

UC Irvine

UC Irvine Electronic Theses and Dissertations

Title

Using Memory Models to Understand Cognitive Impairment Associated with Alzheimer's Disease

Permalink

<https://escholarship.org/uc/item/7rc193ds>

Author

Westfall, Holly Anne

Publication Date

2023

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA,
IRVINE

Using Memory Models to Understand Cognitive Impairment Associated with Alzheimer's
Disease

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Cognitive Sciences

by

Holly A. Westfall

Dissertation Committee:
Professor Michael D. Lee, Chair
Professor Aaron M. Bornstein
Professor Mark Steyvers

2023

DEDICATION

To my family and friends.

TABLE OF CONTENTS

	Page
LIST OF FIGURES	v
LIST OF TABLES	viii
ACKNOWLEDGMENTS	ix
VITA	x
ABSTRACT OF THE DISSERTATION	xii
Introduction	1
1 A Model-Based Analysis of the Impairment of Semantic Memory	5
1.1 Introduction	6
1.2 Behavioral Data	11
1.3 Modeling Analysis	13
1.3.1 A Basic Model of Odd-One-Out Comparison	13
1.3.2 Restructured Representation Analysis	16
1.3.3 Attention Change Analysis	19
1.3.4 Noisy Access Analysis	22
1.3.5 Modeling Conclusions	24
1.4 Discussion	26
1.5 Publication Note	28
2 A Model-Based Analysis of Changes in the Semantic Structure of Free Recall Due to Cognitive Impairment	29
2.1 Introduction	30
2.2 Behavioral Data	31
2.3 Model Description	35
2.3.1 Regression Function	35
2.3.2 Luce Choice Rule	36
2.4 Results	38
2.4.1 Model Descriptive Adequacy	39
2.4.2 Regression Weights	39
2.4.3 Response Determinism	41

2.5	Components-of-Similarity Model	41
2.5.1	Regression Weights	43
2.5.2	Response Determinism	43
2.6	Discussion	44
2.7	Publication Note	47
3	An Extension and Clinical Application of the SIMPLE Model of Memory to Item Repetition in Free Recall	48
3.1	Introduction	49
3.2	Behavioral Data	51
3.3	An Extension of the SIMPLE Model	54
3.3.1	Model Description	56
3.3.2	Application to Clinical Data	59
3.4	Results	61
3.4.1	Descriptive Adequacy	61
3.4.2	Modeling Results	63
3.4.3	Alternative Memory Probing via Temporal Context	64
3.5	Discussion	68
3.6	Publication Note	70
	Conclusion	71
	References	74
	Appendix A Model-Based Inference of Pairwise Similarities	82

LIST OF FIGURES

	Page	
1.1	Overview of three hypotheses for changes in odd-one-out choices. The left side shows the representation of “cow”, “elephant”, and “giraffe” in terms of a set of six features, where the area of circles corresponds to the weight given to each feature. The cow is chosen as the odd one out, as indicated by the encompassing dashed line. The right side of the figure shows how this choice changes according to the three hypotheses. At the top, the representational restructuring hypothesis involves only a subset of the features, and the elephant is chosen. In the middle, the attention change hypothesis involves different feature weights, and the giraffe is chosen. At the bottom, the noisy access hypothesis involves degraded access to the representation and more random choice.	9
1.2	The distributions of marginal posterior expectations of z_k^t for all features, within each FAST stage.	17
1.3	The prior (yellow) and posterior (blue) distributions for the feature transition base-rate ψ for all transitions between successive FAST stages.	18
1.4	The posterior mean and 95% credible intervals for the average attention weight for the physical, abstract, and thematic feature types across FAST stages.	21
1.5	Posterior distributions for response determinism γ and recency bias β for each FAST stage.	24
1.6	Posterior predictive analysis of the descriptive adequacy of the restructured representation, attention change, noisy access, and basic model for all FAST stages.	25
2.1	First-order transition probabilities for the animal free recall task. Animals are positioned according to a multi-dimensional scaling (MDS) plot calculated from the model-based similarity data, where more similar animals appear closer together. First-order transitions are represented as red arrows, where higher transition probabilities are represented by thicker arrows and lower transition probabilities are represented by thinner arrows. As impairment increases, transitions become less consistent with animal similarity.	34

2.2	Posterior distributions of regression weights α for each FAST stage. The similarity between items plays a large role in the order of free recall for the FAST 1 & 2 group only. The number of encounters with an animal name as well as the emotional valence of the animal name have relatively less influence. The weights for the other item-dependent measures are near zero for all groups.	40
2.3	Posterior distributions of response determinism γ for each FAST stage. Response determinism is consistent with probability matching for the FAST 1 & 2 group, but the other groups respond in a way that is less consistent with the model.	42
2.4	Posterior distributions of regression weights α in the components-of-similarity model for each FAST stage. The similarity of the abstract features between animal names seems to be a strong predictor of the order of free recall for the FAST 1 & 2 group only. Similar to the previous model, the number of encounters with an animal name as well as the emotional valence of the animal name have relatively less influence. The weights for the other item-dependent measures are near zero for all groups.	44
2.5	Posterior distributions of response determinism γ in the components-of-similarity model for each FAST stage. Response determinism is consistent with probability matching for the FAST 1 & 2 group, but the other groups respond in a way that is less consistent with the model.	46
3.1	The probability of item recall as a function of their encounters. The <i>left panel</i> shows that there is no evidence for primacy effects for any group. An animal name that appeared for the first time early in the triadic comparison task was no more likely to be recalled as an animal name that appeared for the first time later in the task. The <i>middle panel</i> shows that there is evidence of recency effects for all groups. An animal that appeared for the last time later on in the triadic comparison task was more likely to be recalled than an animal that appeared for the last time earlier in the task. The <i>right panel</i> shows that the more often an item was encountered during the study phase, and chosen as the odd one out, the more likely that animal name was to be recalled.	53
3.2	An illustration of the triadic comparison task. The task begins on the left side, and each tick on the timeline represents an encounter with an animal name. In this example, the blue tick marks represent all of the traces in which the word “beaver” was encountered. Note that beaver was chosen as the odd one out in the first and the last triad in which it appeared. After the 60 second retention interval, a surprise free recall task is presented.	55

3.3	The role of the temporal distinctiveness λ and temporal weight ω parameters in the extended SIMPLE model. The left side of the figure represents SIMPLE’s logarithmically compressed temporal dimension. Parameter λ represents the acuity with which a person is able to access the memory traces representing encounters with animal names. The right side of the figure is a multidimensional scaling representation of the semantic relationships between the animal names, with more similar animals located closer together. Parameter ω represents the relative weight given to temporal versus semantic similarity as memory cues for recall.	58
3.4	Modeling results. The <i>left panel</i> shows the joint and marginal posterior distributions for temporal distinctiveness μ_λ and temporal weight μ_ω . These posteriors indicate that as impairment increased, the temporal distinctiveness of memory traces decreased and people relied more on cues of temporal similarity than semantic similarity. The <i>top-right panel</i> shows that people used a relatively high threshold μ_τ and a noise μ_β consistent with gradual change from low to high recall probabilities. The <i>bottom right panel</i> compares the model-described probabilities of recall to the observed probability of recall, and indicates the model is descriptively adequate.	62
3.5	Modeling results of the alternative context-probing model. The <i>left panel</i> shows the joint and marginal posterior distributions for temporal distinctiveness μ_λ and temporal weight μ_ω . These posteriors indicate that as impairment increases, the temporal distinctiveness of memory traces decreases and people rely more on cues of temporal similarity than semantic similarity. The <i>top right panel</i> shows that people again had a relatively high threshold μ_τ and a noise μ_β consistent with a gradual change from low to high recall probabilities, but the joint posterior now shows a trade-off between these two parameters. The <i>bottom right panel</i> compares the model-described probabilities of recall to the probability of recall in the data, and indicates the model is descriptively adequate except for over-predicting the recall of infrequently recalled items. .	66

LIST OF TABLES

	Page
1.1 A description of the number of assessments, patients, and identifying characteristics for each FAST stage in the data set.	12
2.1 Identifying characteristics for each FAST stage, the number of participants in each stage, and descriptive statistics for the number of words correctly recalled.	33
2.2 Bayes factors for the comparison of each regression weight across adjacent FAST stages. Bayes factors in favor of a difference are in bold font; those in favor of sameness are in italic font. Inconclusive Bayes factors are colored gray and in standard typeface.	41
2.3 Bayes factors for the comparison of each regression weight in the components-of-similarity model across adjacent FAST stages. Bayes factors in favor of a difference are in bold font; those in favor of sameness are in italic font. Inconclusive Bayes factors are colored gray and in standard typeface.	45

ACKNOWLEDGMENTS

First, thank you to my advisor, Michael Lee. Thanks for taking a chance on me, for being a wonderful mentor, and for being a great friend. Thank you to my doctoral committee members, Mark Steyvers and Aaron Bornstein for their helpful feedback and guidance. Thank you to the other faculty who have supported me throughout my graduate program: Ramesh Srinivasan, Nadia Chernyak, Dan Nation, Jeff Rouder, and Barbara Sarnecka. Thanks also to Clara Schultheiss and MyLee Ryan-Glass for helping me navigate the UC Irvine administration.

This research was supported by funding from Embic Corporation. I would like to thank Dennis Fortier, Junko Hara, Tushar Mangrola, and especially Rod Shankle and Jason Bock for their help in providing the data and their contribution to the research. I also thank members of the Bayesian Cognitive Modeling lab at UC Irvine for helpful discussions.

Thanks to all of the wonderful friends I have made during my time in Irvine: Rena Goldstein, Emily Sumner, Adriana Chávez De la Peña, Jaime Islas Farias, Nidhi Banavar, Lauren Montgomery, Manuel Villarreal Ulloa, Jeff Coon, Paulina Silva, Alex Etz, Beth Baribault, and Helen Braithwaite.

Thank you also to my oldest friends, for supporting me not just for the last five years, but for the last few decades: Suzan Choy, Joelene Merritt, Courtney Johnson, and Leilani Kellogg.

Thank you to my parents, who have always supported my education.

And finally, thank you to my husband, Joachim, who believed in me even when I stopped believing in myself.

VITA

Holly A. Westfall

EDUCATION

Doctor of Philosophy in Cognitive Sciences University of California, Irvine	2023 <i>Irvine, CA</i>
Master of Science in Cognitive Sciences University of California, Irvine	2021 <i>Irvine, CA</i>
Master of Arts in Psychology University of South Florida	2013 <i>Tampa, FL</i>
Bachelor of Science in Psychology University of Florida	2007 <i>Gainesville, FL</i>

RESEARCH EXPERIENCE

Graduate Research Assistant University of California, Irvine	2018–2022 <i>Irvine, CA</i>
Center for Evidence-Based Corrections, Statistician University of California, Irvine	2016–2018 <i>Irvine, CA</i>
Graduate Research Assistant University of South Florida	2010–2015 <i>Tampa, FL</i>
Undergraduate Research Assistant University of Florida	2005–2006 <i>Gainesville, FL</i>

TEACHING EXPERIENCE

Teaching Assistant University of California, Irvine	2018–2022 <i>Irvine, CA</i>
Course Assistant University of South Florida	2015 <i>Tampa, FL</i>
Course Instructor University of South Florida	2013–2014 <i>Tampa, FL</i>
Teaching Assistant University of South Florida	2010–2012, 2014 <i>Tampa, FL</i>

REFEREED JOURNAL PUBLICATIONS

- Model-based wisdom of the crowd for sequential decision-making tasks** 2021
Cognitive Science
- Using the weighted Kendall's distance to analyze rank data in psychology** 2021
Tutorials in Quantitative Methods for Psychology
- A model-based analysis of the impairment of semantic memory** 2021
Psychonomic Bulletin & Review
- Ensemble coding of memory strength in recognition tests** 2019
Memory & Cognition
- The list-length effect does not discriminate between models of recognition memory** 2015
Journal of Memory and Language
- Visual search enhances subsequent mnemonic search** 2013
Memory & Cognition

REFEREED CONFERENCE PUBLICATIONS

- A model of free recall for multiple encounters of semantically-related stimuli with an application to understanding cognitive impairment** July 2022
Proceedings of the 44th Annual Conference of the Cognitive Science Society
- A model-based analysis of changes in the semantic structure of free recall due to cognitive impairment** July 2021
Proceedings of the 43rd Annual Conference of the Cognitive Science Society

AWARDS & HONORS

- Indow Fellowship for Research Excellence** 2022
University of California, Irvine
- Outstanding Scholar Award** 2022
University of California, Irvine
- Yellott Scholar Award Honorable Mention** 2019
University of California, Irvine

ABSTRACT OF THE DISSERTATION

Using Memory Models to Understand Cognitive Impairment Associated with Alzheimer's Disease

By

Holly A. Westfall

Doctor of Philosophy in Cognitive Sciences

University of California, Irvine, 2023

Professor Michael D. Lee, Chair

It is well established that Alzheimer's disease leads to a decline in both episodic and semantic memory. My research focuses on semantic memory and uses cognitive models to try to understand the nature of changes in memory as the disease progresses. Applying cognitive models, I am able to make inferences about underlying memory processes that cannot be measured directly from behavioral data.

