and many people should have few friends.

Scale-free graphs (Barabási & Albert, 1999) are most like true human social networks. In a scale-free graph, there are central hubs (nodes with many links) and more peripheral nodes with fewer links. In a social network, this would be the equivalent of a few popular people who are friends with many others, while most people have just a few friends or none at all.

Scale-free networks have high transitivity, because hub nodes are likely to be connected and form many transitive triads. Additionally, scale-free graphs have a power-law degree distribution, similar to the degree distributions of true human social networks. It has frequently been observed that well-documented human social networks reveal a formal structure best described as a scale-free graph (Csányi & Szendrői, 2004; Ebel, Mielsch & Bornholdt, 2002; Liljeros, Edling, Amaral, Stanley, Åberg, 2001; Wang, Moreno & Sun, 2006).

Task Description

Will the learner of a social network graph structure have more or less success depending upon how this task has been described to them? On the surface, it is a very different task to learn who is friends with whom in a class one is observing versus learning who is friends with whom in a tense survival situation one is personally involved in. The abstract, deep structure of the problem is the same, but the surface framing of the problem is different.

It is difficult for people to abstract away irrelevant details in the description of a problem (Gick & Holyoak, 1983; Rein & Markman, 2010). The framing one uses to

understand a system can affect how accurately it is understood. For example, Gentner and Gentner (1983) used two different analogies to train participants on the principles of electrical circuits involving batteries, wires and resistors. Some subjects were instructed to think of an electrical circuit as a hydraulic system of reservoirs, pipes and constrictions in said pipes. Other subjects were instructed to think of an electrical circuit as a track on which a moving crowd of mice ran away from a loudspeaker and through narrow gates. Among other effects, Gentner and Gentner found that subjects who were given the moving crowd task description were better at solving problems involving resistors wired in parallel. Participants had better natural intuitions concerning what would happen when a crowd must navigate through tight alleyways than they did concerning steady-state hydraulics. Although the abstract structure of the problem did not differ between these groups, their performance did, due to the manner in which the problem had been framed.

In the current study, the description given to the task of social network graph learning was varied not by instructing participants to use different metaphors, but instead by varying the extent to which acquiring information about the graph was personally relevant. The initial hypothesis was that participants would acquire the structure of a social graph slowly when it was described as a class they were observing, more quickly when it was described as a class the participant would themselves be participating in and fastest when it was described as a group of people (including the participant) in a survival situation (Nairne, Pandeirada & Thompson, 2008).

It is not always possible to vary personal relevance in learning and memory tasks. However, when it is done, increasing personal relevance often aids retention. When words are rated as to how easily they bring to mind personal memories, they become

more likely to be recalled later (Challis, Velichkovsky & Craik, 1996; Craik & Tulving, 1975; Rogers, Kuiper & Kirker, 1977). In addition, the higher words are rated on this self-reference scale, the larger the effect (Craik & Tulving, 1975; Rogers, Kuiper & Kirker, 1977).

There is also some neuropsychological support for the idea that information relative to the self receives special processing in the brain. Klein, Cosmides, Murray and Tooby (2004) recently provided provocative evidence for a dissociation between acquiring information about the self and acquiring information about others.

A striking example of personal relevance enhancing acquisition was recently demonstrated with participants learning about stroke symptoms (McDonald, et. al., 2009). In this study, participants were given a pretest and posttest with questions regarding the correct response to stroke symptoms. Between tests they read a pamphlet about stroke symptoms and the appropriate response to each. Although both groups increased their scores from pretest to posttest, subjects who read the statement "Learn about stroke to save someone you love" learned the correct response to two more items on average than subjects in the control condition who did not receive the personal relevance statement.

Survival might be described as the interest one holds most extremely personally relevant. Survival processing has recently been proposed as a mnemonic device effective for promoting retention (Nairne, Pandeirada & Thompson, 2008; Nairne, Thompson & Pandeirada, 2007). The central idea behind survival processing is that the human memory system is tuned to retain information important for survival. Therefore, when stimuli are evaluated in terms of how much they pertain to the individual's survival (i.e.

survival processing) those stimuli stand a higher chance of being marked survival-relevant and retained in memory.

It seems natural to test the efficacy of survival processing in the current study. If there was any domain in which it was important for our ancestors to consider survival, surely it would have been the social domain. It would be difficult to argue that knowing one's place in the social order and knowing who was allied with whom were not among the highest survival priorities in our evolutionary past.

Node Attribute Analysis

The experimental design of the current study calls for two independent variables: graph structure and task description. Thus, conclusions can be drawn concerning the causal effect of varying the deep structure and surface description of social network graph learning. However, we need not stop there. Because some conditions include "You" in the list of social network nodes, tests will reveal if this node becomes special in participants' learning. Will there be an egocentric bias to learn one's "own" friends?

Also, nodes will necessarily differ in centrality – a measure of how integral a node is to the pattern of connections in a graph. A node's centrality reflects whether it is isolated and peripheral or well-connected and central to the community of nodes. One common measure of centrality is eigenvector centrality. The *eigenvector centrality* of a node is determined not just by the number of edges the node participates in, but also by the centrality of the nodes on the other end of each edge. It is best understood by way of an example. The PageRank formula used by the Google search engine computes the eigenvector centrality of a webpage in the graph of webpage hyperlinks (Brin & Page, 1998). More incoming links from other pages increase PageRank, but not all links are

counted equally. The more incoming links of its own that a linking page has, the greater its contribution to PageRank. In a social graph, eigenvector centrality increases with number of friendships, but more so with the number of other friends each friend has. (See Bonacich, 1987 for a derivation of and methods for calculating eigenvector centrality.) Will learners respond to centrality, or will they acquire knowledge about who is connected to whom without regard to the properties of individual nodes?

Method

Participants

135 UCSD undergraduates participated in the study.

Design

The design was a 3 (Graph Structure) x 3 (Task Description) mixed design. Graph Structure was varied within subjects, and all subjects were trained on one graph of each type (Random, Caveman, and Scale-Free). The order in which the structures were presented was chosen randomly for each subject.

Task Description was varied between subjects. Participants received the same task description for each of the three graphs they were to learn. The task description (one of Class, Class + You, or Survival + You) was assigned randomly to each subject.

Stimuli

All graphs consisted of 20 nodes. Edges were undirected and unweighted. Self-edge loops were not allowed. In each graph, there were 190 potential edges (i.e. 20 choose 2). Participants were trained on three graph structures: Random, Caveman and Scale-Free.

Graph Structure. Figure 8 depicts an example of each graph structure. Random graphs were generated by the process specified in Erdős & Rényi (1960). To form edges, two nodes were selected at random, and an edge placed between them. This process was repeated 40 times. If the two nodes drawn were already connected, a replacement draw was performed. As a result, each Random Graph contained 40 edges.

Caveman graph structures (Watts, 1999) were not generated by a stochastic process, unlike the Random and Scale-Free graphs. The Caveman graph consisted of four groups (caves) of five nodes. Each cave was fully connected – meaning all nodes in the cave were connected to *all* other nodes in the cave. To connect the caves, one node in each cave broke one intracave edge and replaced it with one intercave edge. The Caveman graph contained 40 edges. Names were randomly mapped to nodes, so while the structure of every Caveman graph was the same, the actual stimuli presented to subjects varied. It was not the case that the same names appeared in the same structural positions for these graphs or any others.

Scale-Free graphs (Barabási & Albert, 1999) were generated through two rounds of preferential attachment edge formation. In the first round, the graph began as two connected nodes. New nodes were added one at a time. Each new node formed an edge with each existing node with the probability $E_i/2N$, where E_i was the number of edges node *i* participated in and N was the total number of edges in the graph. To illustrate, the first node added to the seed set of two connected nodes had a 1/2 chance to form an edge with the first node and an independent 1/2 chance to form an edge with the second node.

The same process was followed in the second round with minor adjustments. Each node *i* formed a new edge with each other node *j* in the graph with the probability

(E_j + 1)/2N. One was added to each node's edge count so that nodes left without any edges in the first round would have a probability greater than zero of acquiring edges in the second round. The resulting graphs contained a mean of 38.8 edges.

A deliberate choice was made to generate Random and Scale-Free graphs stochastically. Using only one "Random" graph and one "Scale-Free" graph would greatly reduce the generalizability of any results found in that manner. Any results in that case could easily and probably appropriately be attributed to the particular graphs employed. By generating new graphs for each subject based on two different processes, we can make much broader inferences about the learning of the *family* of random graphs versus the *family* of scale-free graphs rather than a very narrow inference about one graph versus one other graph. The Caveman graph structure is much more constrained (given a set number of caves), and most variants are isomorphic, so the same specification was used for all participants.

Task Description. Participants were randomly assigned one of three task descriptions: Class, Class + You, or Survival + You. These conditions differed in the text of the instructions that preceded each training block.

The Class instructions read:

> *Imagine you are a sociologist studying a class of 20 students in order to write a report. The first thing you must do is learn who is friends with whom in this class.*

The Class + You instructions read:

> *Imagine you are taking a class with 19 other students. You want to learn who is friends with whom in this new class. You may already be friends*

with some of the students yourself.

The Survival + You instructions read:

Imagine your plane has crashed and you are stranded in the grasslands of a foreign land, without any basic survival materials. Over the next few months, you'll need to find steady supplies of food and water and protect yourself from predators. 19 other people are stranded with you, some of whom are friends already. It will be important to your survival to learn who is friends with whom.

Name Lists. Three lists containing 10 male names and 10 female names were constructed by choosing names from the Social Security Administration's list of popular baby names in the 1990's (Social Security Administration, 2011). The initial letter of each name was unique within each list. At the beginning of each block, the stimuli names were randomly mapped to nodes in the abstract graph structure. In the Class + You and Survival + You conditions, one name was chosen at random and replaced with the word "You."

Procedure

The experiment was administered online through a PHP/mySQL/Flash Web application. Subjects followed a link to the experiment website from the subject pool administration software.

Subjects first read through several short paragraphs of instructions. They were told they would be learning "who is friends with whom" in a group of 20 people.

Subjects were informed that all links were reciprocal – if A was connected to B, then B was also connected to A. Subjects were also instructed to expect test trials in

which two stimuli would be presented and they were to answer YES the people were friends or NO they were not. Subjects were instructed to give their best guess if they did not know the answer.

In the experiment, training trials alternated with test trials. There were 3 blocks of 80 training and test trials each. In each block, the subjects were trained and tested on a new graph. Participants were trained on one graph of each Graph Structure. The order was determined by randomly permuting the three types for each participant.

Training Trials. In a training trial, the name of one node was presented at the top of the screen. Directly below, in a smaller font, were the words "is friends with." In two columns below these words appeared the names of all the nodes the trained node formed an edge with. If the trained node participated in no edges, the token --NONE-- appeared instead of the list of connected nodes.

Nodes were sampled for training without replacement from the list of 20 nodes until all had been shown. This process then repeated. Thus, over the course of 80 trials, each node was the focus of a training trial exactly four times.

The training trial information remained onscreen for 5 seconds. A blank, white screen followed for 500 milliseconds before the onset of a test trial.

Test Trials. In a test trial, two names appeared next to each other. Below the names was the prompt "Friends?". Below the prompt were two equally-sized light-grey buttons labeled YES and NO.

Feedback was provided after the subject made a response. The button for the correct answer was highlighted, and a 300 millisecond sound clip of a bell was played for a correct response or a buzzer for an incorrect response.

For each graph, an equal number of valid (existing) and invalid (non-existent) test edges were generated. All valid edges (approximately 40, depending on structure) and an equal number of invalid edges (chosen at random from all invalid edges) comprised the set of test edges. These edges were randomly sampled without replacement until all had been shown. At that point, all of the constructed test edges became available again for re-testing. This procedure was followed so that responses should be distributed evenly between YES and NO, even though most edges in the graph were invalid.

Results

d' for each combination of conditions is graphed in Figure 9. In all conditions, d'was statistically significantly greater than the value of zero expected from uninformed guessing. A mixed-design ANOVA with Greenhouse-Geisser correction was performed to test for effects of graph structure and task description. The main effect of Graph Structure was significant $F(2, 264) = 14.88$, $MSE = .22$, $p < .001$. Subjects' d' in Scale-Free blocks was .75 and was reliably higher than the .53 d' in Caveman blocks and the .46 d' scored in Random blocks. The slight apparent advantage in Caveman blocks was not significant $t(134) = 1.31$, $p = .19$. No significant effect of Task Description was found nor any interaction with Graph Structure.

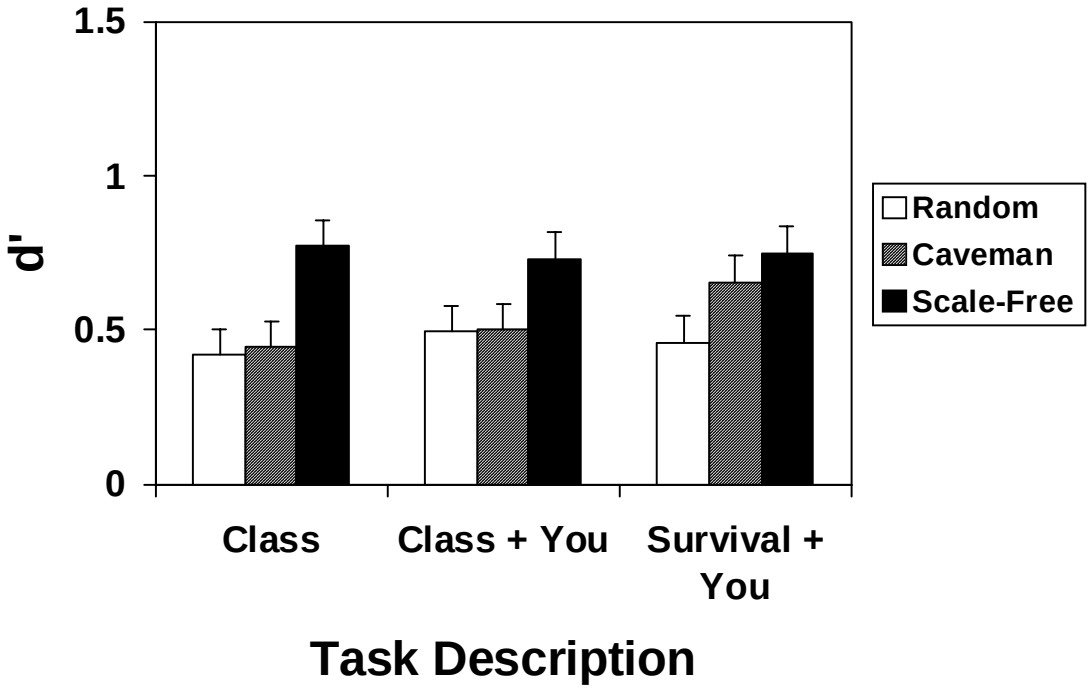

Figure 9. d' in each combination of experimental conditions. Graph Structure varied within-participants. Task Description varied between-participants.

Egocentric Bias. Figure 10 compares d' in test trials in which "You" was a component of the test edge to test trials that did not involve the "You" node. Overall, participants displayed greater discrimination when a test included the "You" node $t(89) = 8.28$, $p < .001$. The size of the effect does not differ significantly as a function of graph structure or task description.

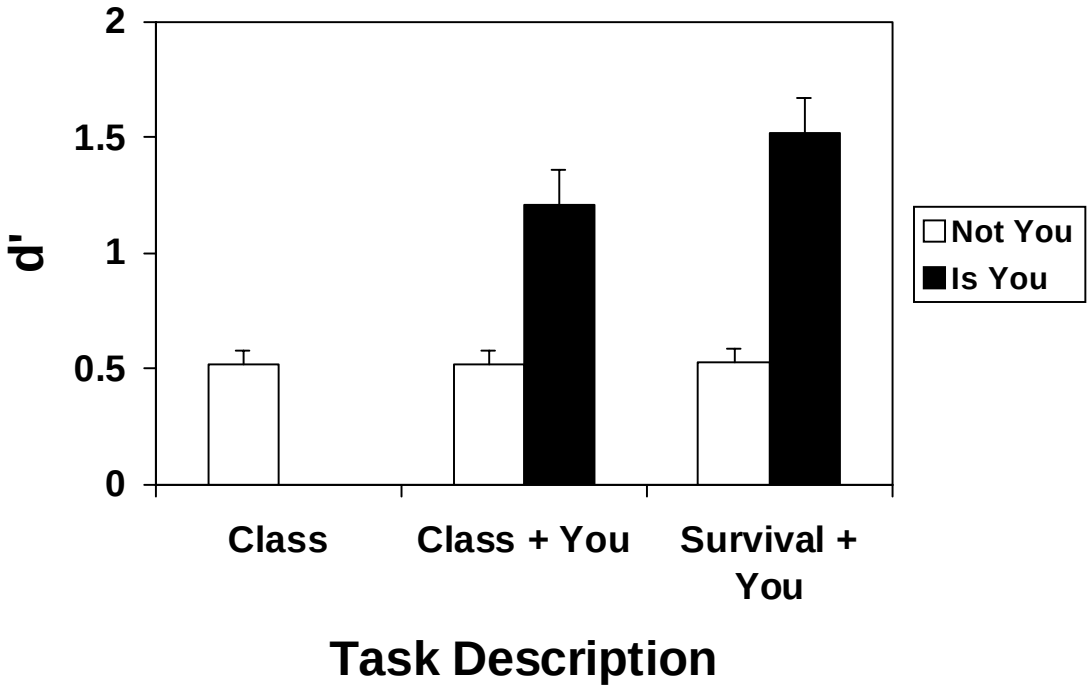

Figure 10. d' for test trials that did or did not include the "You" node, separated by Task Description. The Class condition did not include the "You" node.