In Chapter 1, I evaluate theories of the change in *semantic memory* caused by Alzheimer's disease. Many such theories characterize the cognitive decline associated with Alzheimer's disease as an elaborate combination of systematic changes to underlying mental representations and to attention processes. I conclude that a simpler theory of Alzheimer's disease as interfering with the patient's ability to retrieve information stored in memory can account for the data just as well as more complex theories.

In Chapter 2, I investigate changes in *episodic memory* associated with Alzheimer's disease. I specifically focus on the order of free recall among Alzheimer's patients and develop a regression model that includes semantic similarity and various word-level characteristics as predictors. I demonstrate that the semantic similarity between items and the number of times an item is encountered are potentially important predictors of free recall output.

In Chapter 3, I explore how semantic memory can influence episodic memory recall. I extend the SIMPLE model of memory and apply it to three groups of patients including healthy controls, people with mild cognitive impairment, and people with Alzheimer's dementia. I compare and contrast the relative influence of temporal similarity and semantic similarity on free recall performance within the three patient groups. As cognitive impairment increased, the ability to distinguish between items that were temporally similar decreased, as did the use of semantic similarity as a recall cue.

Collectively, the work in this dissertation shows that cognitive models of memory can provide a more complete understanding of the changes in semantic memory that occur with cognitive impairment.

INTRODUCTION

Alzheimer’s disease is a devastating neurological disorder that is caused by the accumulation of abnormal proteins and the degeneration of neurons. The disease is progressive, and these brain changes can further lead to brain inflammation, tissue atrophy, and eventually, death (Alzheimer’s Association, 2023). In Alzheimer’s disease, the parts of the brain responsible for learning and memory are damaged early on in the course of the disease. Accordingly, memory decline is one of the first noticeable signs of Alzheimer’s disease (McKhann et al., 2011). Currently, there is no cure for Alzheimer’s disease, but various drug treatments and cognitive interventions can help slow progress (Yiannopoulou & Papageorgiou, 2020). For this reason, early diagnosis and intervention is key.

The diagnosis and monitoring of cognitive impairment associated with Alzheimer’s disease is done with a number of standard neuropsychological tests. Such tests include the Boston Naming Task (BNT: Goodglass et al., 1983), the Controlled Oral Word Association Test (COWAT: Benton et al., 1983), the California Verbal Learning Test (CVLT-II: Delis et al., 2000), and the Consortium to Establish a Registry for Alzheimer’s Disease Neuropsychology Assessment Battery (CERAD-NAB: Welsh et al., 1994). These tests include semantic memory tests such as identifying line drawings (i.e., word retrieval), semantic fluency tasks (e.g., “name all of the animals you can think of”), and phonemic fluency (e.g., “name all of the words you can think of that start with the letter ‘s’ ”). Episodic memory tasks include immediate free recall, delayed free recall, cued recall, and recognition memory tasks. How-

ever, the amount of information that can be gained from standard analyses of these types of results is limited. For example, the number of line drawings correctly identified or the number of words correctly recalled from a 10-item list is not a particularly sensitive measure of cognitive impairment.

One approach to improve measurement sensitivity is to use principles of cognitive modeling to measure latent variables not directly observable from behavioral data. By creating a computational account of learning and memory, we can potentially improve the diagnosis, treatment, and monitoring of cognitive impairment associated with Alzheimer’s disease. Advances in computational psychiatry have shown data-driven computational models may be effective in improving patient outcomes (see Huys et al., 2021, for a review). Computational models can be used to bridge neuropsychological foundations to more abstract concepts, like learning and memory (e.g., Maia & Frank, 2011). By quantifying behavioral data within a mechanistic framework, we may be able to distinguish the behaviors and symptoms of those with different levels of cognitive impairment.

This dissertation focuses on one specific neuropsychological test, the Mild Cognitive Impairment Screen (MCIS: Shankle et al., 2009). This screen includes an odd-one-out triadic comparison of animal names as one task and an unexpected free recall of those animal names as another task. The odd-one-out task draws stimuli from a pool of 21 animal names: antelope, beaver, camel, cat, chimpanzee, chipmunk, cow, deer, dog, elephant, giraffe, goat, gorilla, horse, lion, monkey, rabbit, rat, sheep, tiger, and zebra. For each triad of animal names, the participant must choose which animal is least like the other two. For example, if presented with the words “cow”, “elephant”, and “giraffe”, a person might choose “cow” as the odd one out. After a delay, in which participants complete other unrelated tasks, there is an unexpected free recall task of these animal names. The instructions are to try to recall as many of the animal names as possible, in any order.

In the first chapter of this dissertation, I investigate semantic memory for animal names

as inferred from the odd-one-out choices of the triadic comparison task. I compare three competing hypotheses of the change in semantic memory associated with Alzheimer’s disease by developing a cognitive process model of the odd-one-out comparison task. This model includes a common-features measure of similarity between animals, in which animals are considered more similar when they share more features. The choice behavior is assumed to follow a Luce choice rule in which the odd-one-out choice is based on the relative similarities of the three animals in the triad. A person will choose animal A as the odd-one-out, if they judge animals B and C to be more similar to each other than either animal B or C is to animal A . After testing the models that formalize the three hypotheses, I conclude that as symptoms of cognitive impairment increase, access to semantic representation is impaired, but there is no evidence that the structure of the representation changes. Additionally, I identify a decision bias for Alzheimer’s patients to simply respond with the last animal name in the triad.

Following the animal triadic comparison task, people are asked to complete a surprise free recall task of those animal names: a test of episodic memory. In Chapter 2, I develop a regression-type model to try to determine various predictors of the order of free recall. This model also incorporates a Luce choice rule, because one way to think of free recall is as a multinomial choice task. I then incorporate item similarity as one of the predictors, the number of encounters with a word (i.e., how often that particular animal name was chosen as the odd one out), as well as other common word-level features that are typically controlled in a free recall task. These features include word frequency, word length, age of acquisition, and emotional valence. I find that the similarity between the items and the number of encounters are the most important factors in determining the free recall output order.

We learned in Chapter 2 that the semantic similarity and the number of encounters with an item are potentially important influences on whether an item is recalled. In Chapter 3 I extend an existing model that can potentially incorporate these findings. The Scale-

Independent Memory, Perception, and Learning model (SIMPLE: Brown et al., 2007) allows for a two-dimensional representation of memory that incorporates both the order of learning and the semantic similarity between the items. In a standard free recall task, items are presented one at a time, and then after a delay, are recalled. In the MCIS, the animal names are encountered as part of a triadic comparison task. Items are presented in groups of three and are repeated a variable number of times, depending on how many times an animal is chosen as the odd-one-out. The standard version of SIMPLE does not incorporate item repetition, but I extend the model to allow for this by storing every encounter with an item as a separate memory trace. I conclude that as cognitive impairment increases, people find words that were learned closer together in time to be more easily confused. Additionally, while healthy people have a tendency to put roughly equal weight on temporal information and semantic information as memory cues, people with Alzheimer's disease tend to use temporal (i.e., episodic information) to the exclusion of semantic information.

These conclusions could not have been reached by examining recall accuracy by itself; cognitive models were necessary to measure latent variables, such as the reliance on episodic versus semantic memory cues. I hope that research such as this will contribute to computational psychiatry.

Chapter 1

A Model-Based Analysis of the Impairment of Semantic Memory

Abstract

We use cognitive models to evaluate three theories of the change in semantic memory caused by Alzheimer's disease. We use data from 14,096 clinical assessments of 3602 Alzheimer's patients and their caregivers. Each patient completed a semantic memory task involving the odd-one-out comparison of animal names. Each patient was also independently evaluated to determine their level of impairment. Our cognitive models assume a feature-based representation of the animals and odd-one-out choice probabilities based on common-feature similarities. We find no evidence for the restructured representation hypothesis, which claims that impairment causes changes in the features used to represent stimuli. We also find no evidence for the attention change hypothesis, which claims that impairment causes greater attention to be given to concrete features at the expense of more abstract features. We do find evidence for the noisy access hypothesis, which claims that odd-one-out choices become

less determined by semantic similarity and more prone to the simple response strategy of choosing the last option. We conclude that the noisy access hypothesis provides a simple account of odd-one-out choice behavior throughout the progression of Alzheimer’s disease. More elaborate theories involving changes to underlying mental representations and attention processes need to provide evidence they are superior to the noisy access account.

1.1 Introduction

Alzheimer’s disease is a neurodegenerative disease that, according to the World Health Organization (2020), is the leading cause of dementia, with about 10 million new cases worldwide each year. Memory decline is one of the first noticeable symptoms of Alzheimer’s disease (McKhann et al., 2011), and many cognitive measures used to diagnose and monitor this disease rely on tests of episodic memory, such as tests of free and cued recall (Buschke, 1984; Grober & Buschke, 1987) and autobiographical memory (Kopelman et al., 1989; Levine et al., 2002). However, other diagnostic cognitive measures rely on semantic memory, such as category fluency tasks (Newcombe, 1969), verbal fluency tasks (Benton et al., 1983), and picture naming tasks (Kavé, 2005). People diagnosed with Alzheimer’s disease tend to do poorly on semantic memory tasks and have semantic memory deficits over and above those associated with normal cognitive aging (Nebes, 1989).

One way to study semantic memory is by testing people’s odd-one-out choices in a triadic comparison task. This task is part of some well-established Alzheimer’s testing batteries (e.g. Shankle et al., 2009; Trenkle et al., 2007). One of the most common tests involves odd-one-out choices between animal names. In this task, on every trial, people are verbally presented with three animal names and must choose the one that is the least like the other two. For example, out of the names “giraffe”, “elephant”, and “cow”, a person might choose “cow” as the odd one out. This task does not have correct answers but, from a person’s

odd-one-out choices, it is possible to make inferences about their semantic representation of the animals. If a person chooses “cow” as the odd one out, the implication is that they believe “giraffe” and “elephant” are more similar to each other than either is to “cow”.

Previous work has examined changes in stimulus similarity between healthy controls and Alzheimer’s patients using triadic comparisons of odors and colors (Razani et al., 2010), line drawings of common objects (Au et al., 2003), and pictures of animals and tools (Chan et al., 2001). While it is clear that perceptions of similarities do change, the underlying reason for the change is debated. This debate can be summarized in terms of three competing hypotheses. The first, which we call the *restructured representation* hypothesis, is that the underlying semantic representation has changed with the onset of Alzheimer’s disease. In an influential body of work, Chan and colleagues (e.g. Chan, Butters, Paulsen, et al., 1993; Chan, Butters, Salmon, & McGuire, 1993; Chan et al., 1995) argue that patients with Alzheimer’s disease have semantic networks where the associations between items differ from those of elderly healthy controls. One way to quantify people’s latent semantic networks is by using scaling and clustering techniques, such as multidimensional scaling (MDS: Shepard, 1980). In an MDS analysis, the similarity within a set of items is represented as distance between points in k -dimensional space, so that items that are judged to be more similar are located closer together in the space. In elderly healthy controls, an MDS analysis showed a smaller distance between similar animals (e.g., cat and dog) and a larger distance between less similar animals (e.g., cat and sheep). In contrast, for Alzheimer’s patients, the typical clustering of similar items became more diffuse, and MDS analyses revealed an increased distance between similar animals and decreased distance between less similar animals (Chan, Butters, Paulsen, et al., 1993).

A related hypothesis, which we call the *attention change* hypothesis, is that Alzheimer’s patients use different information about the animals to make odd-one-out choices, when compared to healthy controls. In particular, Chan and colleagues argue that Alzheimer’s

patients have a tendency to focus on concrete over abstract information. For example, elderly healthy controls are more likely to make similarity judgments of animals based on domesticity as compared to predation or size. In contrast, Alzheimer’s patients are more likely to make similarity judgments based on size (Chan, Butters, Salmon, & McGuire, 1993; Chan et al., 1995). The attention change hypothesis can be viewed as a special case of the restructured representation hypothesis. Both involve changes in the way people represent stimuli, but the attention change hypothesis assumes that the underlying representation remains the same, and it is only the selective attention to the components of that representation that are impacted by impairment.

A final hypothesis, which we call the *noisy access* hypothesis, is that differences in the triadic comparison task are not due to a change in the underlying semantic representation. Instead, the atypical response pattern seen in Alzheimer’s patients is due to an increasing loss of regularity in their choice behavior, rather than any fundamental representational change. The hypothesis is simply that Alzheimer’s patients have trouble accessing the information their semantic representations provide (Nebes & Brady, 1990). Adopting this position, some researchers (Elvevåg & Storms, 2003; Storms et al., 2003; Voorspoels et al., 2014; White et al., 2014) argue that the conclusions drawn about semantic reorganization from clustering and scaling techniques such as MDS may need to be tempered. Their argument is that differences in MDS output indicate only that the elderly healthy controls and Alzheimer’s patient groups are different in *some* way, but not *how* they are different. It could be that the patient data are simply noisier than the control data, and this noise could have a number of causes, including problems in accessing representations, a failure to understand instructions, or differing response strategies.

Figure 1.1 provides a conceptual overview of how these three hypotheses relate to understanding odd-one-out choices. The example involves choosing between the animals “cow”, “elephant”, and “giraffe”. The left side of the figure shows a cognitively healthy person




Figure 1.1: Overview of three hypotheses for changes in odd-one-out choices. The left side shows the representation of “cow”, “elephant”, and “giraffe” in terms of a set of six features, where the area of circles corresponds to the weight given to each feature. The cow is chosen as the odd one out, as indicated by the encompassing dashed line. The right side of the figure shows how this choice changes according to the three hypotheses. At the top, the representational restructuring hypothesis involves only a subset of the features, and the elephant is chosen. In the middle, the attention change hypothesis involves different feature weights, and the giraffe is chosen. At the bottom, the noisy access hypothesis involves degraded access to the representation and more random choice.

choosing “cow” by relying on a feature-based representation. Each of the animals is represented in terms of whether or not they are fat, are a draught animal, have spots, have fur, come from Africa, or are commonly found in zoos. Different features receive different levels of attention, and the cognitively healthy person is shown as paying the most attention to the African and zoo features. The similarity between each pair of animals is assumed to correspond to how many features they have in common, with the features weighted by their level of attention. For the cognitively healthy person, elephant and giraffe are the most similar, and so cow is chosen as the odd one out.

The remainder of Figure 1.1 shows how this choice can change according to the three hypotheses. Under the restructured representation hypothesis, the features involved in representing animals are different. In the example presented, the African and zoo features are no longer used. Based on the four features that remain, the cow and giraffe are the most similar, and the elephant is chosen as the odd one out. Under the attention change hypothesis, all of the features continue to be used, but the patterns of attention are different. In the example, the fat feature gains attention, making cow and elephant the most similar, and leading to giraffe being chosen as the odd one out. Finally, under the noisy access hypothesis, no aspect of the representation is changed. Instead, access to the information becomes less precise, and the odd-one-out choice becomes less based on similarities and closer to random responding.

The goal of this article is to evaluate these three hypotheses as accounts of the change in odd-one-out performance caused by the progression of Alzheimer’s disease. In the next section, we detail clinical data that assess patients with different levels of impairment. We then introduce a basic cognitive model of odd-one-out choice behavior based on simple assumptions about how stimuli are represented and decisions are made. The model is then extended to capture the specific assumptions of the restructured representation, attention change, and noisy access hypotheses, and each extension is evaluated against the clinical data. We find no evidence for restructured representation and no evidence for the sorts of

changes in attention that have previously been proposed. Instead, we find that the noisy access model provides a good account of the changes in triadic choice behavior. We conclude by discussing the implications and limitations of these findings.

1.2 Behavioral Data

Our behavioral data come from the animal triadic comparison task of the Mild Cognitive Impairment Screen (MCIS: Shankle et al., 2009) at a cognitive disorders clinic, as part of routine cognitive assessment of patients and their caregivers. The task draws from a pool of 21 animals: antelope, beaver, camel, cat, chimpanzee, chipmunk, cow, deer, dog, elephant, giraffe, goat, gorilla, horse, lion, monkey, rabbit, rat, sheep, tiger, and zebra.¹ For each specific assessment, nine animals are chosen from the pool, and each is presented in a triad with every other animal over a total of twelve trials. The triads are presented in accordance with a λ -2 balanced incomplete block design (Burton & Nerlove, 1976). On each trial, patients are verbally presented with the animal names and have to respond by choosing the animal that is least similar to the other two.

The data include 14,096 assessments of 3602 patients and their caregivers, (52% female, age range 16 –103 years, mean age 76 years). At the time of assessment, patients were classified by a physician with the Functional Assessment Staging Test (FAST; Reisberg, 1988). The FAST stages describe the severity of Alzheimer’s disease symptoms in terms of people’s ability to perform daily living tasks, such as managing finances, cooking, and grooming. Table 1.1 provides a summary of the number of unique patients and total number of assessments by FAST stage, and a description of each stage. Patients in stage 1 have no discernible deficits, and those in stage 2 have only a subjective functional deficit. These stages are considered cognitively normal. Patients in stage 3 have symptoms of mild cognitive

¹Animal names are chosen based on work in cultural anthropology that provides evidence for the universality of the semantic structure of the domain (Romney & Moore, 1998).

FAST stage	Description	Patients	Assessments
1 & 2	no deficit or subjective deficit	576	2032
3	objective deficit in complex tasks	908	3685
4	mild dementia evident in daily living tasks	1169	5352
5	moderate dementia	430	1313
6	moderately severe dementia	519	1714

Table 1.1: A description of the number of assessments, patients, and identifying characteristics for each FAST stage in the data set.

impairment, while patients in stages 4, 5, and 6 have mild, moderate, and moderately severe dementia, respectively. We did not include patients in FAST stage 7, who have severe dementia, because their cognitive function has degenerated to the point of an inability to understand simple instructions (Reisberg, 1988). The FAST stage assessments of impairment were made independent of memory test performance and so provide an external measure for grouping patients in order to study changes in the odd-one-out task behavior.

Some basic analyses of the data make clear that odd-one-out choices change as impairment progresses. As an example of change within a specific triad, consider the earlier cow, elephant, and giraffe example. For this triad, cow was chosen 53% of the time by patients in stages 1 and 2. Patients in stage 3, however, chose giraffe more often than cow, choosing giraffe 42% of the time and cow 37% of the time. Patients in stage 5 showed yet another pattern, choosing elephant as often as cow, with both accounting for 42% of all choices. The overall probability of different animals being chosen also often changes. For example, zebra was chosen as the odd one out 22% of the time by patients in stages 1 and 2, but 32% of the time by patients in stage 6, whereas rat was chosen 53% of the time by patients in stages 1 and 2, but only 43% of the time by patients in stage 6.

As a more thorough analysis of change, we examined the changes in similarities between pairs of animals across the FAST stages. Specifically, we compared the rates with which neither of a pair of animals was chosen, when both were presented. The Bayes factor (Lee

& Pope, 2006) favored a change in this rate for 5% of all pairs moving from stages 1 and 2 to stage 3, in 26% of pairs moving from stage 3 to 4, in 18% of pairs from stage 4 to 5, and in 7% of pairs from stage 5 to 6. These differences make it clear that there is widespread change in odd-one-out choices across the FAST stages.

1.3 Modeling Analysis

1.3.1 A Basic Model of Odd-One-Out Comparison

We begin the modeling analysis by developing a basic model of choice behavior in the odd-one-out tasks, which serves as the foundation for evaluating the competing hypotheses. The model has two core components: one for representing the stimuli and their similarities, and one for the decision-making processes that act on the similarities to produce choice probabilities.

Common-features Similarity

Previous research modeling the change in semantic memory with impairment (Chan, Butters, Paulsen, et al., 1993), including modeling focused specifically on animal odd-one-out comparisons (Lee et al., 2016), has relied on spatial representations like MDS. These representations assume that stimuli can be represented in terms of values on a small number of underlying psychological dimensions. An alternative to the dimension-based representational assumptions of MDS is to assume stimuli are represented in terms of features (Goldstone, 1999; Shepard, 1980). We think this is a more natural assumption for the representation of conceptual stimuli like animal names, and it aligns better with the set of hypotheses we aim to evaluate. In particular, the representational restructuring hypothesis assumes that com-

ponents of the representation are added or deleted with impairment. This hypothesis seems far more plausibly expressed in terms of the gain or loss of a few features, which typically apply to only a few stimuli each, rather than the gain or loss of an entire dimension, which always apply to every stimulus.

To represent the animal stimuli in terms of features, we assume that the similarity, s_{ab} , of animal a and animal b is equal to the sum of the weights of their shared features (Shepard & Arabie, 1979; Tversky, 1977).

$$s_{ab} = \sum_k w_k f_{ak} f_{bk}, \quad (1.1)$$

where f_{ak} is a binary indicator variable that determines whether or not animal a has feature k , and w_k represents the salience or weight of feature k . We used a truncated Gaussian² for the priors on the feature weights,

$$w_k \sim \text{Gaussian}(0, 1) \text{ T}(0,). \quad (1.2)$$

This common-features model makes animals similar only to the extent that they share features. Animals do not become more similar by both not having a feature. There is good evidence this is usually a reasonable assumption (Navarro & Lee, 2004), and the common-features model is the basis of widely-used additive clustering and related methods in similarity data analysis (Shepard & Arabie, 1979; Navarro & Griffiths, 2008; Peterson et al., 2018).

Rather than using additive clustering methods to infer the features, we used the Leuven concept database as a set of possible features (De Deyne et al., 2008). This database includes a total 288 features for the domain of animals, and lists whether a set of animals has each

²We parameterize the Gaussian distribution in terms of the mean and precision, consistent with the JAGS software we use to implement the models. The precision is the reciprocal of the variance.

feature based on four independent raters. We adapted the database for our modeling goals by including animals in the MCIS not originally listed, and eliminating features that had identical patterns of presence or absence across the 21 animals. These changes resulted in a final set of 118 possible features.

Luce Choice Rule

The probability, π_a of choosing animal a as the odd one out in a triad of animals a , b , and c is determined by their relative similarities. Intuitively, it is more likely animal a will be chosen if animals b and c are the most similar to each other. This intuition can be formalized by the Luce (1959) choice rule, which defines the choice probabilities as

$$\begin{aligned}\pi_a &= \frac{s_{bc}}{s_{ab} + s_{ac} + s_{bc}} \\ \pi_b &= \frac{s_{ac}}{s_{ab} + s_{ac} + s_{bc}} \\ \pi_c &= \frac{s_{ab}}{s_{ab} + s_{ac} + s_{bc}}.\end{aligned}\tag{1.3}$$

Given these probabilities, the observed choice of participant i on trial t , which we denote y_{it} , is modeled as

$$y_{it} = \text{categorical}(\pi_a, \pi_b, \pi_c).\tag{1.4}$$

This basic model provides an account of how the features used to represent stimuli, and the attention weights for those features, combine to produce similarities, as well as an account of the decision-making processes by which the similarities produce choice probabilities.