Network Centrality. Figure 11 plots response accuracy as a function of the mean network centrality of the nodes in the test edge for random and scale-free graphs. A U-shaped curve is apparent in each plot. Note that more test trials fall in the central area of both plots, because a combination of two nodes in the same extreme degree range was more infrequent than a combination of two nodes both in the medium range of degree values or with opposite extreme scores. Points with fewer than 50 observations were not plotted.

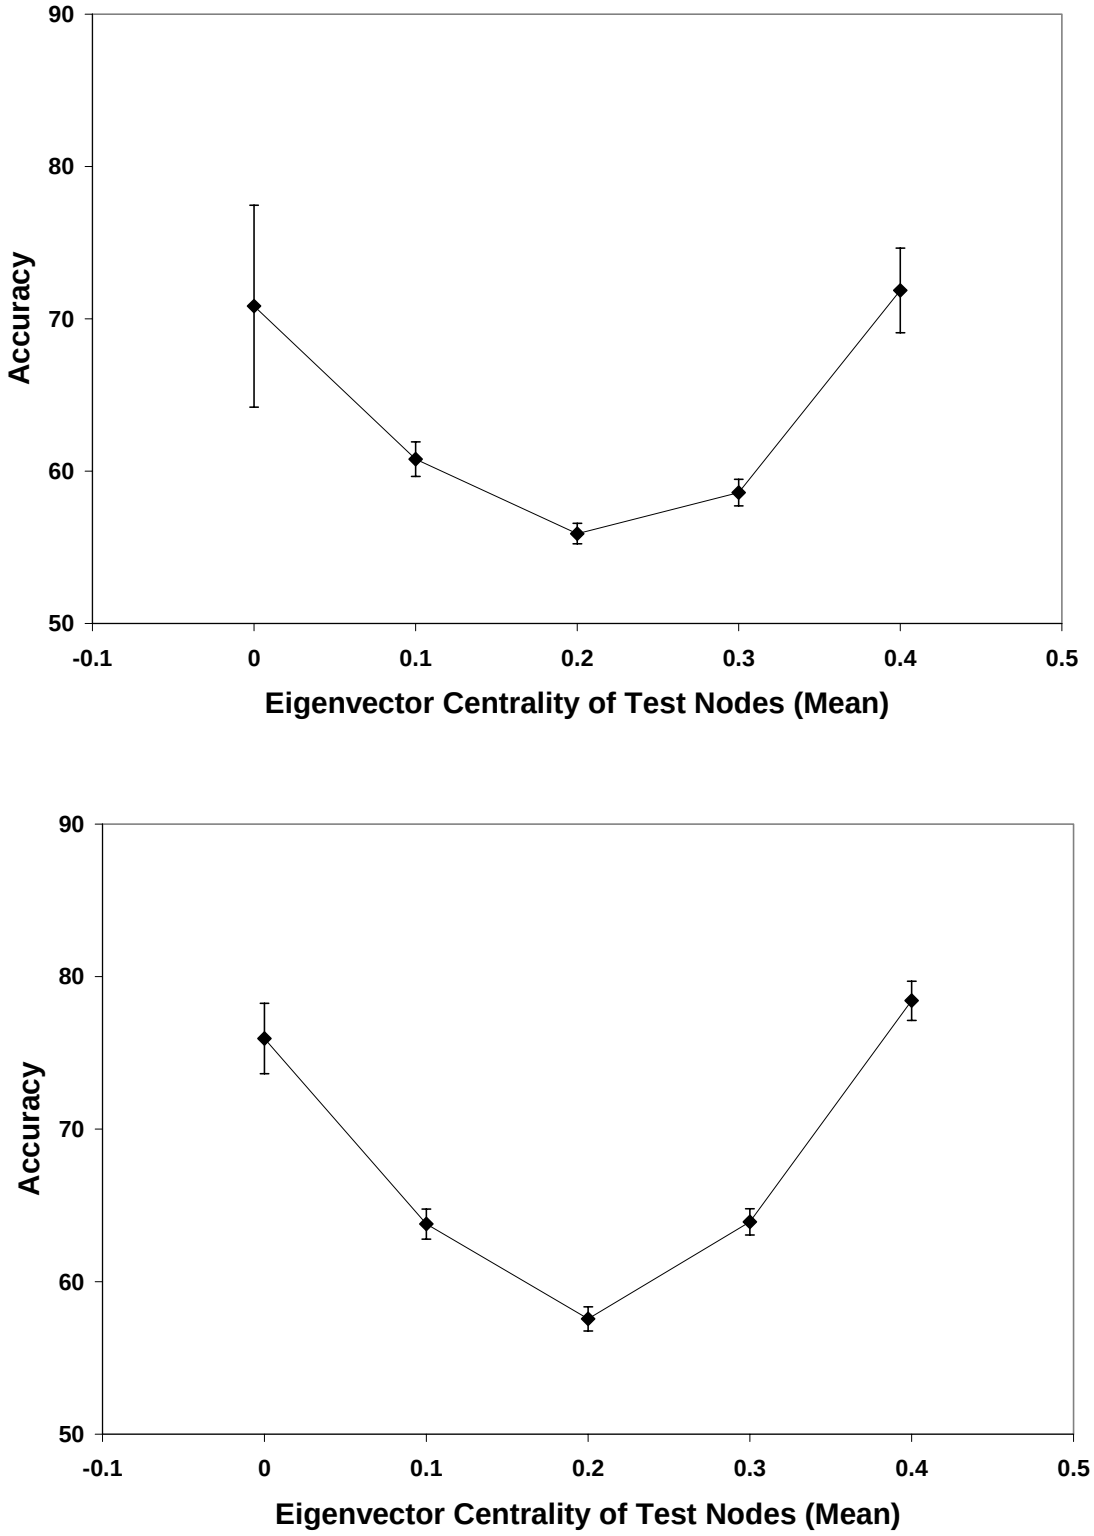

Figure 11. Top: Accuracy as a function of the mean eigenvector centrality of the two nodes in a test trial for the Random Graph condition. Bottom: Accuracy as a function of the mean eigenvector centrality of the two nodes in a test trial for the Scale-Free Graph condition.

Examining sensitivity and bias leads to a deeper understanding of this pattern of results. Figures 12-15 plot hit rate, false alarm rate, d' and c as a function of node centrality for both random and scale-free graphs. The patterns are similar qualitatively across graph structures. Subjects' strategies do not differ between graph structures.

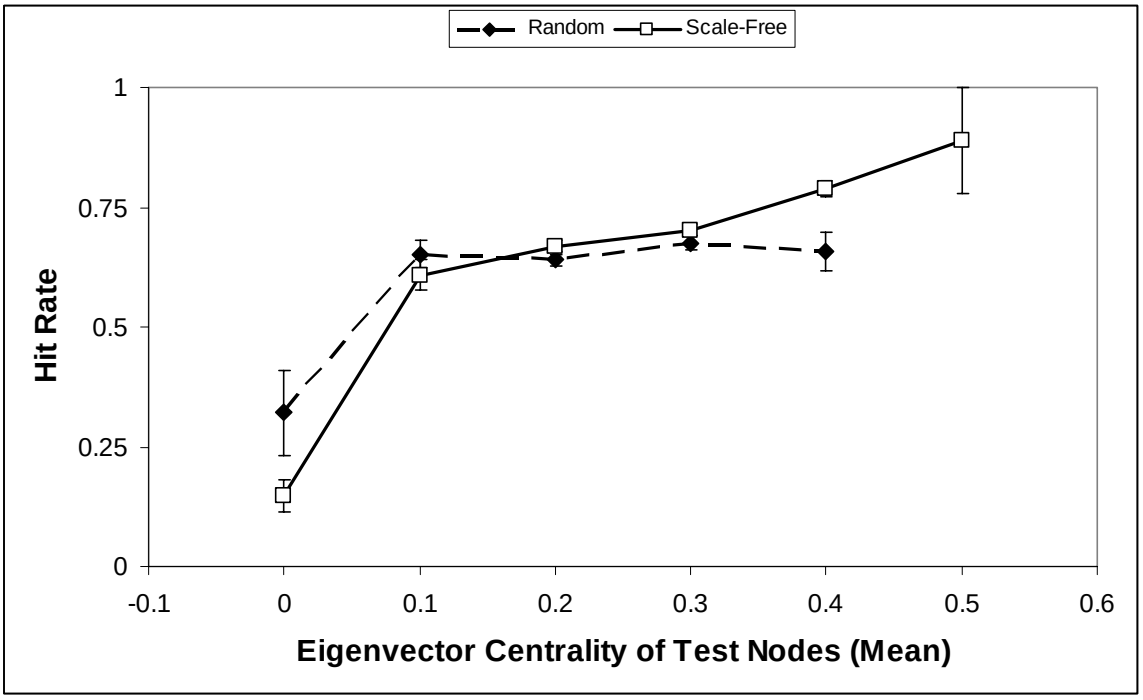

Figure 12. Hit rate as a function of test node centrality.

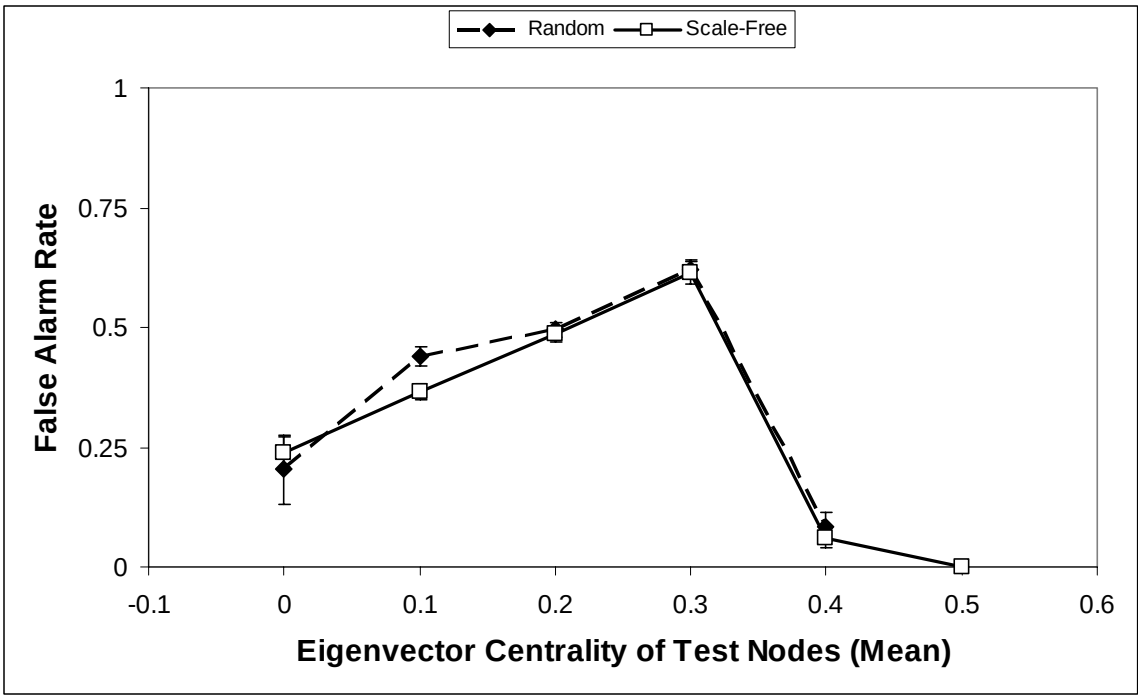

Figure 13. False alarm rate as a function of test node centrality.

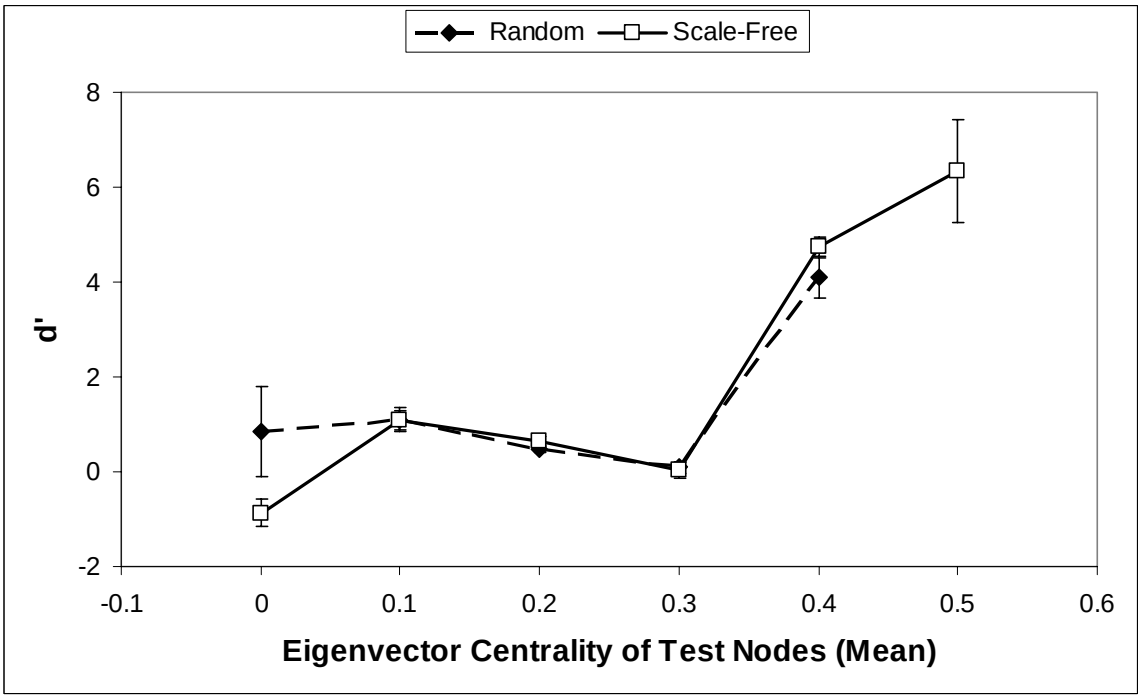

Figure 14. d' as a function of test node centrality.

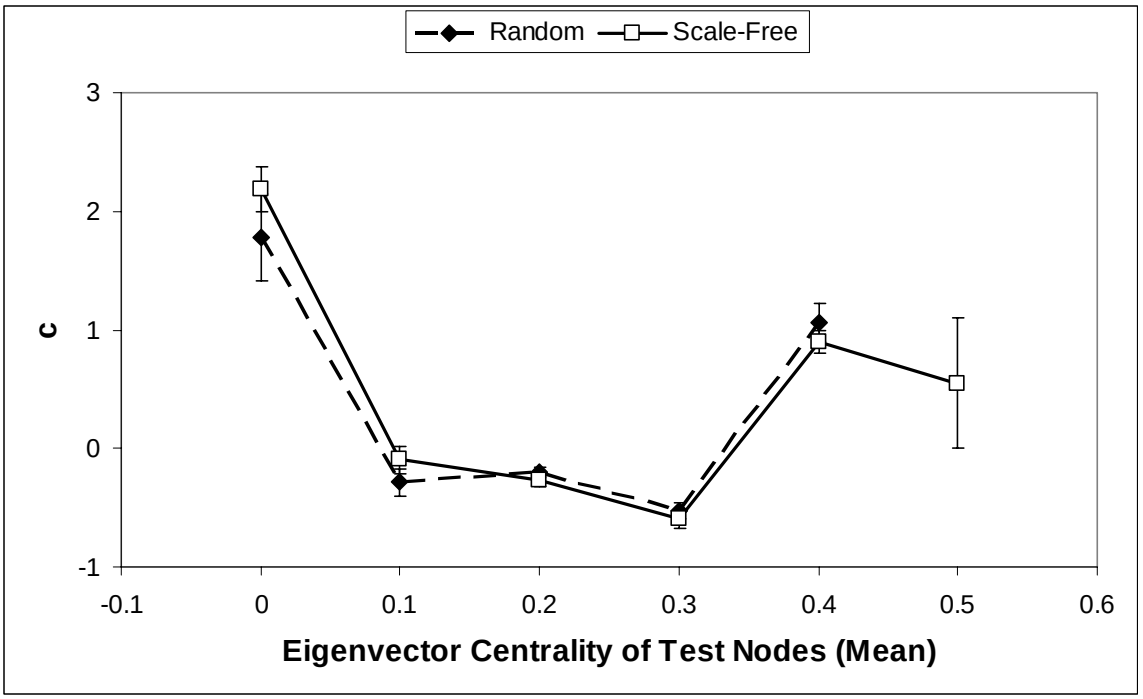

Figure 15. c as a function of test node centrality.

Contrast Figure 14 with Figure 11. With this new analysis, it appears that two different processes are at work at the ends of the centrality spectrum. Actual discriminability is low for pairs of nodes of very low centrality (those less than .10). The high values of c at this point in the graph indicate a strategic choice by subjects to always favor a No response when the edge in question is purported to be between two peripheral nodes. Subjects achieve good accuracy in these cases by betting on a No response (a rational choice). They do not, on average, possess good knowledge about which peripheral nodes are connected.

On the other end of the spectrum, however, discriminability is very high. Subjects' best performance comes when they must judge whether two nodes of high centrality are connected. The low false alarm rate shows that they do not do this by just

responding Yes. In fact, the value of c indicates that subjects become conservative when confronted with central nodes. Their better knowledge about these nodes overcomes this bias, however, and leads to the high accuracy previously observed.

Just like the accuracy results, the discriminability results show that subjects have difficulty retaining knowledge concerning nodes of middling centrality. The values of d' closest to zero occur at eigenvector centrality scores between .20 and .30. To use a social network as a metaphor, the results imply that we assume unpopular people have no friends at all, we know in fine-grained detail who likes whom in the well-connected set, and those in the middle exist in a grey haze of possible connections.

(Another explanation for the apparently changing values of d' and c is that the equal-variance signal detection model is not appropriate to describe subjects' judgments in these tasks. It is possible an unequal variance model of memory (Jang, Wixted & Huber, 2009) would be more appropriate.)

Discussion

Scale-free graph structure is acquired more quickly than caveman or random graph structure. A slight, marginally significant advantage for caveman over random graph structure also appeared in the data. Survival relevance and personal relevance manipulations of the task description failed to produce significant differences in performance. However, there is evidence participants do pay extra attention to the node labeled "You" when it is present. Participants do well when asked about edges that involve nodes of extreme centrality. They develop a sensitivity to the fact that central nodes have many links and peripheral nodes have few links.

The task of learning a network among 20 nodes is a difficult one. Participants do not know how many edges to expect in the graph. There are 190 edges that might exist, so in actuality the participants must learn which one of 2^{190} possible graphs they are confronted with. Despite this level of difficulty, participants do manage to acquire knowledge about who is friends with whom in the several minutes of training and testing devoted to each graph. In the Random Graph condition, in which performance is lowest, participants still achieve an overall accuracy of 58%. This is better than the performance one would expect from uninformed guessing $t(134) = 10.14$, $p < .001$. It is an impressive feat given the scale of the possible-graph space and the fact that random graphs (by design) have no structure.

Caveman graphs, on the other hand, have a well-defined structure. Participants achieve slightly better acquisition of caveman graphs. On average, participants respond correctly to 1.1 more test edges when learning a caveman rather than a random graph. Two properties of caveman graphs work against participants' expectations. The first is the relative homogeneity of degree. Unlike in real social networks or in constructed scale-free networks, in a caveman social graph all individuals have nearly the same number of friends. Thus, learners are not able to use centrality heuristics to pick out popular and unpopular nodes for easy classification.

The second difficulty presented by a caveman graph is the intransitivity introduced during construction of the graph. In this process, one node in each cave breaks an intracave connection, disrupting the pattern of complete connectivity within the cave. The false alarm rate for broken intracave test edges is 61% (where false alarm rate is defined as number of false positives divided by the sum of false positives and true

negatives). Participants make a reasonable assumption of transitivity, but are punished by this peculiarity of the caveman graph.

Scale-free graphs elicited the best acquisition. This finding has been replicated in similar work in which scale-free graph structure was acquired more quickly than random graph structure (Chapter 2). In those studies, acquisition of scale-free graphs was found to be superior whether the graphs were described as social, computer or transportation networks. Clearly, scale-free graphs are more easily acquired. Why this is the case is not as clear, however, and alternative explanations are discussed later in this article.

Unlike graph structure, no significant effects due to task description surfaced. Accuracies in the Class, Class + You and Survival + You conditions were very similar: 60%, 60%, and 61% respectively. Of course, we cannot simply accept the null hypothesis. However, it can be pointed out the experiment was not underpowered. To detect a 5 percentage point difference (the same size as the difference between the Scale-Free Graph and Random Graph conditions) between any two task description conditions, a two-sample t-test had power greater than .80.

Broadly, there are two explanations for the lack of an effect. First, that no effect exists, and second, that the current treatment was too weak to evoke the true effect. In the first case, one would argue that learners successfully abstract away the superfluous personal relevance and survival relevance narrative framing of the instructions. The task is performed by creating associations between individuals, and this is true whether one is instructed to behave as a sociologist, a student or a survivor. The presence of egocentric bias is evidence against learners' ability to completely abstract away irrelevant details,

however. There is no advantage to learning best about the local subgraph around the "You" node, but participants do so anyway.

It may instead be the case that the personal relevance treatment was too weak. In the stroke study discussed in the Introduction (McDonald, et. al., 2009), the personal relevance treatment implied that the participant could use the information learned to actually *save the life* of a loved one. In the current study, the personal involvement was only hypothetical.

It may be too ambitious to ask participants to adopt a novel and strange network as personally relevant in the brief space of an experimental session. Surely, of all possible social networks, people have the best information about their own – the network centered around themselves. However, it may be interesting to compare participants' knowledge of their own personal social networks with their knowledge of those of celebrities, reality show cast members and fictional characters.

Like personal relevance, the survival relevance manipulation failed to produce an effect. It should be noted, however, that the protocol in this study differed in several ways from that used in successful demonstrations of the survival processing effect. Foremost is that participants were not required to evaluate each stimulus in terms of survival relevance. It is also the case that the current task (classify each test edge as a friendship or not a friendship) is not the same as the task in demonstrations of survival processing, which have so far involved straightforward recognition or recall.

Merely describing the task of social network learning in terms of survival did not produce an effect. Future work could present edges one at a time, ask subjects to rate each on a survival relevance scale and require participants to recall or recognize the edges

at test in order to more closely replicate previous successful survival processing protocols.

The inclusion of a node labeled "You" did not have a global effect on accuracy, but it did produce a reliable local effect. Participants displayed an undeniable egocentric bias. Accuracy in trials that included the "You" node was 10 percentage points higher than in trials that did not. Of the 90 participants in conditions that included "You" nodes, 70 consistently scored higher in You trials as compared to other trials.

The egocentric bias is interesting for two reasons. The first involves the literal interpretation of the effect. It would appear that people are more interested in their own relationships than those of others. Even in this hypothetical-universe social network (in which "You" is just as arbitrary a label for a node as the rest of the list of names) the self draws outsized attention. Gilovich, Medvec and Savitsky (2000) have discussed the "spotlight effect" in reference to the belief that others notice us more than they really do. The self cannot get past the outsized salience the self has to the self. This is true even at very early stages of attention. It has been demonstrated that the attentional blink and repetition blindness are both attenuated when the probe stimulus is the subject's own name (Arnell, Shapiro, & Sorensen, 1999; Shapiro, Caldwell, & Sorensen, 1997).

In the current study we see evidence of self-salience yet again. Participants bias their attention and learning toward their own connections. Combined with the spotlight effect, the current results predict people should overestimate the accuracy with which their acquaintances could identify their close friends. Another prediction given the current results would be that knowledge of the structure of one's own social network will drop off severely as the degree from the self increases. One of the reasons for online

social networking's success may be that the user has to manage only those connections directly incident on the self, and knowledge of the extended network is provided as a side effect.

The second interesting aspect of egocentric bias is a more general demonstration of a graph learning phenomenon. Namely, egocentric bias illustrates it is possible to learn local subgraphs to a greater degree than the graph as a whole. It is useful to know that a node can be highlighted during training in such a way that the learner will become especially aware of its connections.

The presence of egocentric bias predicts more generally that an exceptionally salient node will cause uneven learning of the graph structure. To confirm this finding, this prediction could be tested in a number of ways. For instance, one node could be made perceptionally salient by presenting the name in a different color than other names. Or the node could be made semantically salient – by presenting one female name among many male names, for instance.

Examining the effects of nodes themselves on learning also uncovered the effects of network centrality. The graphs in Figures 11 tell a clear story. Specifically, a "popularity position effect" similar to the serial position effect in verbal list learning is apparent. Participants answer correctly more often when they are tested on edges consisting of two nodes both of low centrality or both of high centrality.

This evidence suggests that in addition to having some true knowledge of the social graph structure, learners also employ a simple heuristic: answer no for peripheral nodes and yes for central nodes. Let us choose one set of data to illustrate this point. When participants are learning scale-free graphs, they respond YES 46% of the time

when the test edge has mean node degree less than the median value, and they respond YES 69% of the time when the test edge has mean node degree greater than the median value. Clearly, subjects adjust their propensity to respond yes or no as they rationally should in accordance with node degree.

It is easy to distinguish central nodes from peripheral nodes during training. Central nodes have long lists of connections when they are the trained node, and peripheral nodes have short lists. Central nodes appear often in other nodes' lists of connections, and peripheral nodes do not.

One of the advantages learners have in confronting a scale-free graph is this graph structure's greater variance in node centrality. Learners thus have greater opportunity to employ the network centrality heuristic. Undoubtedly, this contributes to the greater acquisition of scale-free graphs. Further evidence that subjects employ a heuristic based on a node's centrality will be provided by the formal modeling of subjects' learning in Chapter 5.

It is important to acknowledge the limitations of the current study before drawing final conclusions. It may be premature to wholeheartedly endorse the anthropomorphic acquisition hypothesis. As predicted, scale-free graphs were acquired most quickly. The advantage for caveman graphs over random graphs was small and only marginally significant, however. It is clear that graph structure can affect acquisition. The exact reasons and mechanisms are still up for debate. The evidence so far is consistent with the idea that the more a social graph stimulus comports with the properties of true social networks, the easier it is to learn. The ultimate test (and a target for future research) would be to use real-world social networks as training data. Are the graph motifs most

common in real social graphs the most easily learned graph motifs? Do learners respond best to real-world levels of graph properties (moderate transitivity, for example) or is learning optimized by maximal, unrealistic values?

The current study found that describing the social graph learning task as personally relevant or survival relevant had no discernable effect on acquisition. As discussed previously, the manipulations may have been too weak or it may be that the effect does not translate to this domain or this training and testing paradigm.

The field typically frowns on the publication of null results, because they are supposedly uninformative. In cases like this study, however, where the experiment does not lack power and there were legitimate grounds to expect the manipulation to produce an effect, it is useful to see the results. If nothing else, the current study begins to bound the parameters that allow for an effective relevance manipulation.

The author does not expect this work to be the final word on social network acquisition. Indeed, given the current level of interest in social networks, acquisition is a neglected area of study abounding in interesting questions. Which individual differences measured *a priori* could account for performance on this task? Is social graph acquisition more a function of general intelligence or is there some skill orthogonal to *g* that allows for smooth navigation of the social landscape? Of the many ways graphs can differ (in density, transitivity, connectedness, efficiency, etc.) which predict the difficulty of human acquisition? There is a great deal of work in computer science attempting to automate the learning and prediction of link patterns in networks (e.g. Menon & Elkan, 2011). Are there ways in which humans can outperform machines? The answers to these questions await the attention of social cognition scholars.

Chapter 4: Interactions of Graph Structure and Training Regimen: How Best to Teach Differing Graphs

4.1 Introduction and Experiment 1

Network graphs are complex structures. The number of possible configurations of a graph grows exponentially as a function of the nodes in the graph. The relation between graph size and complexity is graphed in Figure 16. A graph of two nodes may exist in only one of two states: connected or not connected. Three nodes allow for three possible edges, each of which may exist or not. Already the space has grown from two potential graphs to eight. Increasing the size of the graph by just one more node yields six possible edges and sixty-four possible graph configurations.

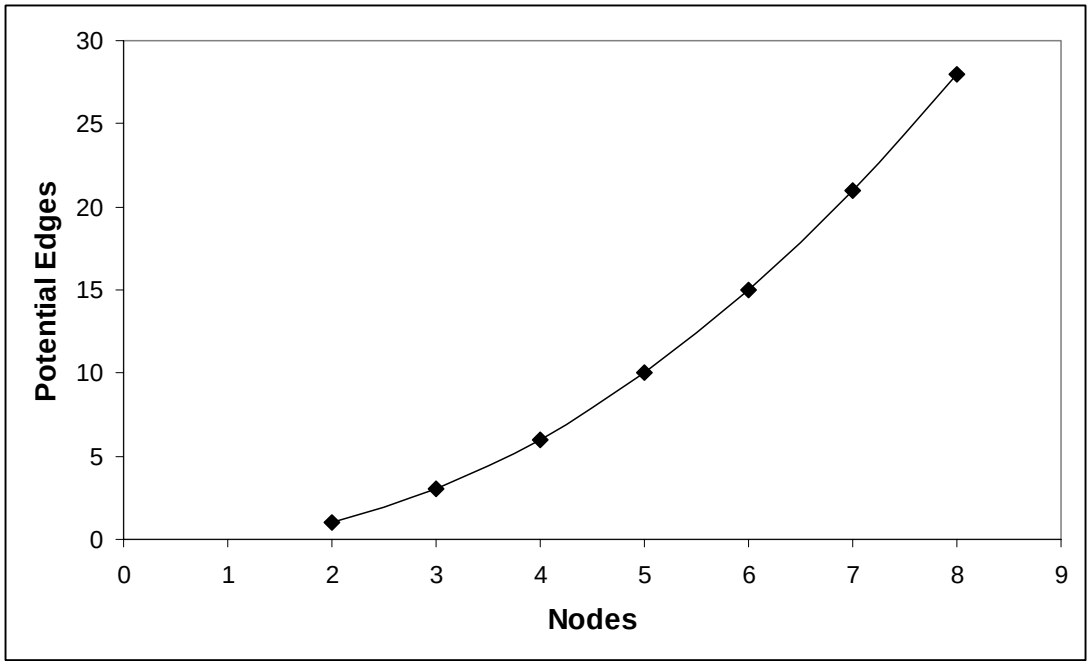

Figure 16. The number of potential edges in a graph as a function of the number of nodes. The value increases as (N choose 2). Note that the number of possible graphs over (N choose 2) edges is $2^{(N \text{ choose } 2)}$. This function rises so steeply, there is nothing to be gained in graphing it, except in log values on the y-axis, which (with a base of 2) would yield the present graph.

In the studies to follow, the stimuli are relatively small graphs of eight nodes. Yet, in each case, the demand placed on subjects is to learn one of 2^{28} or 268,435,456 potential graphs. Given this teeming forest of possibilities, what is the most efficient path to knowledge of the one true structure of a graph?

A simple answer to this question would be satisfying. Alas, as I will demonstrate, the best method for teaching a network graph structure depends upon the structure itself. My previous work (Chapters 2 and 3) has made it clear that graph structure is a determinant of the pace at which knowledge about a graph will be acquired. In this work it will be shown that acquisition rate is affected also by the interaction of structure and training. The first distinction of structure investigated was that between regular and irregular graphs.

Regular and Irregular Graphs

A regular graph is defined as a graph in which a simple pattern of connectivity is repeated again and again at every node. When diagrammed, such graphs form geometric patterns such as a lattice or grid. The ring lattice – depicted in Figure 17 – was chosen as a regular graph structure for this experiment.

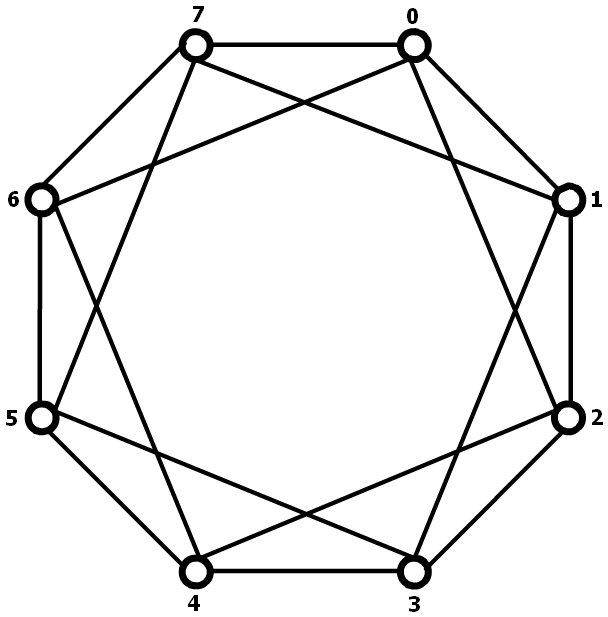

Figure 17. A ring lattice graph of 8 nodes.

An irregular graph is any graph in which the pattern of connectivity is unpredictable from node to node. The slightest perturbation of a regular graph will produce an irregular graph. Figure 18 depicts the irregular graph stimulus to be used in this experiment. Close inspection will confirm that this graph is the ring lattice of Figure 14 after two edges have been "rewired." The edge 7-0 has become 7-2 and the edge 1-3 has become 1-5.

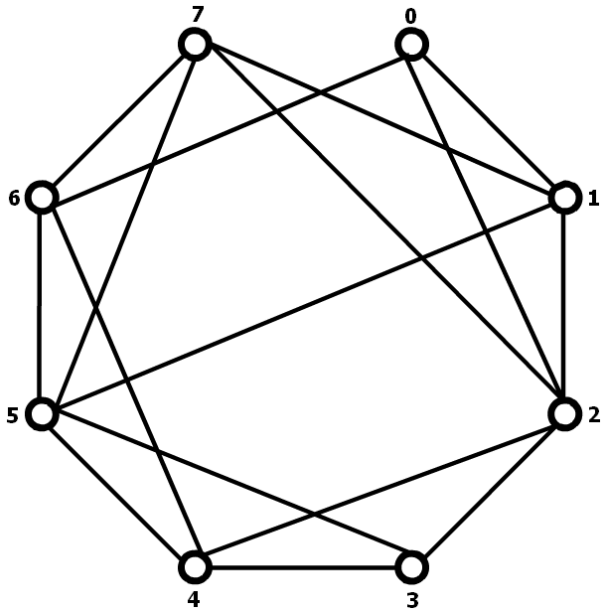

Figure 18. An irregular graph formed by rewiring two edges of the ring lattice.

This particular distinction between a regular graph and irregular graph has been important in the development of network theory and especially in the domain of social networks (Watts, 1999; Watts & Strogatz, 1998). As Watts and Strogatz observed, ring lattice graphs possess several characteristics that make them good models of social networks.