1.3.2 Restructured Representation Analysis

The restructured representation hypothesis assumes that the features to which people attend may change as memory impairment progresses. To create a model consistent with this hypothesis, we extend the basic model in several ways. First, we assume that people in any FAST stage use only a subset of the available features. Following Zeigenfuse & Lee (2010), this assumption is implemented by introducing a latent binary indicator parameter z_k^t that determines whether feature k is considered in the similarity judgments made by people in stage t , so that

$$s_{ab}^t = \sum_k z_k^t w_k f_{ak} f_{bk}. \quad (1.5)$$

The feature-inclusion parameters are given the prior $z_k^t \sim \text{Bernoulli}(\phi_t)$ with a base-rate for each stage $\phi_t \sim \text{beta}(1, 5)$.

Changes in the features z_k across FAST stages would provide evidence in favor of the restructured representation hypothesis. To measure this evidence, we introduce a change process by which the features used to represent the animals can change across FAST stages. This is formalized by binary parameters τ_k^t that indicate whether the inclusion of feature k is the same or different between stage t and stage $t + 1$, such that

$$z_k^{t+1} = \begin{cases} z_k^t & \text{if } \tau_k^t = 0 \\ 1 - z_k^t & \text{if } \tau_k^t = 1. \end{cases} \quad (1.6)$$

The feature-change parameters are given a prior $\tau_k^t \sim \text{Bernoulli}(\psi^t)$ with a base-rate $\psi^t \sim \text{uniform}(0, 1)$ for the transition from stage t to stage $t + 1$.

We implemented the restructured representation model, and all of the models considered in this article, as graphical models in JAGS (Plummer, 2003). JAGS provides a high-level




Figure 1.2: The distributions of marginal posterior expectations of z_k^t for all features, within each FAST stage.

scripting language for implementing probabilistic cognitive models that allows for computational Bayesian analysis using Markov-chain Monte Carlo sampling methods.

Results

Figure 1.2 shows the marginal posterior expectations of the feature-inclusion parameters z_k^t . Most features have posterior probabilities of inclusion close to 0 or 1. Only for a very few features is the inference uncertain. This provides good evidence that participants use about 45 of the 118 features in all of the stages. What remains to be determined is whether these 45 features are the same across the stages.

Figure 1.3 shows, in blue, the posterior distributions of the base-rates ψ for each of the transitions between successive FAST stages. One way to understand this result is in model-




Figure 1.3: The prior (yellow) and posterior (blue) distributions for the feature transition base-rate ψ for all transitions between successive FAST stages.

selection terms, comparing a “null” model that assumes a base-rate of change less than 1% against an alternative model that allows for greater change. Comparing the prior and posterior densities in the interval $0 < \phi_k < 0.01$ allows the Bayes factors between these two models to be estimated using the Savage-Dickey method (Wetzels et al., 2010). These Bayes factors are about 55, 6, 28, and 23 for the four transitions, all favoring the null model that assumes there is negligible change. The other interpretation of the results in Figure 1.3 is in parameter-estimation terms, treating the posterior distributions as measuring the extent of change that is assumed to exist. From this perspective, it is clear that the probability of a feature changing—either being added to a representation, or dropped from a representation—across FAST stages is very small, and almost certainly below 5% in every case.

From either perspective, the modeling results provide clear evidence against the restructured representation hypothesis. While participants use a subset of the possible features to make odd-one-out decisions, they use very close to the same subset in all of the FAST stages.

1.3.3 Attention Change Analysis

The attention change hypothesis assumes that the features used to make odd-one-out choices do not change with impairment, but the weights given to the features do change. The best theoretically-developed version of this hypothesis is provided by Chan and colleagues (Chan, Butters, Salmon, & McGuire, 1993; Chan et al., 1995), who argue that one type of features, called physical features, are given more attention as impairment progresses, while other types of features, called thematic and abstract, are given less attention. Physical features are those related to the animal’s appearance, such as “is fat”, “is specked”, and “has horns”. Abstract features are those related to the animal’s behavior, such as “can swim”, “eats nuts”, and “crawls up trees”. Thematic features are those describing the animal’s role, such as “is a pet”, “is popular among children”, and “is a cartoon figure”. We classified all of our

118 possible features into these three types, using two independent judges to determine the classifications, and a third judge to resolve disagreements.

To create a model consistent with this specific attention change hypothesis, we extend the basic model to include an account of how the weights of the different feature types change across FAST stage. This is accomplished by a hierarchical extension that assumes the weight of each feature is sampled from an over-arching Gaussian distribution that depends on both its type and the stage. Specifically, the feature weight w_k^t in FAST stage t , for a feature k is given by

$$w_k^t \sim \text{Gaussian}(m_{\kappa(k)}t + c_{\kappa(k)}, 1/\sigma_{\kappa(k)}^2) \text{T}(0,) . \quad (1.7)$$

The function $\kappa(\cdot)$ assigns each feature to its appropriate type, so that $\kappa(k) = 1$ means feature k is a physical feature, $\kappa(k) = 2$ means it is an abstract feature, and $\kappa(k) = 3$ means it is a thematic feature. The parameters $m_1, m_2, m_3 \sim \text{Gaussian}(0, 1)$ are slopes for the physical, abstract, and thematic feature types respectively, measuring how the average attention to features of that type increases or decreases with changes in the FAST stage. Similarly, the parameters $c_1, c_2, c_3 \sim \text{Gaussian}(1, 1) \text{T}(0,)$ are intercepts that measure the absolute level of attention. Finally, the parameters $\sigma_1, \sigma_2, \sigma_3 \sim \text{Gaussian}(1, 1/2^2) \text{T}(0,)$ are standard deviations that measure the heterogeneity in the attention weights between different features of the same type in the same FAST stage.

These hierarchical assumptions allow every feature to have its own attention weight in every FAST stage, but provide a measure of the average attention given to the physical, abstract, and thematic feature types. Critically, the model also provides a measure, via the slope parameters, of the change in the mean feature weights for each type as the FAST stages progress. The attention change hypothesis can be expressed concisely in terms of these slopes: attention to physical features increases, so $m_1 > 0$, while attention to abstract and




Figure 1.4: The posterior mean and 95% credible intervals for the average attention weight for the physical, abstract, and thematic feature types across FAST stages.

thematic features decreases, so $m_2, m_3 < 0$.

Results

Figure 1.4 presents the average attention weights for each feature type by FAST stage. The error bars represent 95% credible intervals. The average attention weight for each of the physical, abstract, and thematic features all decrease with impairment. This is inconsistent with the pattern of change predicted by the attention change hypothesis. To quantify this result, we calculated the Bayes factor comparing the specific inequality-constrained attention change hypothesis $m_1 > 0, m_2, m_3 < 0$ against the alternative hypothesis without any constraints. The Bayes factor is about 35 in favor of the alternative hypothesis. Collectively, these results provide strong evidence against the attention change hypothesis. While there is a change in the average attention given to different types of features, there is not a systematic increase in the attention given to more concrete features at the expense of more abstract ones.

1.3.4 Noisy Access Analysis

The noisy access hypothesis assumes that all participants attend to the same features with the same attention weights, but that the ease of access decreases with the progression of impairment. To create a model consistent with this hypothesis we focus on the decision-making assumptions. In particular, we extend the basic model to allow for response determinism and a recency bias in the choice rule.

Response determinism measures how closely choice behavior is determined by the underlying similarities, and can be modeled by extending the Luce choice rule through exponentiation (Lee et al., 2016). Pairwise similarities are now raised to a power, in the form s_{ab}^γ , and different values of the parameter γ determine how consistently people make odd-one-out choices. If $\gamma = 1$, the original Luce choice rule is maintained. If γ decreases towards zero, the probability of choosing animals a , b , and c are all reduced towards $\frac{1}{3}$, and choice behavior becomes more random. As γ increases to values greater than one, responses become more deterministic, and the animal that is least similar to the other two will be more consistently chosen.

The second decision-making extension is to include a recency bias, corresponding to the strategy of choosing the last animal name presented in the sequential verbal presentation. An increase in the recency bias could be explained by a deficit to working memory typically seen in Alzheimer’s patients (see Huntley & Howard, 2010, for a review). To allow for this possibility we assume simply that that last animal is favored by a bias measured by the parameter β . Combining the response determinism and recency bias extensions gives a set of choice probabilities defined as

$$\begin{aligned}
\pi_a &= \frac{1 - \beta}{2} \left(\frac{s_{bc}^\gamma}{s_{ac}^\gamma + s_{bc}^\gamma + s_{ab}^\gamma} \right) \\
\pi_b &= \frac{1 - \beta}{2} \left(\frac{s_{ac}^\gamma}{s_{ac}^\gamma + s_{bc}^\gamma + s_{ab}^\gamma} \right) \\
\pi_c &= \beta \left(\frac{s_{ab}^\gamma}{s_{ac}^\gamma + s_{bc}^\gamma + s_{ab}^\gamma} \right).
\end{aligned} \tag{1.8}$$

We assume the prior $\gamma \sim \text{gamma}(2, 1)$ for response determinism. This distribution has a mode at 1, corresponding to the original Luce choice rule and probability matching, allows for greater values corresponding to deterministic choice behavior, and allows for lesser values corresponding to more random choice behavior. We assume a prior $\beta \sim \text{uniform}(1/3, 1)$ for recency bias. This choice of prior allows for the possibility of unbiased choices corresponding to $\beta = \frac{1}{3}$, but captures our assumption that any bias will be in favor of the last item.

Results

Figure 1.5 presents the posterior distributions for the response bias and response determinism parameters by FAST stage. It is clear that response determinism decreases as impairment progresses across the stages. The Bayes factors testing whether γ is the same or different across successive stages are all greater than 100 in favor of there being a difference. It is also clear that the recency bias generally increases across the FAST stages. The Bayes factor is only 2 in favor of a difference between stages 1 and 2 compared to stage 3, but is greater than 25 for all of the other comparisons. Interestingly, there is a significant recency bias even for cognitively healthy participants in stages 1 and 2, since the Bayes factor is more than 100 in favor of β being greater than $\frac{1}{3}$.




Figure 1.5: Posterior distributions for response determinism γ and recency bias β for each FAST stage.

These results are very consistent with the noisy access hypothesis. The natural interpretation is that changes in choice behavior can be explained by progressively impaired use of the underlying similarities between animals. The choices made are less well determined by the similarities as impairment increases, and there is a compensatory greater reliance on the simple strategy of choosing the last animal name presented.

1.3.5 Modeling Conclusions

Our final analysis of all of four models—the basic model, the restructured representation model, the attention change model, and the noisy access model—involves their descriptive adequacy. Being able to describe the data is a basic requirement for a model to be useful, and an important part of their evaluation. To measure descriptive adequacy, we considered how well model choice probabilities matched behavior. On each trial, a model generates probabilities π_a , π_b , and π_c for the three alternatives. If these probabilities describe the data well, they should match the frequency with which the alternatives are actually chosen. For

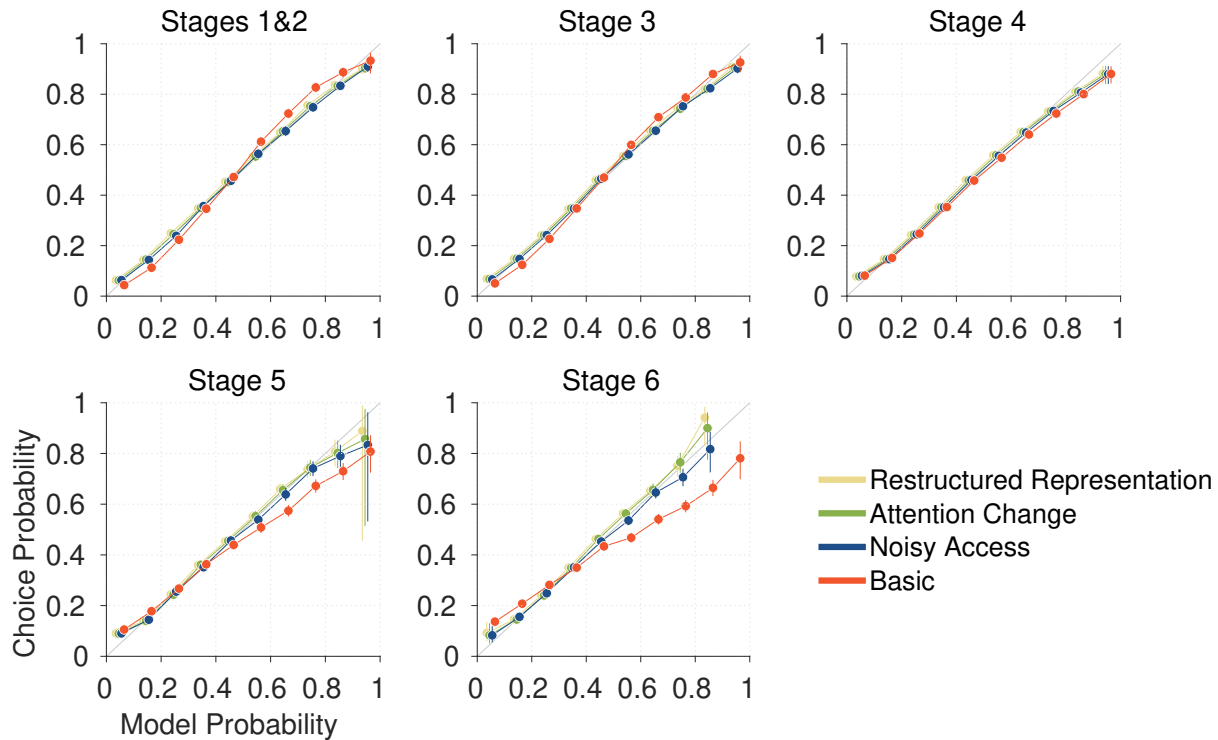


Figure 1.6: Posterior predictive analysis of the descriptive adequacy of the restructured representation, attention change, noisy access, and basic model for all FAST stages.

each of the models we binned the choice probabilities in increments of 0.1, and calculated the frequency with which the alternative was chosen.

The results of this analysis are shown in Figure 1.6. Panels correspond to FAST stages, the lines show the mean observed choice proportions, and error bars represent 95% credible intervals. The basic model is the worst performed. It fails to describe the data in FAST stages 5 and 6, and also has systematic deviations in the other stages. This finding is not surprising, since the basic model assumes there is no change in odd-one-out behavior across the stages, but our basic data analysis found clear evidence of change. The three other models all describe the data well, and are all very similar to each other.

These results show that the restructured representation model, the attention change model, and the noisy access model all meet the basic requirement of descriptive adequacy. Accordingly, our overall evaluation of the models focuses on *how* they achieve descriptive adequacy,

which amounts to asking how they explain the changes in odd-one-out choice behavior. It is clear that the restructured representation model does not capture the changes in the data by changing the features used to represent animals. The attention change model also fails to explain the data in a way consistent with the guiding theory, which requires increases in attention to physical features at the expense of abstract ones. The noisy access model, in contrast, achieves descriptive adequacy in exactly the way predicted by the theory. Response determinism decreases with impairment and is accompanied by an increasing use of the simple response strategy of choosing the last item. Thus, our conclusion is that the noisy access model provides the best account of the data. It is the psychologically and statistically simplest model, and accounts for the data in an interpretable way consistent with its theoretical assumptions.

1.4 Discussion

Using behavior in odd-one-out tasks, we conducted a model-based comparison of three hypotheses previously proposed as accounts of changes in semantic memory performance caused by Alzheimer’s disease. We found no evidence for the idea that semantic representations themselves fundamentally change. We also did not find any evidence for the idea that the attention given to different types of features changes systematically, with more concrete physical properties becoming more prominent at the expense of more abstract features. Instead, we found that the differentiation between these types of features decreases as impairment progresses, suggesting a gradual loss of acuity in the use of the representations rather than a structured shift in attention. Consistent with this interpretation, we did find evidence favoring the idea that access to semantic information becomes noisier as impairment worsens. Choice behavior became less consistently linked to the underlying similarities between stimuli, and participants became more likely simply to repeat the last option. It is possible that

patients in higher FAST stages do not fully understand the instructions of the task—and fail to do a search of semantic memory—but do understand that repeating an animal name back to the clinician is an acceptable response.

We do not believe these results eliminate the possibility that Alzheimer’s disease does cause basic systematic changes in semantic memory. The odd-one-out comparison task provides only one window onto semantic memory, and the clinical data we used considered only basic-level exemplars from the natural kind of animals. This limits the theories that can be tested. Some researchers, for example, have proposed a “bottom-up” degeneration of semantic memory (Henry et al., 2004; Martin & Fedio, 1983; Tröster et al., 1989), where Alzheimer’s patients have more difficulty generating specific exemplars (e.g., broccoli, orange) than categories (e.g., vegetables, fruits). There is also evidence that Alzheimer’s patients have a category-specific semantic memory deficit for living things (Chan et al., 2001; Whatmough & Chertkow, 2002). Our data cannot test either of these possibilities directly.

In addition to the limitations of our data, each of our models made a number of strong assumptions. Our models do not take into account any properties of the animal names themselves, such as their age of acquisition, word length, or word frequency. Word frequency may be particularly important, since previous research has shown that Alzheimer’s patients do not show typical effects of word frequency in tests of episodic memory (Balota et al., 2002; Wilson et al., 1983). Another assumption is that, while we think our use of feature-based representations is appropriate, previous authors have often relied on dimensional representations (Chan, Butters, Paulsen, et al., 1993; Lee et al., 2016). The attention change and noisy access hypotheses probably could be formalized using dimensional representations, which would provide an alternative approach to evaluation. In addition, within the feature-based framework, it is possible to specify other models consistent with the broad idea of attention change. Our model focuses on one very specific influential proposal, but it is possible other systematic changes in attention do occur. What we do believe is that

the noisy access model provides a psychologically and statistically simple and compelling account of our data, and serves as a theoretical safeguard that any more elaborate theory of change in semantic memory must outperform.

One avenue for stronger testing of the adequacy of the noisy access hypothesis, as compared to the possibility of more fundamental changes in representations, is to apply our models to different populations where there are clear expectations about whether and how semantic memory should change. Healthy aging provides a context in which we might expect noisy access to continue to provide a better account than restructured representation or attention change, and thus presents an opportunity for replication. Early child development, in contrast, provides a context in which we *would* expect to observe fundamental changes in the way children represent stimuli. A goal of future work is to apply our models to children's performance on the odd-one-out comparison task.

Understanding how Alzheimer's disease affects semantic memory is a basic theoretical question with important societal implications. We have developed a model-based approach that is capable of expressing and evaluating competing theoretical accounts, and have demonstrated how Bayesian methods allow the models to be applied to a large real-world clinical data set. We hope that future work continues to expand and refine the models, and provides insights into how people's semantic knowledge is impacted by memory impairment.

1.5 Publication Note

This chapter was previously published as Westfall, H. A., & Lee, M. D. (2021). A model-based analysis of the impairment of semantic memory. *Psychonomic Bulletin & Review*, 25(8), 1484-1494. doi: <https://doi.org/10.3758/s13423-020-01875-9>

Chapter 2

A Model-Based Analysis of Changes in the Semantic Structure of Free Recall Due to Cognitive Impairment

Abstract

Alzheimer's disease leads to a decline in both episodic and semantic memory. Free recall tasks are commonly used in assessments designed to diagnose and monitor cognitive impairment, but tend to focus only on episodic memory. Our goal is to understand the influence of semantic memory on the sequence of free recall in a clinical data set. We develop a cognitive process model that incorporates the influence of semantic similarity and other stimulus properties on the order of free recall. The model also incorporates a decision process based on the Luce choice rule, allowing for different levels of response determinism. We apply the model to a real-world data set including free recall data from 2392 Alzheimer's patients and their caregivers. We find that the semantic similarity between items and the number

of encounters with an item are potentially important predictors of the order of free recall output.

2.1 Introduction

Alzheimer’s disease causes changes in memory that are distinct from normal aging. In cognitively healthy adults, memory for past events (episodic memory) tends to decline over time, but memory for facts (semantic memory) typically remains intact (see Balota et al., 2000, for a review). Memory deficits over and above those seen in healthy aging can be a sign of cognitive impairment (Nebes & Brady, 1990). Patients with Alzheimer’s disease display deficits in both episodic and semantic memory early in the disease, with deficits in memory increasing with the severity of cognitive impairment (Mortamais et al., 2017).

A commonly used task in memory assessment is the free recall task, which is widely used to study memory and decision processes. In this task, a participant is presented with a list of items and then – immediately or after some delay – must recall those items, in any order. There are several common ways to analyze free recall data. Techniques such as serial position analysis, in which effects of primacy and recency are typically observed (Murdock, 1962), rely on the order of items presented at study. Conditional response probabilities and lag-recency effects (Howard & Kahana, 1999) are concerned with the temporal relationship of successively recalled items at test. Simple recall accuracy is a quick and easy way of analyzing memory performance, but it is not sensitive to any information about the presentation order of items or the order of recall output.

Often, in a recall task, stimuli are controlled on various dimensions, such as word frequency, word length, age of acquisition, emotional valence, or semantic similarity. These kinds of controls are implemented because semantic memory can influence episodic memory. For

example, people tend to recall words in semantically-related clusters, both between categories (Bousfield, 1953) and within categories (Romney et al., 1993). Because of this, stimuli in free recall tasks tend to be controlled for their semantic relationships, to prevent semantic associations being used in recall. These controls, however, limit the ability of free recall behavior to inform our understanding of a basic feature of human memory, which is to organize stimuli and represent meaning. Understanding the semantic relationships between successively recalled items can provide valuable insight into memory and decision processes, including how semantic memory is affected by impairment in Alzheimer’s patients (Ribeiro et al., 2007).

The goal of this paper is to model the sequence of free recall in a clinical data set in an effort to understand how semantic similarity and other features of stimuli guide the recall process. In the next section, we describe the clinical data set and an initial analysis of first-order transition probabilities. We then describe a cognitive model of free recall output based on the similarity to other recalled words, as well as several other properties of the words themselves. We find that the semantic similarity between items strongly influences the order of free recall, but only for the healthy control group. We also observe that the number of encounters with an item may potentially be an important predictor of free recall order, especially for the group with mild cognitive impairment. In the following section, we develop a similar model that splits the similarity measure into three components. This allows us to explore whether a specific type of feature may be driving the semantic similarity between animal names. Finally, we discuss the implications and limitations of our findings.

2.2 Behavioral Data

The data were collected as part of a routine assessment of Alzheimer’s patients and their caregivers at a clinic specializing in neurodegenerative disorders. All participants completed

the Mild Cognitive Impairment Screen (MCIS: Shankle et al., 2009), which is used to help diagnose and monitor cognitive impairment. This screen includes an odd-one-out comparison of animal names task and an unexpected free recall task of those animal names. The odd-one-out comparison task draws stimuli from a pool of 21 animal names: antelope, beaver, camel, cat, chimpanzee, chipmunk, cow, deer, dog, elephant, giraffe, goat, gorilla, horse, lion, monkey, rabbit, rat, sheep, tiger, and zebra. For each triad of animal names, the participant must choose which animal is least like the other two. For example, if presented with the words “cow”, “elephant”, and “giraffe”, a person might choose “cow” as the odd one out. The clinician does not offer any feedback after each choice, as there is no correct answer for this task. In accordance with a λ -2 balanced, incomplete block design (Burton & Nerlove, 1976), nine animal names are drawn from the pool for each participant, and each of the selected animals is presented verbally in a triad with every other animal over the course of 12 trials. After a delay, in which participants complete other unrelated tasks, there is an unexpected free recall task of these animal names. The instructions are to try to recall as many of the animal names as possible, in any order.