Most importantly, the ring lattice is highly clustered. Connected nodes have many edge partners in common. In a social network, this is the equivalent of having friends who are also your friend's friends. This clustering reflects the propensity of true social networks to display high transitivity and to contain highly-connected cliques. However, a perfectly regular ring lattice is *too* highly clustered to accurately represent true network structures.

Ring lattice graphs have a high diameter – the greatest distance between any two nodes. Here, distance is defined as the number of intermediate nodes one must visit in the shortest path that follows valid edges from one node to another. Ring lattice graphs have a high diameter, because nodes are only connected to their neighbors. It takes many edge traversals through intermediate nodes to reach the opposite side of the lattice.

True networks are often "small-world" networks. As Watts and Strogatz demonstrated with data on the power grid of the western United States, the collaboration graph of film actors, and the neural network of the worm *Caenorhabditis elegans*, these networks are highly clustered and at the same time of low diameter.

Watts and Strogatz provided a model to transform regular graphs into more realistic representations of small-world networks. They introduced the parameter β which takes as its value the probability of rewiring each edge in a graph. Rewiring an edge meant keeping one node (think of this as the source) fixed and choosing a second node (the destination) at random. At $\beta = 0$, a ring lattice would be unchanged. At $\beta = 1$, the lattice would be completely transformed into a random graph. At low values of β, clustering remains high, while diameter falls quickly, creating a small-world network.

The graph in Figure 18 can be considered the result of applying the Watts and Strogatz process to the ring lattice in Figure 17 with $\beta = .125$.

No work has yet explored whether regular or irregular graphs are more easily learned. The predictability of the patterns of connectivity in a regular graph may make these graphs easier to represent in memory. Experiments in visual perception and memory have shown that the brain is biased toward symmetric representations (Freyd & Tversky, 1984; Tversky & Schiano, 1989). Near-symmetric figures are distorted toward

symmetry in memory, but already symmetric figures are not distorted. This work would suggest that regular graphs will be represented with fewer errors in memory. This should be especially true for graphs represented as diagrams. The symmetric or asymmetric nature of the graph will be visually salient in this case.

On the other hand, irregular graphs are a more realistic representation of actual networks, and may therefore be more familiar and easily remembered. An irregular graph is more likely to match one's expectations about the structure of a true social network. In addition, regular graphs have no variability in the number of edges each node participates in. (This is known as the node's degree.) Irregular graphs do have degree variability. In previous experiments, degree variability has correlated positively with acquisition and has been a good predictor of performance. These two factors predict an advantage in learning an irregular graph.

The structure of a graph is one factor likely to affect the rate at which knowledge about it is acquired. Another is the type of training employed to impart that knowledge.

Diagram or Verbal Training

When one works with network data, two forms of representation quickly become familiar: the edge list and the network diagram.

An edge list is the list of dyads (node pairs) that are connected in a graph. The presence of a dyad in the list means the two nodes are connected, and the absence of a dyad means the nodes are not connected. An edge list is a compact representation for sparsely connected networks. In a well-organized edge list, one can quickly look up a particular node and examine all its connections.

A network diagram is a visual representation of a graph. Each node is represented with a circle, and lines connect the circles representing nodes that are connected by edges in the graph. Network diagrams are also somewhat compact representations of sparsely connected graphs, because only existing edges are drawn. One can determine all of the edges a node participates in by finding the appropriate circle and examining all of the lines connected to it. Crossing lines can sometimes make this process difficult, however. One potential benefit of the network diagram is the opportunity to use the spatial layout of the nodes to provide information that must be inferred from edge lists. For example, one can place nodes that are connected and that share other connection partners closer together to represent the metaphorical "closeness" of those nodes.

Given these two common network graph representations, two methods of teaching novel graph structures suggest themselves. Presenting edge lists constituted Verbal training. Names of connected nodes were shown together, as they would appear in an edge list. Presenting network diagrams constituted Diagram training. Circles labeled with node names were connected by lines representing edges.

The aphorism "A picture is worth 1000 words" succinctly paraphrases the hypothesis under test in regard to training. It is often assumed that diagrams make better learning aids than verbal descriptions, but it has proven difficult to demonstrate (Cheng, 2002; Larkin & Simon, 1987). Cheng developed a system for diagramming electrical circuits that represented the relationships between current, voltage, resistance and power with geometric features such as line lengths and rectangle areas. He compared training utilizing diagrammatic representations of circuits to conventional training focusing on the equations that relate theses concepts (e.g. Ohm's Law: $V = Ir$). On most measures,

performance was similar for those taught with diagrams and those given conventional training. However, diagram training conferred reliable advantages on complex problems after a 5-day retention interval. Cheng argued this result was due to diagram students attaining a more coherent understanding of the laws of electricity.

In summary, this experiment sought to address how the characteristics of network training and networks themselves affect acquisition. Regular and irregular graphs were trained through both verbal descriptions and diagrammatic depictions. Previous work weakly suggests that regular graphs should be easier to acquire and diagrams will make better teachers than verbal lists.

Method

Participants

172 UCSD undergraduates participated in the study.

Design

The design was a 2 (Graph Structure) x 2 (Training Regimen) within-subject design. All subjects were trained on one graph of each type (Ring Lattice and Watts-Strogatz) once under each training regimen (Diagram or Verbal). The order in which the four conditions were presented was chosen randomly for each subject, under the constraint that the number of subjects in each ordering should be as near equal as possible.

Stimuli

All graphs consisted of 8 nodes. Edges were undirected and unweighted. Self-edge loops were not allowed. In each graph, there were 28 potential edges (i.e. 8 choose 2). Participants were trained on two graph structures: Ring Lattice and Watts-Strogatz.

Graph Structure. Figure 17 depicts the Ring Lattice graph. The Ring Lattice graph had a regular structure. Every node participated in four edges. The nodes could be arranged in a ring such that each node was connected to its two nearest neighbors and its two second-nearest neighbors and no other nodes.

The Watts-Strogatz graph was identical to the Ring Lattice graph except for two perturbations. Two of the sixteen edges in the Ring Lattice graph were "rewired." Two nodes lost connections and thus participated in only three edges, and two nodes gained connections and thus participated in five edges. As can be seen in Figure 18, it remains likely that neighbors are connected, but some neighbor connections are missing and new long-distance connections have taken their place.

Name Lists. Four lists containing four male names and four female names were constructed by choosing names from the Social Security Administration's list of popular baby names in the 1990's (Social Security Administration, 2011). The initial letter of each name was unique within each list. At the beginning of each block, the stimuli names were randomly mapped to nodes in the abstract graph structure.

Procedure

The experiment was administered online through a PHP/mySQL/Flash Web application. Subjects followed a link to the experiment website from the subject pool administration software.

Subjects first read through several short paragraphs of instructions. They were told they would be learning "who is friends with whom" in four different groups of eight people.

Participants were told to expect one of two types of training for each group of people. The instructions stated:

> *Your job is to learn who is friends with whom. To teach you, you will be shown either diagrams or pairs. For pairs, you will see the names of two people who are friends. For diagrams, you will see 8 circles with a name by each one. If two circles are connected by a line, it means those two people are friends.*

Subjects were informed that all links were reciprocal – if A was connected to B, then B was also connected to A. Subjects were also instructed to expect test trials in which two stimuli would be presented and they were to answer YES the people were friends or NO they were not. Subjects were instructed to give their best guess if they did not know the answer.

In the experiment, training trials alternated with test trials. There were 4 blocks of 64 training and test trials each. Subjects completed one block of training and test trials for each combination of Graph Structure and Training Regimen. Each subject received a randomly chosen permutation of conditions, with the constraint that the number of subjects in each permutation remained as equal as possible.

Diagram Training Trials. In a diagram training trial, a visual representation of the graph to be learned was presented to the subject for 10 seconds. The diagrams appeared as depicted in Figures 17 and 18, with the addition of a name next to each node. The words "Friends are connected by lines." appeared above the diagram.

Verbal Training Trials. In a verbal training trial, one node was selected to be the focus of training. Nodes were sampled without replacement from the list of 8 nodes until

all had been selected. This process then repeated. Thus, over the course of 64 trials, each node was the focus of training exactly 8 times.

All of the edges that involved the node in focus were presented to the participant. An edge presentation began with a blank, white screen presented for 500 milliseconds. Next, the name of the training node appeared alone for 500 milliseconds. The words "is friends with" appeared to the right of the training node, and 500 milliseconds later, the name of a node connected to the training node. The complete stimulus (e.g. *Amanda is friends with Benjamin*) remained onscreen for 1 second. This process repeated for every edge involving the node to be trained.

Test Trials. In a test trial, two names appeared next to each other. Below the names was the prompt "Friends?". Below the prompt were two equally-sized light-grey buttons labeled YES and NO.

Feedback was provided after the subject made a response. The button for the correct answer was highlighted, and a 300 millisecond sound clip of a bell was played for a correct response or a buzzer for an incorrect response.

For each graph, subjects were tested on an equal number of valid (existing) and invalid (non-existent) test edges. The set of valid edges and (separately) the set of invalid edges were randomly sampled without replacement until all in a set had been shown. At that point, all of the edges in the set became available again for re-testing. This procedure was followed so that responses should be distributed evenly between YES and NO. In each block, there were 32 valid test edge trials and 32 invalid test edge trials.

Results

d' for each combination of conditions is graphed in Figure 19. In all conditions, d' was statistically significantly greater than the value of zero expected from uninformed guessing. A repeated-measures ANOVA with Greenhouse-Geisser correction was performed to test for effects of graph structure and training regimen. The main effect of Graph Structure was significant $F(1, 171) = 7.60$, $MSE = .39$, $p = .006$. Subjects' d' in Ring Lattice blocks was 1.04 and was reliably higher than the .92 d' in Watts-Strogatz blocks.

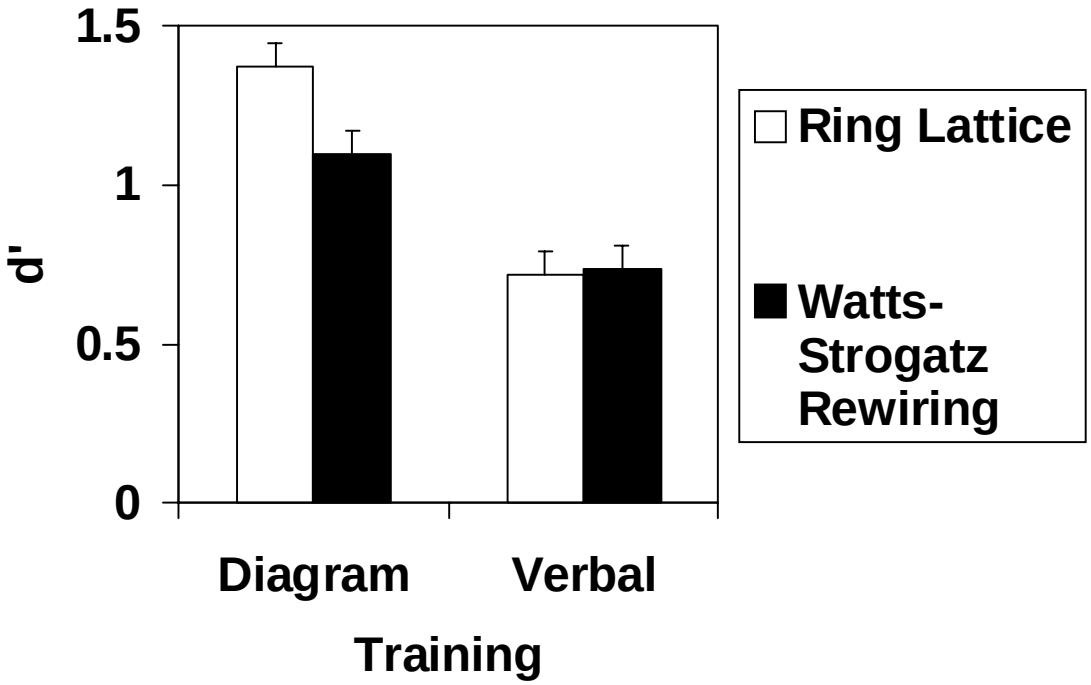

Figure 19. d' in each condition in Experiment 1. The main effect of graph structure, the main effect of training and the interaction of training and graph structure are significant.

The main effect of Training Regimen was significant $F(1, 171) = 59.50$, $MSE = ..74$, $p < .001$. Subjects' d' in Diagram training blocks was 1.23 and was reliably higher than the .73 d' in Verbal training blocks.

The interaction between Graph Structure and Training Regimen was significant $F(1, 171) = 10.08$, $MSE = .37$, $p = .002$. Paired-sample t-tests reveal that the d' advantage conferred by Diagram training is greater for Ring Lattice graphs as opposed to Watts-Strogatz graphs. Under Verbal training, no reliable difference between graph structures is observed.

Discussion

For both graphs, diagram training led to better acquisition than verbal training. A regular, ring-lattice graph was easier to learn than the same graph slightly perturbed into an irregular form. However, this was only the case under one form of training. Verbal training leads to poorer performance no matter the structure of the graph. In all cases, there is learning. Participants do learn to discriminate existing friendship edges from non-edges.

The results of this experiment strongly suggest a practical recommendation accompanied by a theoretically interesting caveat. The practical recommendation is to represent a network graph with a diagram whenever possible if the goal is to help others understand its structure. However, it is necessary to take care in constructing the proper diagram.

In their theoretical treatise on the utility of diagrams in problem-solving, Larkin and Simon (1987) were adamant that it is not enough to create just any arbitrary diagram. Instead, they argued diagrams provide advantages by following three conventions:

- *Diagrams group together all information that is used together, thus avoiding large amounts of search for the elements needed to make a problem-solving inference.*
- *Diagrams typically use location to group information about a single element, avoiding the need to match symbolic labels.*
- *Diagrams automatically support a large number of perceptual inferences, which are extremely easy for humans.*

Network diagrams meet all three conditions. They group together all information into one figure. In fact, all nodes and edges must be present or else a network diagram would provide misleading information. Network diagrams use location (circles) to group information about elements (nodes). Although the difference may seem trivial at first glance, the results of this experiment imply that a node labeled Alice with lines reaching out to Bob and Carol is a more effective representation than the two statements "Alice is friends with Bob" and "Alice is friends with Carol." Network diagrams support perceptual inferences, if care is taken creating the diagrams. In the experimental stimuli, nodes placed near each other were likely to be friends. In addition to examining the lines connecting node markers, simple proximity of the markers could be used as a heuristic for inferring whether nodes were connected by an edge. For Watts-Strogatz graphs, proximity was not the perfect cue it was for Ring Lattice graphs and performance in the Diagram condition suffered.

The observation that learning a Watts-Strogatz graph benefits less from a diagram than a Ring Lattice graph leads to an important caveat. Though the Ring Lattice graph was the "regular" graph and the Watts-Strogatz graph the "irregular" graph in this study,

the Watts-Strogatz graph is still mostly regular. The value of β for the Watts-Strogatz graph was only .125. Thus, 87.5% of the edges are the same as those in a perfectly regular ring lattice. This difference was enough to create a reliable decline in the advantage of a diagram over a verbal description. The advantage is likely to continue to decline as β increases. At $\beta = 1$, the graph would be completely randomly determined, and it is possible the benefits of a diagram would no longer apply and the advantage over a verbal description could be lost. If one imagines the spaghetti-like mess of a diagram for a random graph, a diagram might even provoke a disadvantage over a verbal description. These predictions should be tested empirically.

The previous caveat might be reframed as an admonition to shape a diagram to reflect the network. The arrangement of nodes in the experimental stimuli was perfect for the Ring Lattice graph, but not for the Watts-Strogatz graph. A reshaping that reflected the more central position of the nodes that gained edges in the Watts-Strogatz perturbation may have lessened the decrease in the efficacy of diagram training. There are many methods to shape network diagrams, and one should choose the method that best represents the underlying structure of the graph.

As a case in point, it is interesting to examine a particular set of errors that subjects made while acquiring Watts-Strogatz graph structure. There are four errors that could be called "symmetry errors," because they imply the graph was perceived or remembered as more symmetrical than it actually was. The four edges marked with broken grey lines in Figure 20 depict these potential errors. In two cases, the error involves a false-alarm response, i.e. the subject mistakenly reports an edge exists when it does not. In the other two cases, the error involves a miss response, i.e. the subject

mistakenly reports an edge does not exist when it does. These errors occurred in roughly equal numbers in Verbal and Diagram training conditions. Keep in mind, however, that subjects made reliably fewer errors in Diagram training. The *proportion* of errors that are symmetry errors was higher in the Diagram condition (17.8%) than it was in the Verbal condition (15.6%). This difference can be examined within-subject, and a paired-sample t-test approached significance $\underline{t}(171) = 1.91$, $\underline{p} = .057$.

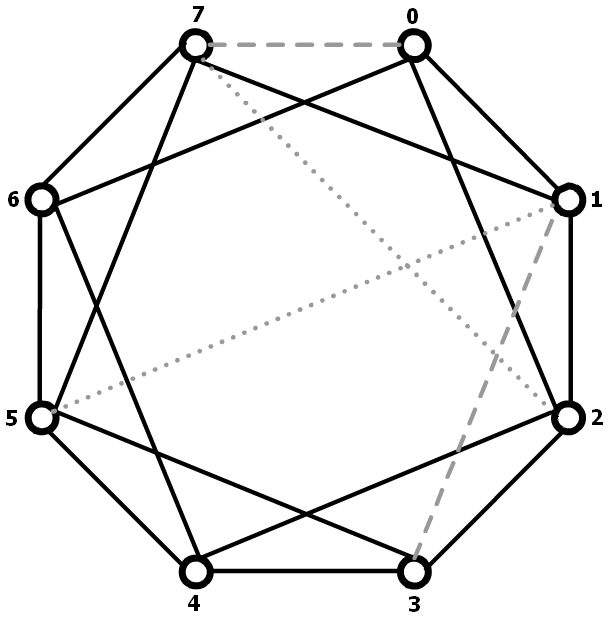

Figure 20. A Watts-Strogatz graph illustrating potential symmetry errors. Dashed grey lines indicate false-alarm symmetry errors. Dotted grey lines indicate miss symmetry errors.

When trained with a diagram, subjects remembered (or perceived) irregular network graphs to be more symmetrical than they actually were. This conclusion accords well with findings of symmetry bias in visual form perception and memory (Freyd & Tversky, 1984; Tversky & Schiano, 1989). Tversky and Schiano showed subjects asymmetric curves that looked like normal curves skewed to various degrees either left or right. Subjects were asked to draw from memory the curve they had studied. Subjects'

drawings, on average, were more symmetrical than the studied stimulus. The same effect (to a lesser degree) was present when subjects merely copied a stimulus they had in front of them. The curves were both misperceived and misremembered as more symmetrical than they were. Provisionally, it can be suggested that the same process led to the greater proportion of symmetry errors in the Diagram condition (in which this effect could operate) than the Verbal condition (in which it could not).

One should be mindful of the brain's symmetry bias when representing network graphs that are not symmetrical. Symmetry bias is greatest for near-symmetric stimuli and weakens the further from symmetry a form strays. When constructing a diagram for an irregular graph, it may be wise to accentuate the asymmetries in the diagram to avoid the errors that would arise due to the brain's tendency to erase small deviations from symmetry.

4.2 Experiment 2

The previous experiment compared the acquisition of graphs at two points on the regularity-to-randomness dimension. There are many other dimensions on which graphs may vary. One of the most fundamental is density.

Graph density is defined as the proportion of possible edges that are valid edges. Density takes a value between zero and one. An empty graph contains no edges and has density zero. A fully connected graph contains all possible edges and has density one. In a low density graph, few of the edges that might exist do exist. In a high density graph, most of the edges that might exist do exist.

Density is perhaps the most informative single number that describes a graph. With just the density, one can place a subjective probability on the existence of each edge in a graph. (All probabilities would be equal and equal to the density.)

One definition of graph entropy relies solely on graph density and the number of nodes in the graph (Li, Wang, Wang & Zhou, 2008). Graph entropy is a measure of the information content of a graph. Calculating graph entropy is an attempt to quantify – in bits – the minimum size a message would need to be to completely describe a graph. There are competing definitions of graph entropy (see Solé & Valverde, 2004 and Gadouleau & Riis, 2011). I use the definition of Li, Wang, Wang and Zhou exclusively here, because it is the most straightforward to apply and explain.

Li entropy is directly analogous to binomial entropy. Figure 21 plots entropy as a function of the probability of an event X. If event X is sure to never happen and thus has probability zero, there is no doubt about the outcome and no information gained when X does not occur. Thus entropy is zero when the probability of X is zero. A very similar argument explains why entropy is zero when the probability of X is one. It is sure to happen, and no information is gained when it does. Maximal entropy exists when the probability of X is .50. This is the least predictable system, and thus uncertainty as to the outcome is greatest. When the outcome is revealed, uncertainty is reduced to zero. The greatest change in information occurs in this case.

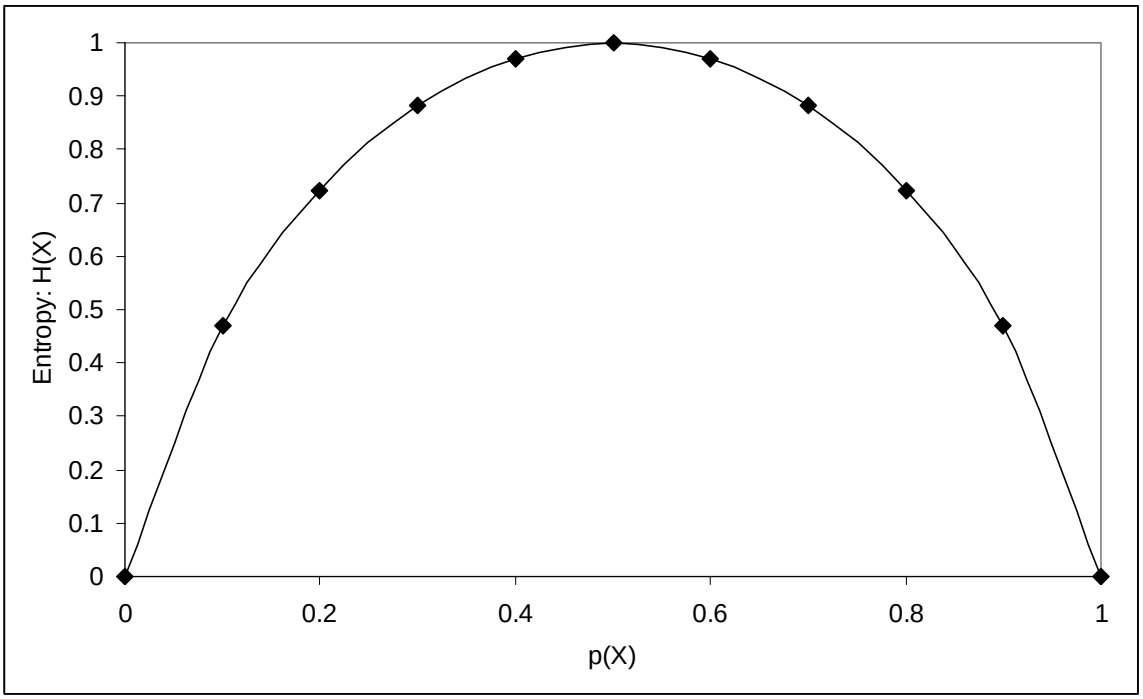

Figure 21. Entropy as a function of the probability of an event X.

The Li entropy of a graph is nearly identical, except that graph density is substituted for event probability. Figure 22 plots Li entropy as a function of graph density for graphs of three different sizes. The shape of the function is the same for all three graphs – zero at density zero or one and maximal at density = .50. As one would expect, larger graphs are defined to contain more information, because there are more independent events (i.e. edges) occurring with a given probability (i.e. density).

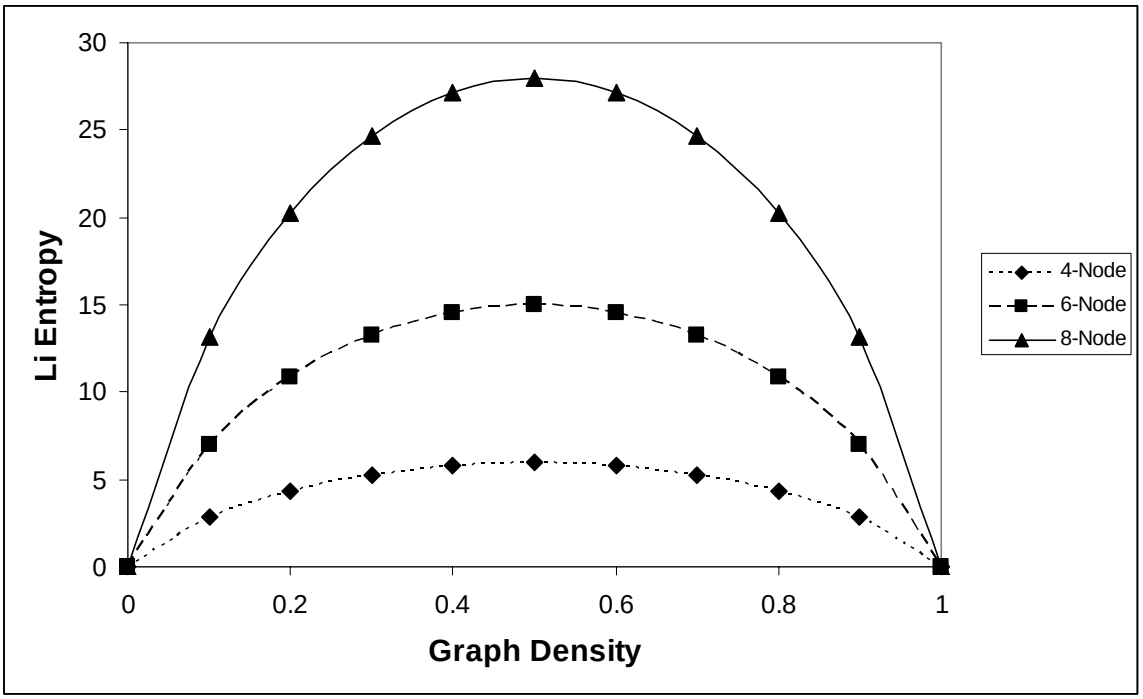

Figure 22. Li entropy as a function of graph density for graphs of 4, 6 and 8 nodes.

Li entropy makes a strong prediction regarding density and acquisition. Given the assumption that people acquire information at a constant rate (on average) it should take longest to approach total knowledge of a medium-density graph, because medium-density graphs contain the most information. Low and high density graphs contain less information. Low and high density graphs should be more easily learned because there is a default state for edges (invalid and valid, respectively) and learning can be focused on the exceptions to the general rule. Either state (valid or invalid) is equally likely in a medium-density graph.

Separate from the effects of density on acquisition, this experiment will again address the practical question of how best to communicate graph structure. One fact to keep in mind is that many interesting networks are sparsely connected. For example, for

all the many hyperlinks on the World Wide Web, it is absolutely the case that most websites *do not* link to most other websites. The same is true for social networks. Most people will never in their life interact with most other people.

Because many interesting networks are sparsely connected, it is tempting to teach the structure of a network by enumerating only its valid edges. If all unnamed edges can be assumed to not exist, naming only valid edges would seem the most efficient method to communicate the existing structure of a graph. However, the acquisition of high-density graphs may suffer in this context. If most edges are valid, and one is trained on a list of valid edges, learning depends on noticing which edges are *absent* from the list – a difficult task. To distinguish it from other types of training, training on valid edges only will be referred to as Positive training in this manuscript.

Another method of training would be to spend equal time training the set of valid edges and the set of invalid edges. Consider pairs of training edges in which one edge is always a valid edge and the other is always an invalid edge. This type of training seems like a natural fit for medium density graphs. In addition to efficiently revealing the state of edges in a medium density graph, this method suggests a medium density state by alternately presenting valid and invalid edges. For the same reasons, this method may be ill-suited to low and high density graphs. Edges of the rare class will necessarily be oversampled and repeated, which may be a waste of training time while edges in the common class have yet to be revealed. Additionally, training on equal numbers of valid and invalid edges may give the learner the false impression that a graph is of medium density when it is not. Training on valid and invalid edge pairs will be referred to as Equal training.

Finally, rather than selecting edges for training in any systematic manner, one might simply sample edges and reveal their state to the learner. Researchers and pollsters randomly sample populations in order to efficiently gather information about general tendencies. Similarly, a random sample of edges will quickly provide information about density, because the ratio of valid to invalid edges trained should converge toward the true ratio. A random sample avoids the redundant presentation of the rare class that is a problem with Equal training. The learner will not need to infer that an edge is invalid through its absence as he would under Positive training. Invalid edges will be explicitly named as such. Training on randomly sampled edges will be referred to as Sample training.

This experiment compares three types of training applied to graphs of three different densities. The results will have implications for theories of graph information complexity and will provide practical recommendations regarding how best to teach network structure.

Method

Participants

225 users of Amazon Mechanical Turk participated in the study. Participants ranged in age from 18 to 60 years old with a mean age of 32. 60.4% of the subjects were female. Participants were paid $0.45 upon successful completion of the experiment.

Design

The design was a 3 (Graph Density) x 3 (Training Regimen) between-subject design. Each subject was trained and tested on one graph under one training regimen.

Subjects were randomly assigned to conditions, under the constraint that the number of subjects in each condition should be as near equal as possible.

Stimuli

All graphs consisted of 8 nodes. Edges were undirected and unweighted. Self-edge loops were not allowed. In each graph, there were 28 potential edges (i.e. 8 choose 2). A participant was trained on one of three graphs varying in density.

Graph Density. Figure 23 depicts the three graph stimuli. In an 8-node, undirected graph there are 28 potential edges. The Low Density graph contains 7 edges and has density 7/28 = .25. The Medium Density graph contains 14 edges and has density 14/28 = .50. The High Density graph contains 21 edges and has density 21/28 = .75.

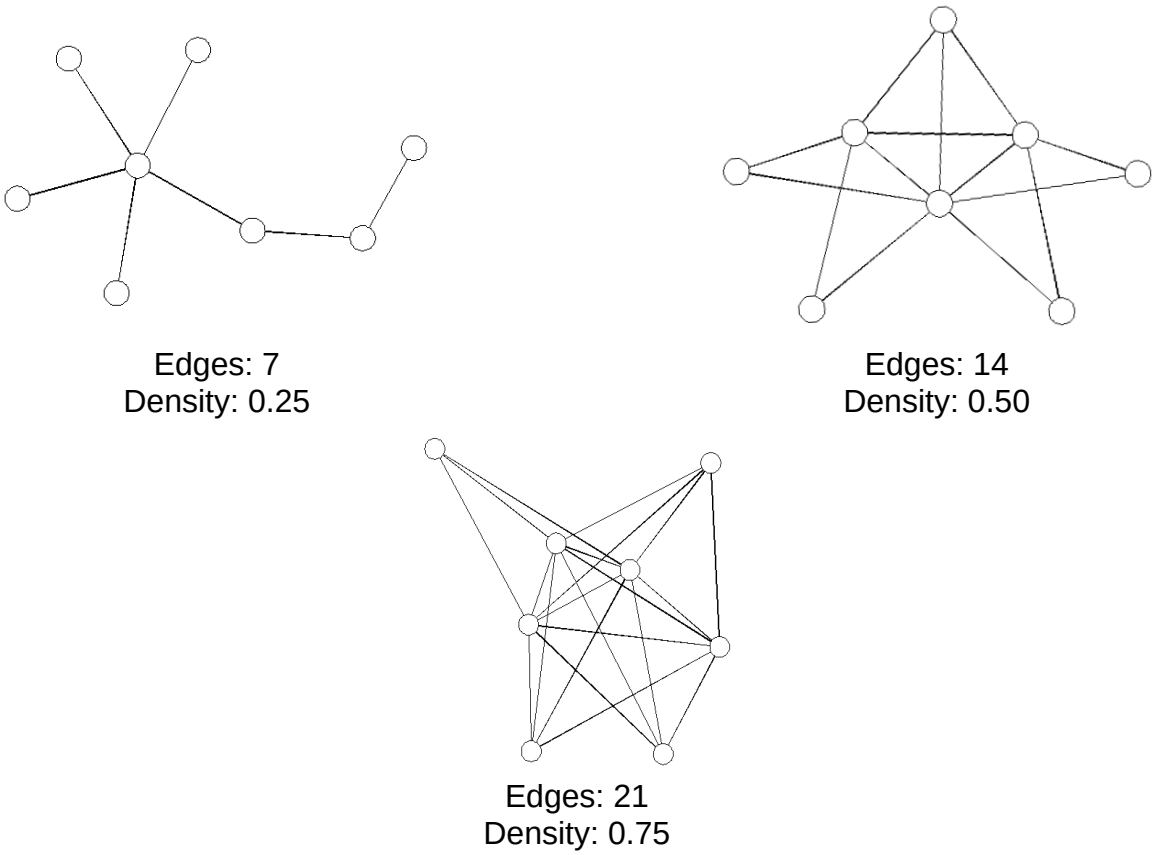

Figure 23. Network diagrams of the Low, Medium and High Density graph stimuli.

Previous work (Chapters 2 and 3) has shown that subjects learn scale-free graphs – graphs in which few nodes have many connections and many nodes have few connections – most quickly. The stimuli graphs were constructed to be scale-free in their structure, so that subjects would not find structure to be an impediment to their learning.

It is worth noting that the High Density graph contains subgraphs of valid edges equivalent to the valid edges of the Medium and Low Density graphs, and the Medium Density graph contains a subgraph of valid edges equivalent to the valid edges of the Low Density graph. Also, the Low Density graph contains subgraphs of *invalid* edges equivalent to the invalid edges of the other two graphs and the Medium Density graph

contains a subgraph of invalid edges equivalent to the full set of invalid edges within the High Density graph.

Name Lists. Four lists containing four male names and four female names were constructed by choosing names from the Social Security Administration's list of popular baby names in the 1990's (Social Security Administration, 2011). The initial letter of each name was unique within each list. At the beginning of each block, the stimuli names were randomly mapped to nodes in the abstract graph structure.

Procedure

The experiment was administered online through a PHP/mySQL/Flash Web application. Subjects followed a link to the experiment website from the Amazon Mechanical Turk user interface.

Subjects first read through several short paragraphs of instructions. They were told they would be learning "who is friends with whom" in a group of eight people. It was explained they would learn by observing pairs of names, some labeled FRIENDS and some labeled NOT FRIENDS. Subjects were informed that all links were reciprocal – if A was friends with B, then B was also friends with A. They were told to consider two people friends only if the experiment explicitly named them as friends.