The data set includes assessments from 2392 participants (52% female, 48% male, age range 16–101 years, mean age 74 years). At the time of assessment, all participants were also classified using the Functional Assessment Staging Test (FAST: Reisberg, 1988). The FAST assessment is an evaluation of a person’s ability to perform Instrumental Activities of Daily Living (IADLs: Lawton & Brody, 1969), such as cooking, cleaning, and managing finances, as well as Activities of Daily living (ADLs: Katz et al., 1963), such as dressing, bathing, and grooming. Participants in FAST stages 1 and 2 have either no functional deficit or only a subjective deficit and are considered to be cognitively healthy for the purposes of this analysis. Those in stage 3 have mild cognitive impairment (MCI) and are beginning to show an objective deficit in accomplishing more complex tasks. Participants in stages 4, 5, and 6, have been diagnosed with mild, moderate, and moderately severe Alzheimer’s dementia, respectively. The FAST assessment is made independently of the MCIS, and so

Table 2.1: Identifying characteristics for each FAST stage, the number of participants in each stage, and descriptive statistics for the number of words correctly recalled.

Stage	Description	n	M	SD
1 & 2	no deficit or subjective deficit	518	6.8	1.6
3	objective deficit in complex tasks	782	5.6	1.9
4	mild dementia evident in IADLs	770	4.1	2.0
5	moderate dementia	152	3.5	2.0
6	moderately severe dementia	170	3.0	1.8

it provides a way to group participants by impairment, in order to study changes in free recall. A summary of the number of participants grouped into each FAST stage is presented in Table 2.1.

In an initial analysis of the data, we wanted to determine whether the semantic similarity of the animal names influenced the order of free recall output, and if so, whether this influence varied with cognitive impairment. First we inferred the pairwise semantic similarity of the 21 animal names from the odd-one-out comparison choices of the FAST 1 & 2 group using the model described by Westfall & Lee (2021).¹ We chose this particular model-based measure of similarity, because it allows us to account for non-similarity-based influences on odd-one-out choice behavior, such as recency bias and response determinism. These pairwise similarity values were standardized to range from 0 to 1 with each animal having maximal similarity with itself. Further details are provided in the Appendix. Then we calculated first-order transition probabilities to measure the probability of the next recalled word given the most recently recalled word, similar to semantic conditional response probabilities (Howard & Kahana, 2002b).

Using the model-based similarities, we performed a multi-dimensional scaling analysis (MDS: Shepard, 1980). The resulting plot depicts more similar animals closer together in space and

¹There are many ways to calculate similarity, including free association measures (De Deyne et al., 2019; Nelson et al., 2004), latent semantic analysis (LSA: Landauer & Dumais, 1997), vector cosine similarity via an algorithm such as word2vec (Mikolov et al., 2013), and even from the odd-one-out comparison task itself (Romney et al., 1993).

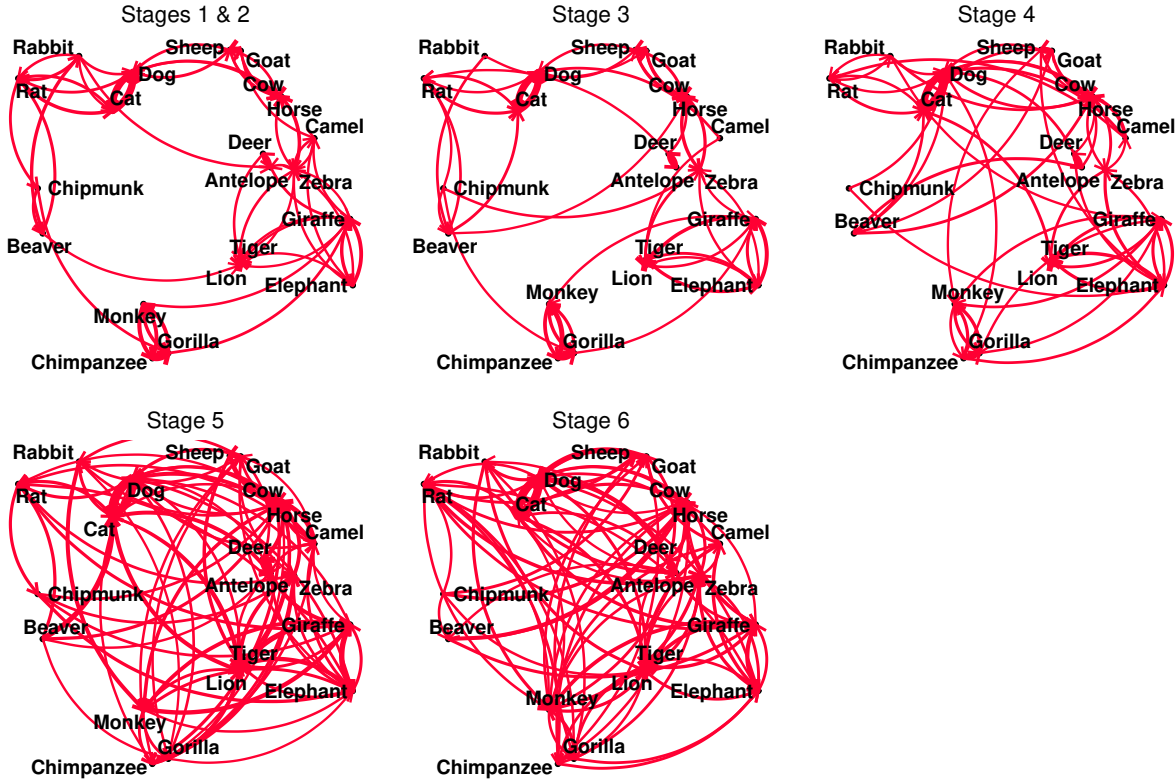


Figure 2.1: First-order transition probabilities for the animal free recall task. Animals are positioned according to a multi-dimensional scaling (MDS) plot calculated from the model-based similarity data, where more similar animals appear closer together. First-order transitions are represented as red arrows, where higher transition probabilities are represented by thicker arrows and lower transition probabilities are represented by thinner arrows. As impairment increases, transitions become less consistent with animal similarity.

less similar animals farther apart. In Figure 2.1, for each stage, all 21 animals are presented in accordance with the results of the MDS analysis. Transition probabilities are represented by arrows of varying width. Higher transition probabilities are represented by thicker arrows and lower transition probabilities are represented by thinner arrows. Transitions occurring between animals that are located near each other are indicative of similarity-based memory search, while transitions between animals that are far apart indicate some other process, inconsistent with similarity. In stages 1 & 2, most of the transitions occur between animals that are semantically similar to each other. Transitions between similar animals such as “gorilla”, “chimpanzee”, and “monkey” are common, while transitions between dissimilar animals, such as “camel” and “dog”, are much less common. In stages 3 and 4, many of

the transitions occur between semantically similar animals, but there is a relatively greater number of transitions that link dissimilar animals. In stages 5 and 6, there appear to be as many transitions between similar animals as dissimilar animals. From these plots, it is clear that as FAST stage increases, transitions become less consistent with the similarity of the animal names.

2.3 Model Description

A description of the model follows and is broken into two parts: a regression function, allowing for the influence of between- and within- stimuli factors on memory, and a decision-making process that produces choice probabilities. We applied the model to each FAST stage separately, rather than attempting to estimate a linear or otherwise monotonic progression across FAST stages. While our purpose was to measure change across FAST stage, we did not want to incorporate such a strong assumption into the model itself.

The model was written in R (R Core Team, 2018) and implemented in JAGS (Plummer, 2003) via the rjags package (Plummer, 2016). JAGS provides a high-level scripting language for implementing probabilistic cognitive models that allows for computational Bayesian analysis using Markov-chain Monte Carlo sampling methods (Lee & Wagenmakers, 2013). The results are based on three chains of 10,000 posterior samples collected after 1000 discarded burn-in samples, with a thinning factor of 10. We assessed convergence of chains by visual inspection and through the \hat{R} statistic (Brooks & Gelman, 1998).

2.3.1 Regression Function

The regression function contains weights of theoretical interest α , which express the influence of various psychological features on recall probabilities. Specifically, we assume that the

memory strength Q allowing for the recall of animal j on trial k given that animal i has just been recalled on trial $k - 1$ (i.e., $y_{k-1} = i$) is a function of both the similarity of animal j to animal i , and specific features of word j :

$$Q_k(j|y_{k-1} = i) = \alpha_1 \times \text{similarity}_{ij} + \alpha_2 \times \text{encounters}_j \\ + \alpha_3 \times \text{frequency}_j - \alpha_4 \times \text{length}_j - \alpha_5 \times \text{aoa}_j - \alpha_6 \times \text{valence}_j. \quad (2.1)$$

For each animal name, we obtained the word frequency, word length, age of acquisition, and emotional valence from The English Lexicon Project (Balota et al., 2007). Word frequency corresponds to the log-transformed Hyperspace Analogue to Language (HAL: Lund & Burgess, 1996) frequency norms. Word length is defined as the number of syllables. We also include the number of encounters, which is the number of times a participant chose an animal name in the odd-one-out comparison task, as a word-level feature. All of the predictors were re-scaled to be in the range of 0 to 1.

We assume each FAST stage has its own regression weights, and they are given a Dirichlet prior. This choice of prior ensures that the weights sum to one and allows us to interpret the weights as the relative importance of each predictor on memory:

$$\boldsymbol{\alpha} \sim \text{Dirichlet}(1, 1, 1, 1, 1, 1). \quad (2.2)$$

2.3.2 Luce Choice Rule

Since there is a discrete number of possible animal names, the free recall task can be thought of as a multinomial choice task. We use the Luce choice rule (Luce, 1959) extended to incorporate response determinism, which allows us to assign a probability to each of these

response options. According to the Luce choice rule, choice probabilities S are defined as:

$$S(j) = \frac{\exp(\gamma Q(j))}{\sum_m \exp(\gamma Q(m))}, \quad (2.3)$$

where j refers to one of the recalled animal names, the summation in the denominator is over all recalled animal names, and the function Q is defined in Equation 2.1 (and the conditioning on the previous observation y_{k-1} is now implicit). Finally, the response determinism parameter γ determines the degree to which a participant chooses a response consistent with the regression function Q . If $\gamma = 0$, Q is ignored, and the probability of choosing option j becomes $\frac{1}{m}$. If $\gamma = 1$, decisions become consistent with probability matching. As γ increases, the decision becomes more deterministic based on Q , and a participant will eventually always pick the option with the highest Q . In this way, the determinism parameter can be interpreted as the consistency of decisions. Following Lee et al. (2016), we assume each FAST stage has its own response determinism, and assume a gamma prior:

$$\gamma \sim \text{gamma}(2, 1), \quad (2.4)$$

which has a mode corresponding to probability matching, but allows for higher and lower values.

Censoring repeated recall.

For the purposes of this analysis, we removed extra-list intrusions, but it is still possible for the participant to recall the same word repeatedly. Because this type of task error could be related to cognitive impairment, we extend our model to include a censoring component that captures the ability to inhibit repeated recalls.

The probability of choosing a response option on a particular trial is determined by the Luce

choice rule and a censoring index δ , which is a binary indicator that takes the value 1 for words that have not been previously recalled. This index can be pre-calculated and treated as observed data.