Subjects were instructed to expect test trials in which two stimuli would be presented and they were to answer FRIENDS if the people were friends or NOT FRIENDS if they were not. Subjects were instructed to give their best guess if they did not know the answer.

In the experiment, training trials alternated with test trials. There were 112 training and test trials.

Training Trials. Each training trial consisted of two edges presented to the subject. If an edge to be trained was valid (that is, if the two nodes were connected in the graph), the names associated with each node were presented in green, close together, with a line between them and the word FRIENDS printed above. If the edge to be trained was invalid (not connected in the graph), the names were presented in red, far apart, with empty space between them and the words NOT FRIENDS printed below. Figure 24 depicts a training trial for a valid edge and Figure 25 an invalid edge.

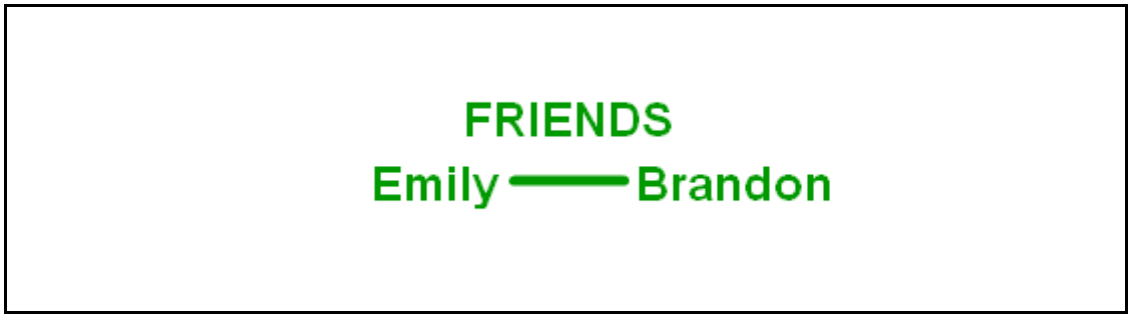

Figure 24. A valid edge training trial.

Figure 25. An invalid edge training trial.

An edge presentation began with a blank, white screen presented for 500 milliseconds. Next, the name of a node appeared alone for 750 milliseconds. Then the name of the second node appeared. (At this point, the color and spacing of the names already identified the edge as valid or invalid.) After 750 milliseconds, the label

FRIENDS or NOT FRIENDS appeared along with a connecting line in the case of FRIENDS. The complete stimulus remained onscreen for 1 second.

In the Positive training condition, subjects observed valid edges only. All valid edges were sampled without replacement until all had been presented. This process then repeated. The number of times each edge was trained depended on the density of the graph. The lower the density, the fewer valid edges, and thus the more training repetitions there were of each valid edge.

In the Equal training condition, subjects observed one invalid edge for each valid edge. Valid edges were sampled without replacement and paired with invalid edges sampled (separately) without replacement. When a sampling process ran out of edges, all edges of that type (and only that type) were made eligible for sampling again. Consequently, valid edges were repeated during training more often than invalid edges in the Low Density graph condition. The converse was true in the High Density graph condition. In the Medium Density graph condition, each edge was the object of training an equal number of times.

In the Sample training condition, subjects observed randomly sampled potential edges. All potential edges were sampled without replacement until all had been presented. This process repeated exactly 8 times. (28 potential edges x 8 presentations = 224 = 112 training trials x 2 edges per trial.)

Test Trials. Potential edges were randomly sampled for testing without replacement until all had been tested. This process then repeated. All 28 edges were tested 4 times each.

In a test trial, two names appeared next to each other. Above and below the names were two equally-sized light-grey buttons. The upper button was labeled FRIENDS and the lower button was labeled NOT FRIENDS.

Feedback was provided after the subject made a response. The button for the correct answer was highlighted, and a 300 millisecond sound clip of a bell was played for a correct response or a buzzer for an incorrect response. Near the button that the user clicked, a smiling face appeared if the response was correct, or a frowning face if the response was incorrect. Figure 26 depicts a test trial in which the subject incorrectly chose the NOT FRIENDS response.

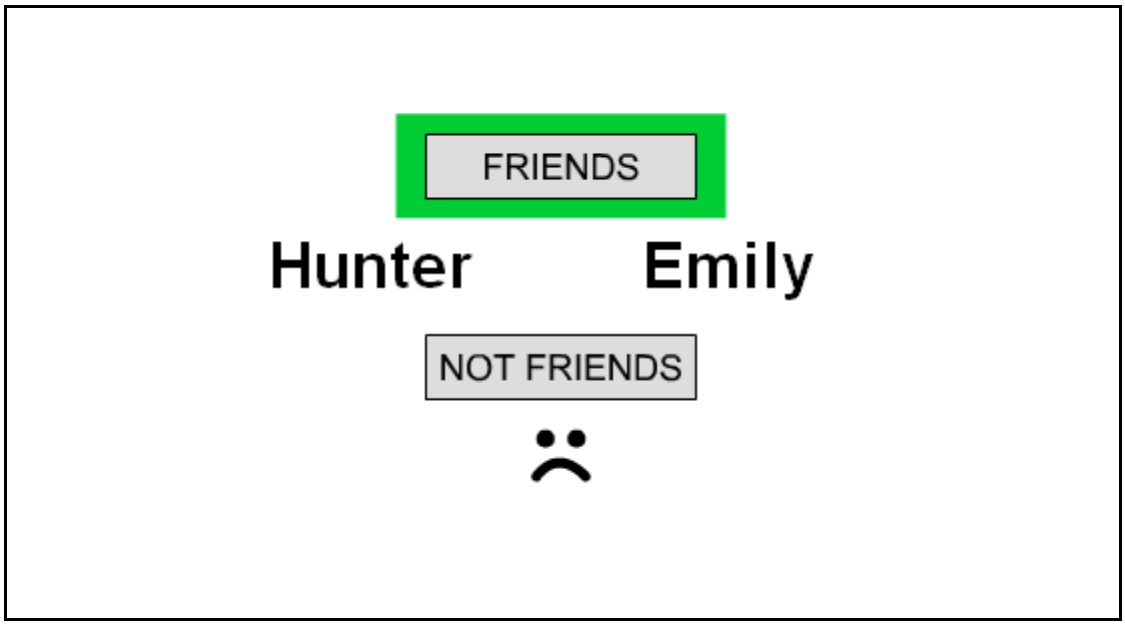

Figure 26. A test trial in which the subject chose NOT FRIENDS as a response when the correct answer was FRIENDS.

Results

d' for each combination of conditions is graphed in Figure 27. In all conditions, discrimination was statistically significantly greater than the d' value of zero expected

from uninformed guessing. ANOVA was performed to test for effects of graph density and training regimen. The main effect of Graph Density was marginally significant F(2, 216) = 2.25, MSE = .49, p = .10. Mean d' for Low Density graphs was .99, while Medium Density d' was .80 and High Density d' was .77.

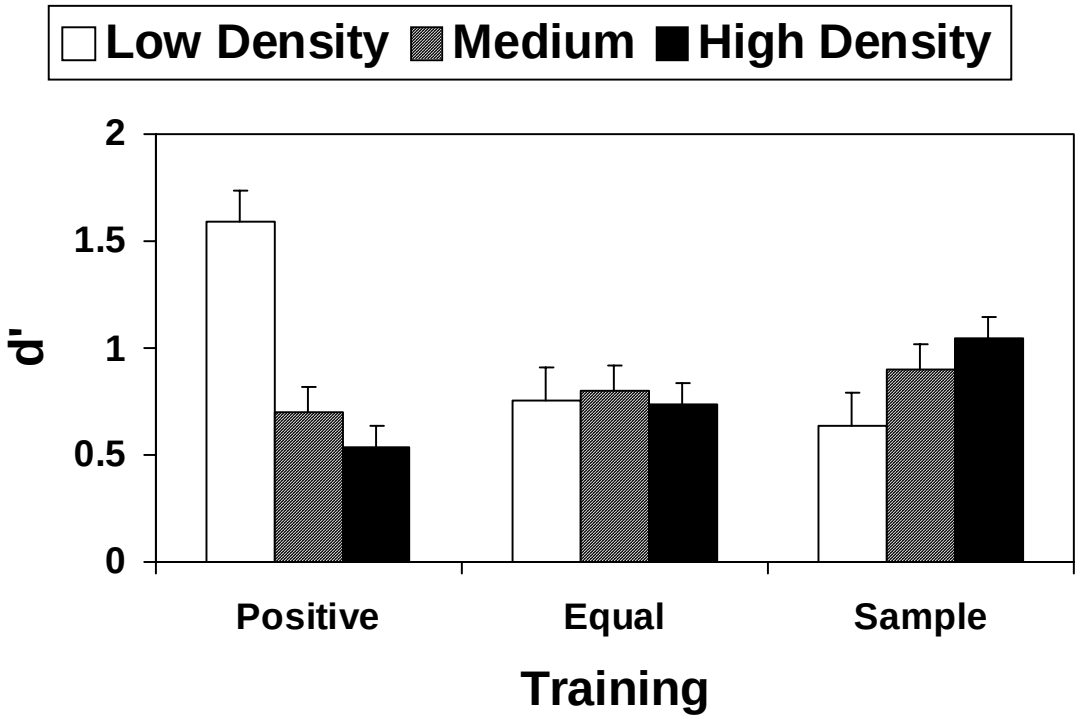

Figure 27. d' in each combination of graph density and training regimen condition.

There was not a consistent main effect of Training Regimen F(2, 216) = 1.21. Instead, the influence of training was apparent in the interaction of Training Regimen and Graph Density F(4, 216) = 8.15, MSE = .49, p < .001. For Medium Density graphs, subject performance is similar across training conditions. Subjects excel when a Low Density graph is trained exclusively with valid edges. High Density graphs are better learned through Sample training than Positive training.

Discussion

The predictions made based on calculations of Li entropy did not hold. It was not the case that Low and High Density graphs were consistently acquired more quickly than the Medium Density graph.

Still, the results demonstrate that it is possible to construct training such that subjects learn a low density graph quickly. The results also imply that acquisition of high density graphs varies by training regimen and is best when training consists of sampling edges and revealing their state. It would be interesting to attempt to train a high density graph through "negative" training – that is, teach the learner only which edges are invalid. It is possible the extraordinary performance seen for positive training of the low density graph could also be attained for high density graphs in this way. It is not clear how (or if) one could design training so that learning would be fastest for graphs of medium density.

Training regimen on its own had no reliable effect. However, the interaction makes clear the practical recommendations to be made regarding how to teach network graph structure. If the network to be trained is of low density, it is best to train using only positive examples. Choice of training regimen seems to matter less for graphs of medium to high density. One might wish to favor random edge sampling. Discrimination performance trends higher the more training shifts away from systematic selection of edges and toward simple random sampling.

Conclusion

There is no one-size-fits-all best method for the training of graph structure. Making general recommendations without sounding vacuous or adding a string of qualifications is not easy. Nonetheless, following is a short list:

- Diagram the network. This makes all information immediately available to the learner. The learner can focus his attention as necessary to learn about specific elements in the graph.
- Arrange the nodes in the diagram to emphasize the information you wish the learner to acquire. If the graph is regular, make the diagram representation symmetrical. If it is near-regular, emphasize the asymmetry or risk symmetry errors.
- If it is not feasible to diagram the network, select the appropriate form of edge training based on graph density. Low density graphs should be trained by enumerating the valid edges. All other graphs will likely benefit from a simple strategy of randomly sampling potential edges and revealing their state.

The current work illuminates a critical shortcoming in network theory. It is not well-defined how we should quantify the complexity or information content of a graph. Graph entropy failed to consistently predict acquisition rates. There exists a great multitude of quantifiable graph features (density, connectedness, efficiency, transitivity, etc.). No one variable has yet asserted itself as the best predictor of acquisition difficulty. Finding the right variable (or minimum set of variables) to adequately describe the informational complexity of a network would be very useful.

More work is necessary to advance on this goal. Better models will aid with practical applications such as predicting how long it will take to learn a given network and how to construct networks and hierarchies that are easier for humans to internalize and understand. Additionally, better models will provide insight into how the brain forms and maintains associations in large, potentially very complex spaces.

Chapter 5: Modeling Learning Performance

In all the experiments in this work, learning the structure of a network has followed a number of patterns: accuracy increases with more training, accuracy is higher when learning scale-free graphs as opposed to random graphs, and accuracy is highest when tested on edges containing nodes of extreme centrality. Here, I test two simple models of graph acquisition. These learning systems were given the same input in the same sequence as human learners and tested on the same edges. How the systems' output matches and strays from subject performance sheds light on how humans approach the problem of learning a graph structure.

Perfect Memory Model

The first class of model rests on one simple assumption: at each learning opportunity, the system (with some probability) forms a perfect representation in memory of the information presented and never forgets that information. Because encoding is perfect, and there is no loss or interference at storage or retrieval, the model is called the Perfect Memory Model.

One parameter – p – represents the probability that a bit of information presented by the experiment is stored in memory. If that bit of information is stored, the information is stored without distortion and for the duration of training and testing. The model responds to test edge prompts for which it has information with perfect accuracy. If the state of the edges is unknown to the model, it will respond Yes with probability .50. Versions of the model at varying levels of p were trained and tested using training and testing histories identical to those of subjects in Experiment 2 of Chapter 2.

Figure 28 plots the learning curves produced by the model and the data from subjects in the experiment. The best fitting model is one with a low value of p. A value of .05 produces a close match. Interpreting this result provides a convenient rule-of-thumb. Recall that in this experiment, graphs had twenty nodes. In each training trial, one node was chosen and presented together with a list of all nodes it formed an edge with. By implication, the states of all possible edges involving the trained node (19 edges in all) were revealed. Names present formed valid edges and names absent invalid edges.

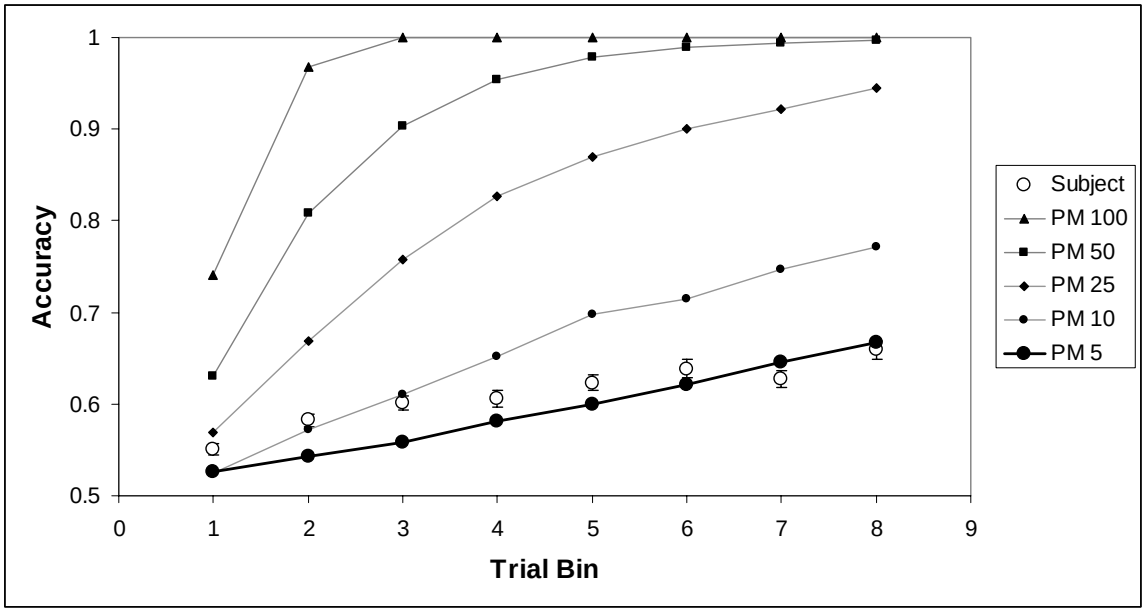

Figure 28. Accuracy as a function of training for subjects and five versions of the Perfect Memory Model. PM5, the version with probability of .05 to remember each trained edge is the best fit.

A value of p at .05 implies that the state of approximately one edge was learned per training trial. (More precisely, this is the expected value of edges learned per training trial.) The model suggests that despite the fact that all of the information about one node's edges was available on each training trial, what subjects took away was knowledge about one edge.

Another fact Figure 28 makes clear is that all of the information about each graph has been delivered to the model (and subjects) by the end of the second trial bin. Trial bins are 10 trials wide, so this means after 20 trials the complete structure of the graph has been made known. This makes sense, because each graph has 20 nodes, and one node is trained per trial. If subjects had perfect memory (like a Perfect Memory Model where p = 1), they would know all there is to know about the graph. For both subjects and models, all training after the first 20 trials repeats information they have already had at least one opportunity to learn.

Figure 29 plots accuracy over the course of the experiment separated by graph structure. It can be seen that unlike human learners, the model does not distinguish between random and scale-free graphs. The model grabs bits of information regarding edges as stochasticity allows and is not affected by graph structure.

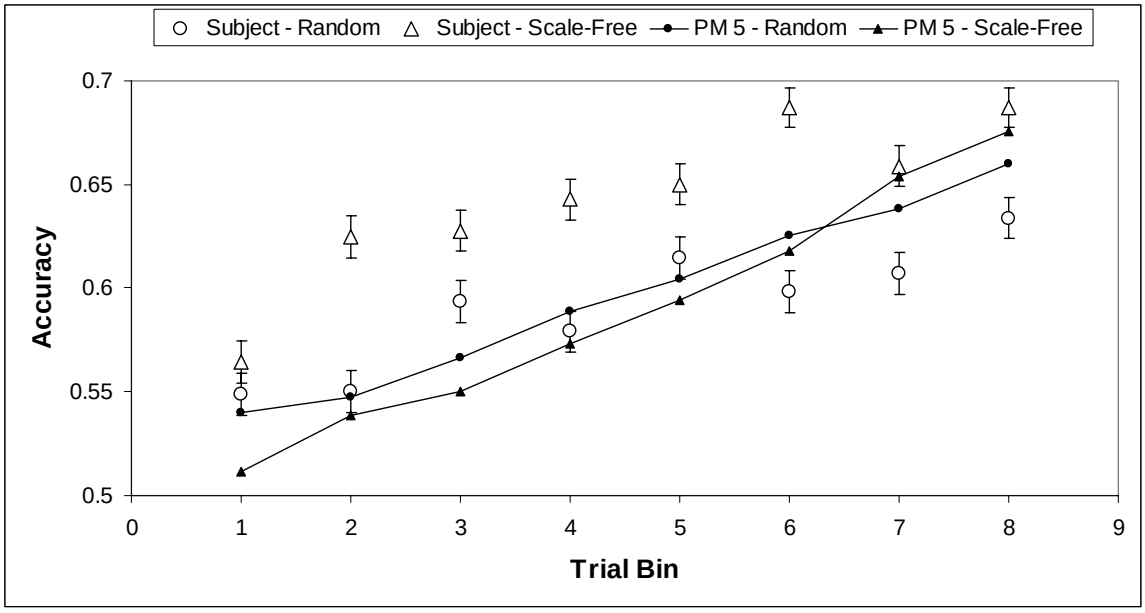

Figure 29. Accuracy as a function of training for subjects and PM5. The subject data (open markers) shows separation by graph structure. The model data (filled markers) does not differ between graph structures.

Figure 30 plots accuracy as a function of mean eigenvector centrality of the edge nodes. The model curve is flat, but human performance (as it has been in repeated experiments) is best for the most peripheral and the most central nodes.

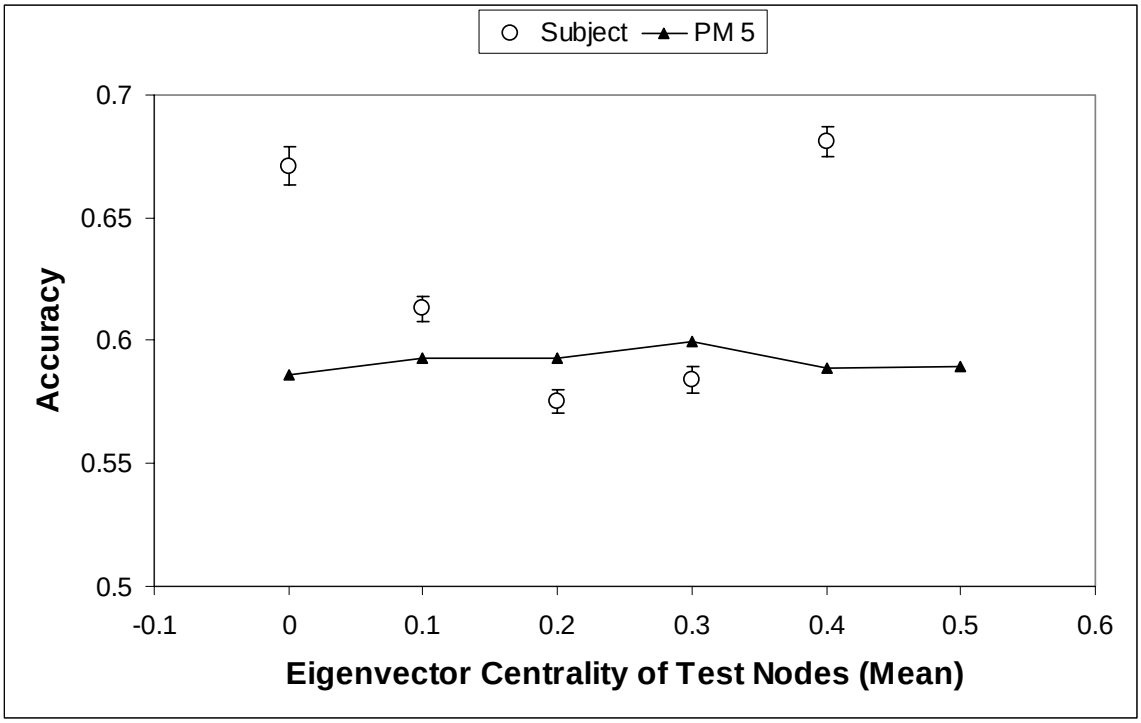

Figure 30. Accuracy as a function of test node centrality for subjects and PM5. The subject data (open markers) rises at both extremes. The model (filled markers) does not respond to node centrality.

In addition, the model responded Yes only 50% of the time, whereas subjects responded Yes 55% of the time, showing the usual yes-bias in responses.

At the highest level of abstraction, the Perfect Memory Model does well at producing a learning curve similar to that of human subjects. It also quantifies just how much information subjects take in at each training trial and indicates it is just a fraction of the available information. The model fails to capture the more subtle patterns of human performance, however. The model is insensitive to graph structure and node properties. It cannot reproduce the yes-bias ever present in human responses.

Frequency Heuristic Model

The Perfect Memory Model assumed that learning was focused on edges. The Frequency Heuristic Model instead focuses on nodes. This model does not even bother to represent beliefs about the states of possible edges. The Frequency Heuristic Model simply counts the number of occurrences of each node during training and relies on the (correct) assumption that highly connected nodes appear more often to make educated guesses regarding the existence of edges.

Specifically, the model stores 20 counts, one for each node. Each time a node was presented for training and each time it appeared in a list of connected nodes, the node's counter was incremented. When an *edge* was presented for testing, the model made two calculations. First, it summed the incidence counts of the two nodes in the test edge. Then it calculated the mean summed incidence counts over all possible two-node combinations. Thus, it calculated 190 sums. For every *possible* edge in the graph, it added the incidence counts of the two nodes. The model then compared two numbers. The first was the summed incidence for the two nodes in the test edge. The other number was the mean of the summed incidences for the other 189 possible edges. If the summed incidence of the test edge was greater than the mean for all other possible edges, the model responded Yes. If the test edge value was less than the mean, the model responded No. If the values were equal, the model responded Yes with probability .50.

(This model was compared to a model that used the median as the criterion and to a model that used the incidence counts in a softmax decision process. No substantial differences in model fit were found. Being the simplest to explain and implement, the mean model is presented exclusively.)

Figure 31 plots the learning curve produced by the model and the data from participants in the experiment. This simple model ignores the information about specific edges the training trials are meant to convey. This simple model contains no free parameters that are tuned to this particular learning task. And yet, this simple model performs better than subjects do.

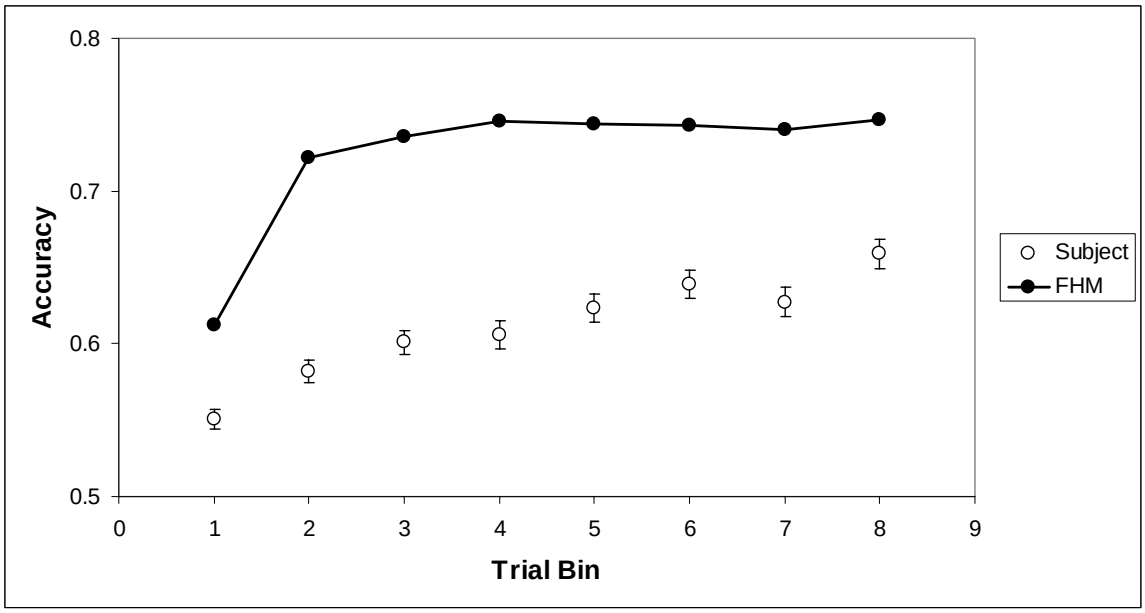

Figure 31. Accuracy as a function of training for subjects and the Frequency Heuristic Model. FHM achieves higher accuracy, but asymptotes, whereas subject performance is still improving.

In defense of the human subjects, the model does possess complexity not immediately obvious. The model keeps perfect track of 20 independent values at all times. It averages 190 sums each time it answers a test prompt. Introducing reasonable levels of noise to the model's observations and calculations would necessarily decrease its performance.

Another trend to notice is that learning plateaus after trial bin 2. This happens because the first 20 trials provide all of the useful relative frequency information. After

this point, sets of training trials only multiply the incidence counts by a constant value. The different permutations of trained nodes subjects received introduces some random noise per trial bin (as can be seen in Figure 31), but no additional information can be learned nor any improvement in performance made.

Figure 32 plots accuracy over the course of the experiment and separated by graph structure. The model possesses the same advantage for scale-free graphs over random graphs that human subjects have. This is due to the greater variability in node degree present in scale-free graphs. The frequency heuristic is simply more effective for scale-free structure than random structure.

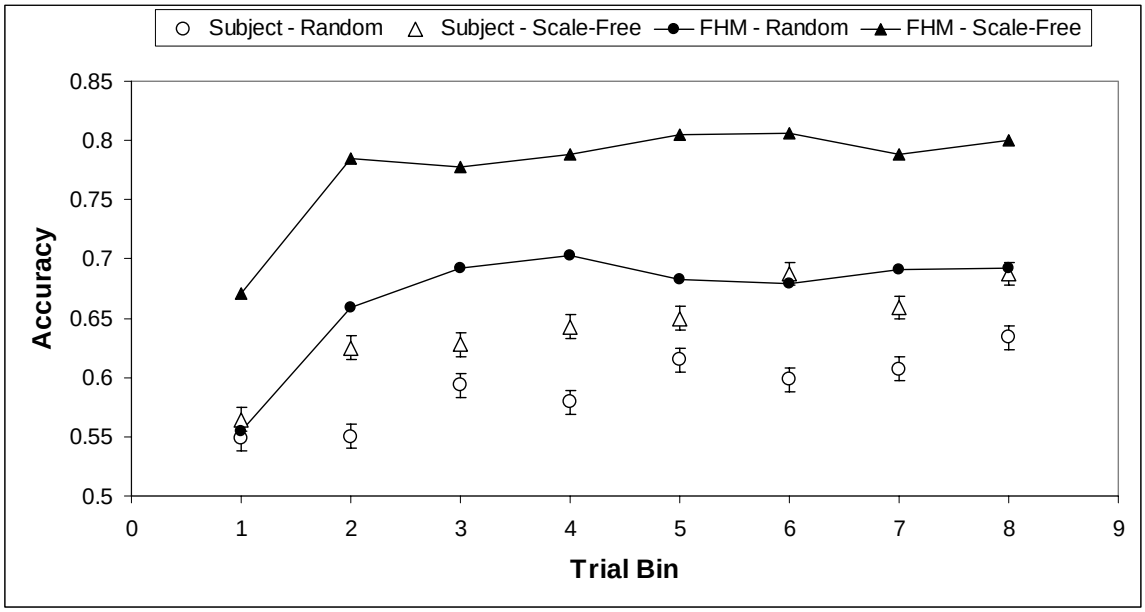

Figure 32. Accuracy as a function of training for subjects and the Frequency Heuristic Model. The model performs better in the scale-free condition, just like subjects.

Figure 33 plots accuracy as a function of node centrality. The model displays the same U-shaped accuracy curve as human subjects. The frequency heuristic is more effective at classifying edges in which nodes are of extreme centrality. Peripheral nodes

appear less often during training, and thus the model is likely to choose a No response when an edge contains these nodes. Two peripheral nodes are unlikely to be connected, so the model's bet is often a good one. Similar logic explains why the model often responds correctly for central nodes.

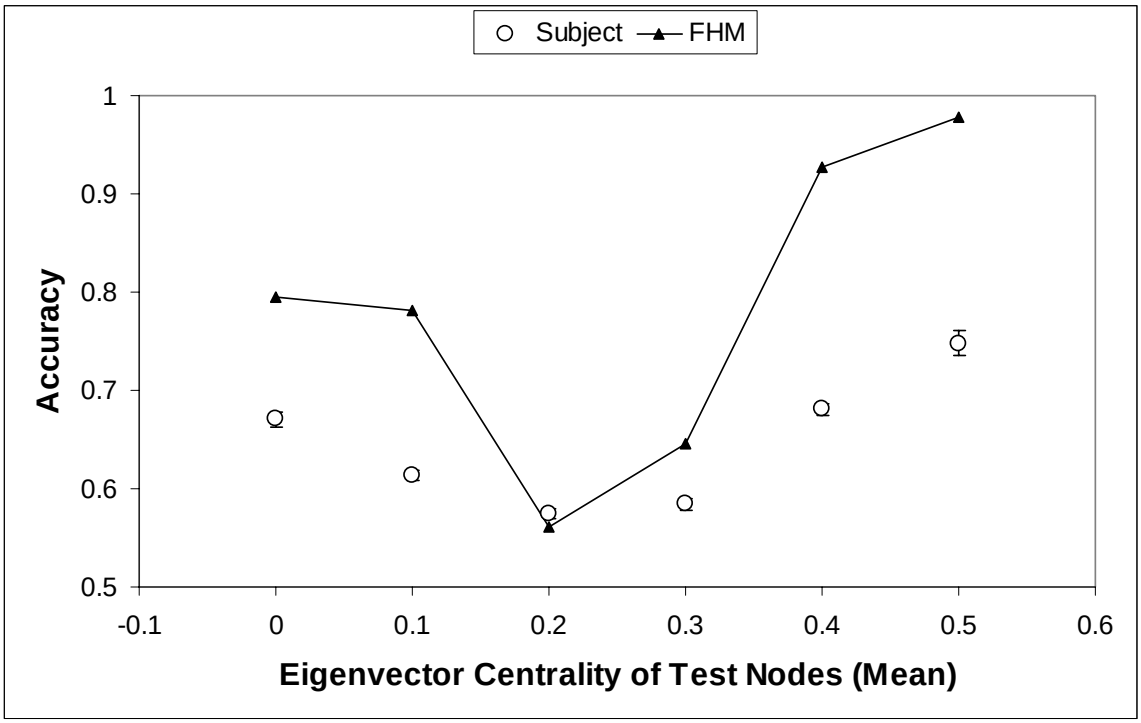

Figure 33. Accuracy as a function of test node centrality for subjects and FHM. The model data displays the same U-shaped pattern as the subject data.

The yes-bias of the Frequency Heuristic Model is 63% and is even more severe than subjects' yes-bias.

The Frequency Heuristic Model produces all the patterns of performance evident in human data. This suggests that human learners rely on a similar heuristic. The failure of the Perfect Memory Model to capture these patterns, and the low value of p that produced a good fit suggests that human learners rely on the heuristic more than they rely on storing specific edge information in memory.

Figure 34 sketches a multinomial model that subsumes the two models implemented here.

- The probability p is that from the perfect memory model. It represents the probability the state of the test edge has been represented in memory. p is likely to be low, at least during early training.
- Probability q extends the perfect memory model. It allows for the possibility that the learner stored the wrong value in memory or that the learner makes a response error. Fitting a value to q would reveal how likely these events are.
- Probability (1-p) is the probability no specific edge information is stored for a test edge. It is likely to be high, at least during early training.
- Probability r is the probability that the learner resorts to the frequency heuristic. This probability is likely to be high, making this path the most likely in the model.
- Probability (1-r) is the probability the learner does not have enough information to trust the frequency heuristic or is not willing to expend the effort. The subject will simply guess.