Whether a participant is in a state of censoring on a trial is determined by parameter z . If $z = 0$, the participant *is not* in a censoring state, and the probability of recalling alternative j is determined by the Luce choice rule. If $z = 1$, then the participant *is* in a censoring state, and the probability of recalling alternative j is equal to δ times the Luce choice rule:

$$P_k(y_k = j) = \begin{cases} S(j), & \text{if } z_k = 0 \\ \delta_j S(j), & \text{if } z_k = 1. \end{cases} \quad (2.5)$$

The parameter z has a Bernoulli prior with hyperparameter ϕ :

$$z_k \sim \text{Bernoulli}(\phi). \quad (2.6)$$

In our application of the model ϕ is a person-specific parameter that captures the individual's ability to censor previously recalled words. ϕ can range from 0 (no censoring) to 1 (perfect censoring), and is given a standard uniform prior:

$$\phi \sim \text{uniform}(0, 1). \quad (2.7)$$

2.4 Results

We have a clear expectation for response determinism. As impairment increases, free recall output should become less consistent with the memory of the items. In other words, response determinism should decrease as FAST stage increases. However, we are also interested in the relative influence of each of the regression predictors on the structure of free recall, and

whether they change across impairment.

2.4.1 Model Descriptive Adequacy

We quantified model fit by creating a confusion matrix for each FAST stage that compares the true word recall of the participants to the model-described word recall. From these matrices we were able to calculate the overall descriptive accuracy of the model, which was between 38% and 57% for each stage. A model that predicted outcomes completely randomly would have an accuracy of only $\frac{1}{21}$, or 5%. While far from perfect, this suggests that the model provides a reasonably accurate description of the data, and performs an order of magnitude better than a random chance model.

2.4.2 Regression Weights

The posterior distributions for the regression weights α of each of the predictors are presented in Figure 2.2. The regression weights for similarity are relatively high for the FAST 1 & 2 group, meaning that the similarity between items is particularly important in the sequence of free recall for this group. For the remaining groups, similarity seems to play a very small role in the order of free recall. The number of encounters with an animal name (i.e., the number of times an animal name was chosen as the odd-one-out) seems to be at least somewhat influential on recall order for all FAST groups. The number of encounters may be particularly important for the FAST 3 group, which is the group with MCI. Emotional valence has relatively less influence on recall order compared to similarity and number of encounters. The weights for the other item-dependent properties, word frequency, word length, and age of acquisition, are all very close to zero for all groups.

To quantify the change in the regression weights, if any, across FAST stage, we calculated

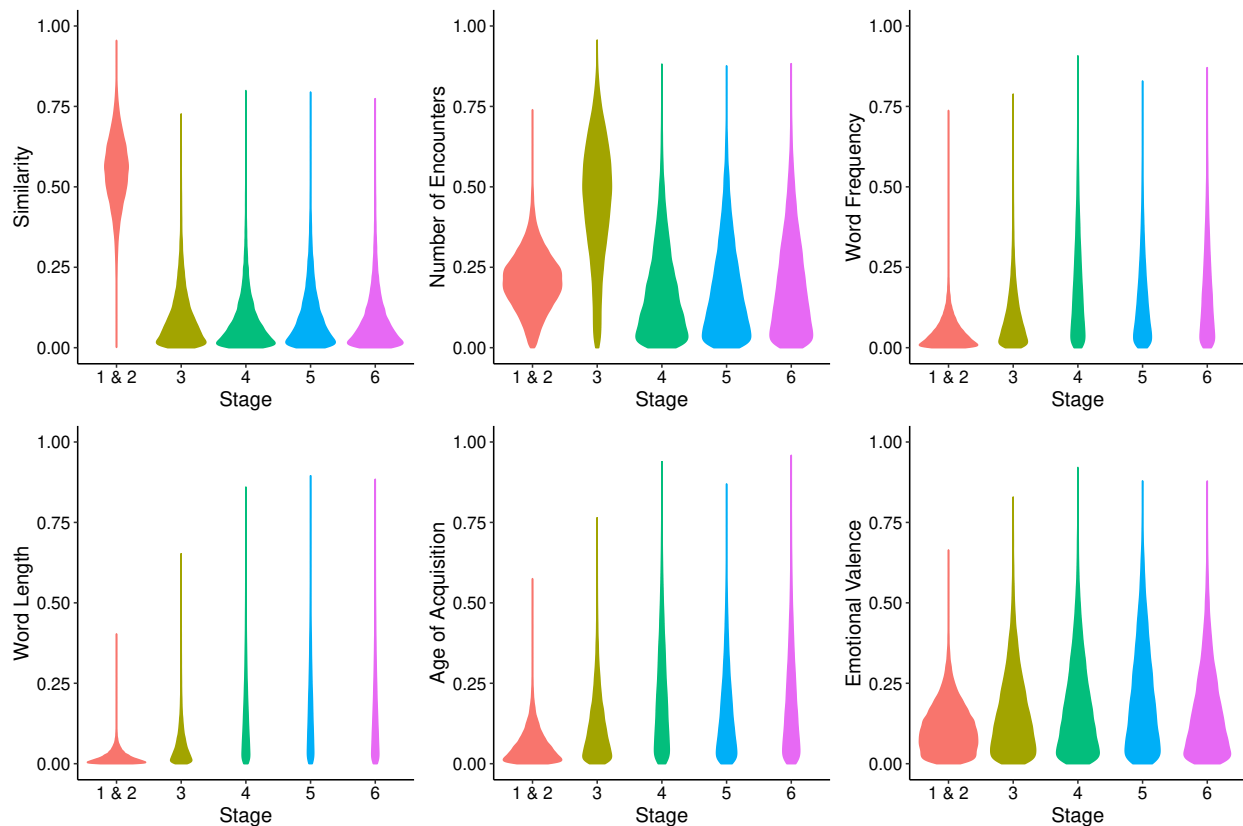


Figure 2.2: Posterior distributions of regression weights α for each FAST stage. The similarity between items plays a large role in the order of free recall for the FAST 1 & 2 group only. The number of encounters with an animal name as well as the emotional valence of the animal name have relatively less influence. The weights for the other item-dependent measures are near zero for all groups.

Bayes factors between adjacent stages for each of the regression weights. Our results are presented in Table 2.2. Bayes factors in favor of a difference are in bold font, while Bayes factors in favor of sameness are in *italic font*. With the exception of similarity, and perhaps the number of encounters, there is not a large change in the regression weights across FAST stage. It is clear that these two predictors, in particular, may have an influence on the order of free recall.

Table 2.2: Bayes factors for the comparison of each regression weight across adjacent FAST stages. Bayes factors in favor of a difference are in bold font; those in favor of sameness are in italic font. Inconclusive Bayes factors are colored gray and in standard typeface.

	1&2 vs. 3	3 vs. 4	4 vs. 5	5 vs. 6
Similarity	37.0	1.96	1.91	1.82
Encounters	3.24	3.40	1.08	1.22
Frequency	2.28	1.14	1.01	1.06
Length	<i>4.12</i>	1.51	1.03	1.03
AOA	1.80	1.02	1.22	1.08
Valence	1.14	1.03	1.25	1.35

2.4.3 Response Determinism

The posterior distributions for response determinism γ are presented in Figure 2.3. Response determinism for the FAST 1 & 2 group is near 1.0, while the distributions for the other groups are lower. Again, we computed Bayes factors to quantify any differences between adjacent FAST stages. In this case, the Bayes factor comparing FAST 1 & 2 to FAST 3 was 1.39 in favor of a difference. For the other comparisons, stage 3 vs. stage 4, stage 4 vs. stage 5, and stage 5 vs. stage 6, the Bayes factors in favor of sameness were 7.84, 6.45, and 4.39, respectively. These results suggest that people in the FAST 1 & 2 group respond in a way that is consistent with probability matching, while those in other groups do not.

2.5 Components-of-Similarity Model

In a previous publication (Westfall & Lee, 2021), we tested a specific hypothesis regarding semantic memory and the odd-one-out task. This hypothesis, which we referred to as the *attention change* hypothesis, posits that there is a systematic and predictable change in semantic memory for people with cognitive impairment associated with Alzheimer’s disease. The assumption is that memory for physical features of an animal should remain relatively salient for people with Alzheimer’s disease, while memory for more abstract or thematic




Figure 2.3: Posterior distributions of response determinism γ for each FAST stage. Response determinism is consistent with probability matching for the FAST 1 & 2 group, but the other groups respond in a way that is less consistent with the model.

features should decrease with impairment (Chan, Butters, Salmon, & McGuire, 1993; Chan et al., 1995). In our study of the odd-one-out choices, we did not find a preference for odd-one-out choices based on physical features at the expense of abstract or thematic features for people with cognitive impairment. While we did not find evidence in favor of this hypothesis for semantic memory accessed to make the odd-one-out choices, it is possible that different features categories may have differing influence on episodic memory recall.

In this section, we split the 118 features we used to calculate overall similarity and categorized each one as physical, abstract, or thematic. Physical features are concrete features of an animal related to its appearance such as, “is fluffy”, “is brown”, and “has short legs”. Abstract features are those related to the animal’s behavior. These features include, “digs holes”, “is carnivorous”, and “can swim”. Thematic features refer to the role the animal might play, and include features such as, “is a beast of burden”, “is a pet”, and “is a cartoon figure”. The only change to the regression model is that now there are three separate

predictors for similarity that correspond to physical, abstract, and thematic features. The Luce choice rule and the censoring component of the model are identical to those described above.

2.5.1 Regression Weights

The posterior distributions for the regression weights α of each of the predictors for the components-of-similarity model are presented in Figure 2.4. For similarity, it does appear that the type of feature – in this case, abstract features – makes a difference on item recall, but only for the FAST 1 & 2 group. For all other groups, the influence of similarity, of any feature type, is much lower. However, this result is not consistent with the *attention change* hypothesis, because we do not see an accompanying increase in the influence of physical features for people with MCI or Alzheimer’s disease. Similar to the previous model, the number of encounters with an item and emotional valence have relatively less influence on recall order, while the weights for the other item-dependent properties, word frequency, word length, and age of acquisition, are all very close to zero.

To quantify any change in the regression weights across FAST stage, we calculated Bayes factors between adjacent stages. Our results are presented in Table 2.3. Again, Bayes factors in favor of a difference are in bold font, while Bayes factors in favor of sameness are in italic font. With the exception of abstract similarity, regression weights do not tend to shift across FAST stage.

2.5.2 Response Determinism

The posterior distributions for response determinism γ are presented in Figure 2.5. Not surprisingly, we find similar results to the previous model for response determinism, where




Figure 2.4: Posterior distributions of regression weights α in the components-of-similarity model for each FAST stage. The similarity of the abstract features between animal names seems to be a strong predictor of the order of free recall for the FAST 1 & 2 group only. Similar to the previous model, the number of encounters with an animal name as well as the emotional valence of the animal name have relatively less influence. The weights for the other item-dependent measures are near zero for all groups.

the distribution for the FAST 1 & 2 group seems to be near 1.0, while the distributions for the other group are lower. Again, we computed Bayes factors to quantify any differences between adjacent FAST stages. The Bayes factor comparing FAST 1 & 2 to FAST 3 was 1.29 in favor of a difference. For the other comparisons, stage 3 vs. stage 4, stage 4 vs. stage 5, and stage 5 vs. stage 6, the Bayes factors in favor of the null were 11.2, 8.67, and 6.44, respectively. Again, these results suggest that only the FAST 1 & 2 group is responding in a way consistent with probability matching.

2.6 Discussion

We created a cognitive process model to try to understand the sequence of free recall output in a clinical data set. The model has two main components: a regression-type equation that

Table 2.3: Bayes factors for the comparison of each regression weight in the components-of-similarity model across adjacent FAST stages. Bayes factors in favor of a difference are in bold font; those in favor of sameness are in italic font. Inconclusive Bayes factors are colored gray and in standard typeface.

	1&2 vs. 3	3 vs. 4	4 vs. 5	5 vs. 6
Similarity - Physical	1.82	1.63	1.45	1.40
Similarity - Abstract	11.5	1.30	1.21	1.32
Similarity - Thematic	2.31	1.93	1.53	1.40
Encounters	2.46	2.63	1.22	1.29
Frequency	1.68	1.06	1.17	1.14
Length	2.76	1.17	1.04	1.11
AOA	1.35	1.15	1.22	1.19
Valence	1.02	1.24	1.31	1.23

describes the influence of within- and between-item factors, and a decision process based on the Luce choice rule. We found that the semantic similarity between animal names was an important predictor for the order of recall, but only for the FAST 1 & 2 group, i.e., the healthy control group. The number of encounters with an animal name may also be an important predictor of order of recall, especially for the FAST 3 group, i.e., the MCI group. Word frequency, word length, and age of acquisition seem to matter very little in terms of word choice for this particular set of stimuli. The fact that these word-specific characteristics had little influence on the order of recall is not surprising in this particular context. These animal names were chosen for this task to be similar on several dimensions.