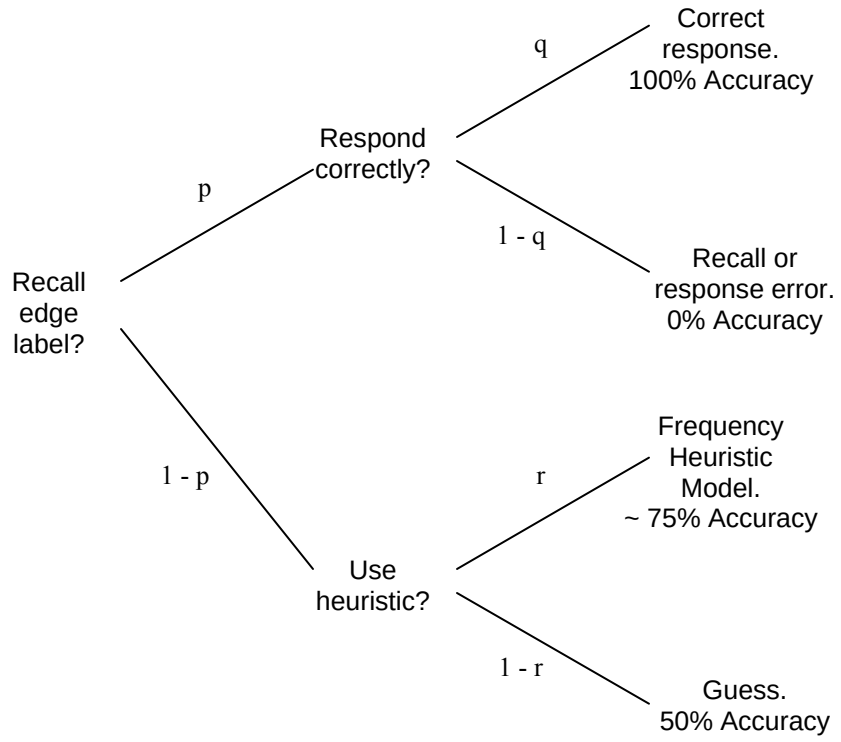

Figure 34. A multinomial model that subsumes the Perfect Memory and Frequency Heuristic Models.

Fitting values to these parameters after varying amounts of training would reveal the degree to which each branch is employed and how this varies as the learner accrues knowledge. It is necessary to contemplate such things if one imagines how learning would continue past the amount of training involved in the current experiment.

Surely, with enough training one could learn even a complex network with complete fidelity. The Frequency Heuristic Model is successful at matching the current data, but it is not capable of achieving full knowledge. Indeed, the curves suggest the model reaches its upper asymptote of accuracy approximately halfway into training. Human learning appears to continue through the last trial, albeit at a slow rate. The full multinomial model, though somewhat unwieldy, is capable of representing a learner at any stage of learning – from novice to infallible expert.

Chapter 6: Concluding Remarks

Network graphs describe how everything and everyone is connected. It is essential that we understand them. Hopefully, this work will add to our knowledge of how the brain comprehends complex systems and point the way toward techniques to make that comprehension more efficient.

I have introduced the Anthropomorphic Acquisition Hypothesis – which states that people use their own social network graph as a model for learning new networks. The hypothesis receives some support. People learn networks with the same scale-free structure as human relationship networks more quickly than random networks. This is true regardless of the surface description of the network or the training task. However, it is not necessarily true that the reason scale-free graphs are learned efficiently is because of their resemblance to social network structure. As they do in many areas, it may be that human learners are simply responding to the statistics of the environment. The fact that the frequency heuristic model mimics human performance so well suggests that (at least initially) learners identify nodes as generally well-connected or poorly-connected and make educated guesses about the validity of edges based on these observations.

Some results are not consistent with the Anthropomorphic Acquisition Hypothesis. A perfect, symmetric network was learned more readily than a less regular, more realistic network. This suggests more spatial, perhaps perceptual processes are at work representing networks in the brain rather than those concerned with social cognition.

A cavalcade of future studies suggest themselves in this line of research. For example, what accounts for individual differences in acquisition? Will intelligence predict who will acquire a new network structure quickly? Do people with larger and denser social graphs learn network structure more efficiently?

It is clear that graph structure affects acquisition. It is encouraging that the current experiments found no evidence that the surface description of a network graph had any effect. If it is indeed the case surface description matters little or not at all, general theories of graph acquisition may be simpler and yet still powerful.

As the world becomes more connected and complex, an understanding of network graph acquisition will become more and more essential. It is already necessary today in fields as disparate as integrated circuit design and international counterterrorism. Studying acquisition of graph structure will give us the means to learn natural networks more efficiently and to design new networks and hierarchies in a way that best suits our cognitive architecture.

References

Alexa Internet, Inc. (2011). Top Sites, *Alexa*. Retrieved June 1, 2011, from http://www.alexa.com/

Arnell, K. M., Shapiro, K. L., & Sorensen, R. E. (1999). Reduced repetition blindness for one's own name. *Visual Cognition, 6(6)*, 609-635.

Barabási, Albert-László & Albert, Réka (1999). Emergence of scaling in random networks. *Science, 286,* 509–512.

Bhalla, U. S., & Iyengar, R. (1999). Emergent properties of networks of biological signaling pathways. *Science, 283(5400)*, 381-387.

Bonacich, P. (1987). "Power and Centrality: A Family of Measures." *American Journal of Sociology, 92,* 1170-1182.

Brin, S., & Page, L. (1998). The anatomy of a large-scale hypertextual web search engine. *Computer Networks and ISDN Systems, 30(1-7),* 107-117.

Bunke, H., & Shearer, K. (1998). A graph distance metric based on the maximal common subgraph. *Pattern Recognition Letters, 19(3-4)*, 255-259.

Burt, R. S., & Ronchi, D. (2007). Teaching executives to see social capital: Results from a field experiment. *Social Science Research, 36(3),* 1156-1183.

Challis, B. H., Velichkovsky, B. M., & Craik, F. I. M. (1996). Levels-of-processing effects on a variety of memory tasks: New findings and theoretical implications. *Consciousness and Cognition, 5(1-2),* 142-164.

Cheng, P. C. H. (2002). Electrifying diagrams for learning: Principles for complex representational systems. *Cognitive Science, 26(6)*, 685-736.

Clark, S. E. (1992). Word-frequency effects in associative and item recognition. *Memory and Cognition, 20(3),* 231-243.

Cottrell, N. B. (1975). Heider's structural balance principle as a conceptual rule. *Journal of Personality and Social Psychology, 31(4),* 713-720.

Craik, F. I. M., & Tulving, E. (1975). Depth of processing and retention of words in episodic memory. *Journal of Experimental Psychology-General, 104(3),* 268-294.

Csányi, G., & Szendrői, B. (2004). Structure of a large social network. *Physcial Review E, 69,* 036131.

De Soto, C. B. (1960). Learning a social structure. *The Journal of Abnormal and Social Psychology, 60(3),* 417-421.

De Soto, C. B., Henley, N. M., & London, M. (1968). Balance and the grouping schema. *Journal of Personality and Social Psychology, 8(1),* 1-7.

De Soto, C. B., & Kuethe, J. L. (1959). Subjective probabilities of interpersonal relationships. *The Journal of Abnormal and Social Psychology, 59(2),* 290-294.

Ebel, H., Mielsch, L.I., Bornholdt, S. (2002). Scale-free topology of e-mail networks. *Physcial Review E, 66,* 035103(R).

Erdős, P. & Rényi, A. (1960). On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci, 5,* 17-61.

Flynn, F. J., Reagans, R. E., Amanatullah, E. T., & Ames, D. R. (2006). Helping one's way to the top: Self-monitors achieve status by helping others and knowing who helps whom. *Journal of Personality and Social Psychology, 91(6),* 1123-1137.

Freyd, J., & Tversky, B. (1984). Force of symmetry in form perception. *American Journal of Psychology, 97(1)*, 109-126.

Gadouleau, M., & Riis, S. (2011). Graph-theoretical constructions for graph entropy and network coding based communications. *IEEE Transactions on Information Theory, 57(10)*, 6703-6717.

Gentner, D., & Gentner, D. R. (1983). Flowing waters or teeming crowds: Mental models of electricity. In D. Gentner & A. L. Stevens (Eds.) *Mental Models.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Gick, M. L., & Holyoak, K. J. (1983). Schema induction and analogical transfer. *Cognitive Psychology, 15(1),* 1-38.

Gilovich, T., Medvec, V. H., & Savitsky, K. (2000). The spotlight effect in social judgment: An egocentric bias in estimates of the salience of one's own actions and appearance. *Journal of Personality and Social Psychology, 78(2),* 211-222.

Heider, F. (1958). *The psychology of interpersonal relations.* New York: Wiley.

The Internet Movie Database. (2011). Box office business for The Social Network, *IMDB.com*. Retrieved June 1, 2011, from http://www.imdb.com/title/tt1285016/business

Jang, Y., Wixted, J. T., & Huber, D. E. (2009). Testing signal-detection models of Yes/No and two-alternative forced-choice recognition memory. *Journal of Experimental Psychology: General, 138(2),* 291-306.

Janicik, G. A., & Larrick, R. P. (2005). Social network schemas and the learning of incomplete networks. *Journal of Personality and Social Psychology, 88(2),* 348-364.

Kelley, R., & Wixted, J. T. (2001). On the nature of associative information in recognition memory. *Journal of Experimental Psychology: Learning Memory and Cognition, 27(3),* 701-722.

Koch, C., & Laurent, G. (1999). Complexity and the nervous system. *Science, 284(5411)*, 96-98.

Krackhardt, D., & Kilduff, M. (1999). Whether close or far: Social distance effects on perceived balance in friendship networks. *Journal of Personality and Social Psychology, 76(5),* 770-782.

Klein, S. B., Cosmides, L., Murray, E. R., & Tooby, J. (2004). On the acquisition of knowledge about personality traits: Does learning about the self engage different mechanisms than learning about others? *Social Cognition, 22(4),* 367-390.

Larkin, J. H., & Simon, H. A. (1987). Why a diagram is (sometimes) worth 10000 words. *Cognitive Science, 11(1)*, 65-99.

Li Ji, Wang Bing-Hong, Wang Wen-Xu, & Zhou Tao. (2008). Network entropy based on topology configuration and its computation to random networks. *Chinese Physics Letters, 25(11)*, 4177-4180.

Libes, D., "Choosing a Name for Your Computer", *Communications of the ACM, Vol. 32*, No. 11, Pg. 1289, November 1989.

Liljeros, F., Edling, C. R., Amaral, L. A. N., Stanley, H. E., & Åberg, Y. (2001). The web of human sexual contacts. *Nature, 411(6840),* 907-908.

McDermott, R., Fowler, J.H., & Christakis, N.A. Breaking Up is Hard to Do, Unless Everyone Else is Doing it Too: Social Network Effects on Divorce in a Longitudinal Sample Followed for 32 Years. Manuscript submitted for publication.

McDonald, D., Monaco, A., Guo, R., Fiano, J., Matney, L., Turner, G., et al. (2009). The effect of personal relevance on learning stroke Symptoms/Response. *Western Journal of Nursing Research, 31(2),* 141-152.

Menon, A. K., & Elkan, C. (2011). *Link prediction via matrix factorization.* Manuscript submitted for publication.

Nairne, J. S., Pandeirada, J. N. S., & Thompson, S. R. (2008). Adaptive memory - the comparative value of survival processing. *Psychological Science, 19(2),* 176-180.

Nairne, J. S., Thompson, S. R., & Pandeirada, J. N. S. (2007). Adaptive memory: Survival processing enhances retention. *Journal of Experimental Psychology-Learning Memory and Cognition, 33(2),* 263-273.

Poitou, J. P. (1970). Status congruence as a cognitive bias. *Journal of Personality and Social Psychology, 16(4),* 592-597.

Rein, J R, & Markman, A B. (2010). Assessing the concreteness of relational representation. *Journal of Experimental Psychology. Learning, Memory, and Cognition, 36(6),* 1452-1465.

Rogers, T. B., Kuiper, N. A., & Kirker, W. S. (1977). Self-reference and encoding of personal information. *Journal of Personality and Social Psychology, 35(9),* 677-688.

Scherer, C. W., & Cho, H. C. (2003). A social network contagion theory of risk perception. *Risk Analysis, 23(2),* 261-267.

Shapiro, K. L., Caldwell, J., & Sorensen, R. E. (1997). Personal names and the attentional blink: A visual "cocktail party" effect. *Journal of Experimental Psychology-Human Perception and Performance, 23(2),* 504-514.

Shepard, R.N., Hovland, C.I., Jenkins, H.M. (1961). Learning and memorization of classifications. *Psychological Monographs, 75,* 42.

Social Security Administration. (2011). Top names of the 1990s, *Social Security Online*. Retrieved June 1, 2011, from http://www.ssa.gov/oact/babynames/decades/names1990s.html

Sole, R. V., & Valverde, S. (2004). In Ben-Naim E., Frauenfelder H. and Toroczkai Z. (Eds.), *Information theory of complex networks: On evolution and architectural constraints.*

Thomson Reuters. (2011). *ISI Web of Knowledge*. Retrieved June 1, 2011, from http://pcs.isiknowledge.com/analyze/ra.cgi

Tsujimoto, R. N., Wilde, J., & Robertson, D. R. (1978). Distorted memory for examplars of a social structure: Evidence for schematic memory processes. *Journal of Personality and Social Psychology, 36(12),* 1402-1414.

Tversky, B., & Schiano, D. J. (1989). Perceptual and conceptual factors in distortions in memory for graphs and maps. *Journal of Experimental Psychology-General, 118(4)*, 387-398.

Walker, C. J. (1976). The employment of vertical and horizontal social schemata in the learning of a social structure. *Journal of Personality and Social Psychology, 33(2),* 132-141.

Wang, F., Moreno, Y., & Sun, Y. (2006). Structure of peer-to-peer social networks. *Physical Review E, 73,* 036123.

Watts, D. J. 1999. *Small Worlds: The Dynamics of Networks between Order and Randomness.* Princeton, N.J.: Princeton University Press.

Watts, D. J., & Strogatz, S. H. (1998). Collective dynamics of 'small-world' networks. *Nature, 393(6684)*, 440-442.

Appendix

Experiment 1 Name Stimuli

Students

Andrew	Anthony	Ashley	Amanda
Brittany	Brooke	Brandon	Benjamin
Christopher	Cody	Courtney	Chelsea
Danielle	Diana	Daniel	David
Eric	Evan	Emily	Elizabeth
Faith	Frances	Frank	Forrest
Gregory	Gabriel	Grace	Gina
Hannah	Heather	Hunter	Henry
Ian	Isaac	Isabel	Irene
Jessica	Jennifer	Joshua	Jacob
Kyle	Kevin	Kelsey	Katherine
Laura	Lindsey	Logan	Luke
Michael	Matthew	Megan	Melissa
Nicole	Natalie	Nicholas	Nathan
Oscar	Omar	Olivia	Octavia
Paige	Patricia	Patrick	Paul
Ryan	Robert	Rachel	Rebecca
Sarah	Samantha	Steven	Samuel
Thomas	Timothy	Tiffany	Tara
Vanessa	Veronica	Victor	Vincent

Cities

Alchevsk	Ardee	Abingdon	Abiko
Bibrka	Balbriggan	Bacup	Bando
Chasiv	Carlow	Caistor	Chiba
Debaltseve	Dundalk	Dagenham	Daito
Enerhodar	Edenderry	Ealing	Ena
Fastiv	Fermoy	Fairford	Fujiidera
Hadiach	Gorey	Gorleston	Gamagori
Ichnia	Hillsborough	Harlow	Habikino
Kaharlyk	Inver	Ilford	Ibaraki
Makiivka	Johnstown	Jarrow	Joso
Nadvirna	Kells	Keighley	Kadoma
Obukhiv	Leixlip	Langport	Matsubara
Pavlohrad	Macroom	Nailsea	Nagareyama
Radekhiv	Naas	Oakengates	Obu
Saky	Omagh	Paddock	Ryugasaki
Talne	Passage	Ramsey	Saijo
Uhniv	Shannon	Saffron	Tahara
Valky	Templemore	Tadcaster	Urayasu
Yahotyn	Westport	Uckfield	Wajima
Zhovti	Youghal	Ventnor	Yachimata

Computer Servers

A646.net	A131.com	A508.net	A505.com
B269.com	B764.com	B316.com	B686.net
C828.net	C979.com	C907.com	C023.com
D893.com	D097.net	D137.com	D034.com
E951.net	E177.com	E590.net	E270.com
F230.com	F481.net	F212.net	F766.net
G133.com	G814.com	G372.com	G801.net
H035.net	H979.com	H077.net	H427.com
I864.net	I723.net	I522.net	I230.com
J830.net	J597.com	J470.net	J973.net
K156.com	K470.com	K518.com	K159.com
L969.net	L174.net	L268.com	L858.com
M415.com	M243.com	M233.net	M435.net
N625.com	N575.net	N249.com	N611.net
O655.com	O167.com	O720.net	O759.com
P119.com	P968.net	P221.com	P196.net
Q495.net	Q731.net	Q838.com	Q594.net
R119.net	R766.net	R730.net	R637.com
S707.net	S734.net	S194.net	S703.com
T743.com	T299.com	T737.net	T547.net

Experiment 2 Name Stimuli

Alchevsk	Adrasmon	Asahan	Abiko
Bibrka	Buston	Burmeso	Bando
Chasiv	Chkalovsk	Cilacap	Chiba
Debaltseve	Dushanbe	Demak	Daito
Enerhodar	Farkhor	Elelim	Ena
Fastiv	Ghafurov	Fef	Fujiidera
Hadiach	Hisor	Gunung	Gamagori
Ichnia	Isfara	Ilaga	Habikino
Kaharlyk	Jomi	Jepara	Ibaraki
Makiivka	Khujand	Keerom	Joso
Nadvirna	Mastchoh	Lotu	Kadoma
Obukhiv	Nurak	Mappi	Matsubara
Pavlohrad	Panjakent	Nias	Nagareyama
Radekhiv	Qayroqqum	Oksibil	Obu
Saky	Rumi	Puncak	Ryugasaki
Talne	Sharora	Rantau	Saijo
Uhniv	Tursunzoda	Sugapa	Tahara
Valky	Vose	Tegal	Urayasu
Yahotyn	Yovon	Wonogir	Wajima
Zhovti	Zafarobod	Yalimo	Yachimata