We further investigated the effect of semantic similarity on recall by creating another model, the components-of-similarity model, that categorized the different animal features into three different types: physical, abstract, and thematic. We found that abstract features seemed to be driving the effect of item similarity on order of recall for the FAST 1 & 2 group. Again, we also found that the number of encounters with an animal name may be an important predictor of recall order, especially for the FAST 3 group.

For both models, we found that the posterior distributions for response determinism for




Figure 2.5: Posterior distributions of response determinism γ in the components-of-similarity model for each FAST stage. Response determinism is consistent with probability matching for the FAST 1 & 2 group, but the other groups respond in a way that is less consistent with the model.

the FAST 1 & 2 group was near 1.0, or probability matching. The posterior distributions for the other groups were lower. A decrease in response determinism means that decisions become less consistent with the regression model. However, Bayes factor comparisons between successive stages revealed that any differences were not large enough to make strong claims. Our use of a response determinism parameter is similar to the model described in Lee et al. (2016). We assume that all participants have the same underlying semantic representation, but the ability to access that information decreases as impairment increases (Nebes & Brady, 1990). In other words, cognitive impairment limits the ability to access stored semantic information in memory. This is essentially the same conclusion about the impacts of impairment reached by Westfall & Lee (2021) in our model-based analysis of the odd-one-choice behavior that we used to determine semantic similarity.

A major assumption of this model is that we know what items will be recalled. The model

is only concerned with the order of free recall output given the words that were recalled successfully. A complete account of the sequence of free recall should be able to predict the words that will be recalled, given the memory stimuli, as well as providing an account of first-word recall.

Nevertheless, the modeling results advance our understanding of the factors that influence the free recall of within-category items. In particular, the use of a cognitive model allows us to identify patterns and relationships not observable in standard data analysis. Measuring latent psychological parameters potentially provides a more precise measurement of semantic clustering than other more common behavioral methods, such as the California Verbal Learning Test (CVLT-II: Delis et al., 2000). The model allows us access to information that would otherwise be lost in a typical account of free recall accuracy, and helps us understand recall order when information on study order is unavailable or ambiguous. Most importantly, our modeling provides an insight into how semantic memory and the number of encounters with an item drives free recall, and how this interaction changes with memory impairment. This knowledge will be helpful in future model development.

2.7 Publication Note

A previous version of this chapter was published as Westfall, H. A., & Lee, M. D. (2021). A model-based analysis of changes in the semantic structure of free recall due to cognitive impairment. In T. Fitch, C. Lamm, H. Leder, & K. Teßmar-Raible (Eds.), *Proceedings of the 43rd Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.

Chapter 3

An Extension and Clinical Application of the SIMPLE Model of Memory to Item Repetition in Free Recall

Abstract

The surprise free recall of triadic comparisons, a task used in clinical settings, presents a unique analysis challenge for many memory models, because learning occurs incidentally and items are presented multiple times. To account for this design, we extended the SIMPLE (Brown et al., 2007) model of memory, which assumes to-be-remembered items are stored as separate logarithmically-compressed temporal traces. The ability to retrieve these traces depends on the acuity of memory probes and the semantic similarity between the items represented by the traces. We applied this model to a real-world clinical data set includ-

ing healthy controls, people with mild cognitive impairment, and people with Alzheimer’s dementia. We found that as memory impairment progressed, both the acuity with which temporal memory could be accessed and the use of semantic information decreased relative to temporal information.

3.1 Introduction

Within research and clinical settings, a common test of memory is the free recall task (Healey & Kahana, 2014). In its standard form, a list of items is presented to a participant, and they are asked to recall as many items as possible in any order, either immediately or after some delay. While the free recall task is a valuable tool for both basic and applied research, the design of the task does not always correspond to the way memory is used in real-world decision making. In everyday life, there are only some situations in which one needs to memorize a clear-cut list of items, either because the source material is not available or because there is high time pressure. The rote learning of material for scholastic examinations is an example of the first type of situation. Pilots recalling sequences of actions to take in an emergency is an example of the second type.

More often, learning occurs incidentally; people encounter information that they need to remember while accomplishing other unrelated tasks. As an example, consider how people remember which food items to buy when they shop. Sometimes, shopping may be based on recalling a previously-prepared shopping list of required items, consistent with the rote learning of a study list in a standard memory experiment. More often, people have to recall what grocery items to buy from their memories of cooking, during which ingredients are encountered repeatedly in loosely-structured sequences over many episodes. While preparing a meal, if someone looks for paprika in their pantry and sees that they have run out, they must simultaneously think of a substitute ingredient while also remembering to pick up more

paprika the next time they are at the store. Once at the store, finding paprika can be a memory cue itself for other spices that need to be replenished.

In this article, we examine how people recall the items they encounter in a triadic comparison task that is a component of the Mild Cognitive Impairment Screen (MCIS: Shankle et al., 2009) used to diagnose cognitive impairment associated with Alzheimer’s disease. During the triadic comparison task, animal names are presented multiple times in groups of three and people must choose the “odd one out”. People’s choices provide a measure of semantic memory that itself is useful for understanding impairment (Westfall & Lee, 2021). Although they are not told this before the triadic comparison task, people are later asked to recall the animal names. People’s ability to recall the animal names provides another measure of memory that is sensitive to both episodic memory, given the temporal sequence of encountering the animal names, and semantic memory, given the similarity relationships between the animals.

An important step towards understanding any cognitive capability is to develop models of people’s behavior. However, the design of the triadic-recall task creates complications that are not addressed by most traditional memory models. One model that does potentially have the required flexibility to model complicated patterns of stimulus encounters is the Scale-Independent Memory, Perception, and LEarning model (SIMPLE: Brown et al., 2007). When applied to standard study-test free recall tasks, SIMPLE assumes that each study item is stored as a separate memory trace, and these traces are logarithmically compressed along a dimension of time. This logarithmic compression captures common observations in free recall, such as primacy and recency. SIMPLE also allows for traces to differ on multiple dimensions, which allows for the additional influence of semantic similarity on the memory traces. The standard version of the SIMPLE model does not explicitly incorporate mechanisms such as rehearsal or item repetition. However, given the basic assumption that items are represented as memory traces, the model can be extended to store each repeated encounter with a study

item as a separate memory trace.

In this article, we develop an extended version of the SIMPLE model and demonstrate its ability to account for free recall behavior in the triadic-recall task by applying it to a real-world clinical data set. In the next section, we describe the data set and the details of the triadic comparison and free recall tasks. Then we describe our extension to the SIMPLE model and how we account for the design of the triadic-recall task. We used a hierarchical Bayesian implementation of the extended model to compare memory performance across three groups of people: healthy controls, patients with MCI, and patients with moderately severe Alzheimer’s dementia. We found that as impairment increased, not only did accuracy of recall decrease, but the temporal distinctiveness of memory traces decreased and people came to rely more on temporal context rather than the semantic similarity of the items as cues.

3.2 Behavioral Data

We used data from the MCIS administered at a clinic specializing in neurodegenerative disorders. The triadic comparison task uses 21 animal names as stimuli: antelope, beaver, camel, cat, chimpanzee, chipmunk, cow, deer, dog, elephant, giraffe, goat, gorilla, horse, lion, monkey, rabbit, rat, sheep, tiger, and zebra. Based on a balanced incomplete block design (Burton & Nerlove, 1976), nine animal names are drawn from the pool of 21 animal names for each test, and each of the selected animals is presented verbally in a triad with every other of the nine animals over the course of 12 trials. For each triad, the person must choose which animal is least like the other two. For example, someone presented with the words “cow”, “elephant”, and “giraffe”, might choose “cow” as the odd one out. There is no correct answer for this task, and so the clinician does not offer any feedback after each choice. The unexpected free recall of these animal names occurs after a delay during which

people complete other unrelated tasks. The free recall instructions are to try to recall as many of the animal names as possible, in any order.

The data we used come from a larger clinical data set which contains 398 tests completed by healthy controls, 3808 completed by people with MCI, and 1154 completed by people with Alzheimer's dementia. The control group shows no functional cognitive impairment, while patients in the MCI group are beginning to show objective deficits in accomplishing more complex tasks, such as managing finances (Reisberg, 1988). The patients with Alzheimer's disease have been diagnosed with moderately severe dementia and are beginning to show difficulty accomplishing tasks from the Activities of Daily Living (ADLs: Katz et al., 1963), such as dressing, bathing, and grooming.

Due to the potential influence of the surprise free recall task, we limited analysis to the first assessment completed by an individual. We included the data from all 95 unique subjects in the Healthy group and randomly sampled 95 tests from each of the MCI and Dementia groups for analysis. The 95 healthy controls (52% female, 48% male, mean age 61 years) recalled an average of 6.2 animals, with individuals recalling a minimum of 1 animal name and a maximum of all 9. The 95 MCI patients (59% female, 41% male, mean age 75 years) recalled an average of 4.7 animals, with a minimum of 0 and a maximum of 8 animal names. The 95 dementia patients (36% female, 64% male, mean age 80 years) recalled an average of 1.7 animal names with the full range of 0 to 9 animals recalled. The distribution of ages in this data set is a reflection of the true time course of Alzheimer's disease. In the real world, the older a person is, the more likely the onset of Alzheimer's disease, and once diagnosed, the symptoms increase in severity over time.

As a first analysis of the data, we looked for regularities commonly found in free recall, such as effects of primacy and recency. Primacy is an advantage to recall for items occurring at the beginning of a list, while recency is an advantage to recall for items presented at the end of a list (Murdock, 1962). The left panel of Figure 3.1 shows no evidence of a primacy




Figure 3.1: The probability of item recall as a function of their encounters. The *left panel* shows that there is no evidence for primacy effects for any group. An animal name that appeared for the first time early in the triadic comparison task was no more likely to be recalled as an animal name that appeared for the first time later in the task. The *middle panel* shows that there is evidence of recency effects for all groups. An animal that appeared for the last time later on in the triadic comparison task was more likely to be recalled than an animal that appeared for the last time earlier in the task. The *right panel* shows that the more often an item was encountered during the study phase, and chosen as the odd one out, the more likely that animal name was to be recalled.

effect for any of the three groups. An animal name may be encountered for the first time in triads 1 through 9, according to the test design. As each animal name appears in a total of 4 triads, it cannot appear for the first time in triads 10 through 12. Figure 3.1 shows that animal names appearing for the first time early in the triadic comparison task were no more likely to be recalled than animal names that appeared for the first time later in the task. We quantified this observation by performing a Bayesian regression analysis and calculating a Bayes factor comparing a null model with slope equal to 0 to an alternative model with non-zero slope. The Bayes factors in favor of the null model were 10, 1.5, and 41 for the Healthy, MCI, and Dementia groups, respectively. In the middle panel of Figure 3.1, there is evidence of a recency effect for all three groups. The Bayes factors in favor of a non-zero regression line slope were over 400 for Healthy and MCI groups and 12 for the Dementia group. These findings are consistent with other research involving incidental learning in which effects of recency but not primacy were found in surprise free recall (Marshall & Werder, 1972).

Another common finding in memory research is that the probability of item recall increases when that item is repeated during study (see Toppino & Gerbier, 2014, for a review). As people make odd-one-out judgments, they say their choice out loud, and choosing an animal as the odd one out may have some effect on whether that animal is later successfully recalled. Throughout the triadic comparison task, an animal name can be encountered a minimum of four (i.e., the animal name was only ever said by the clinician and was never chosen as the odd one out) and a maximum of eight times (i.e., the animal name was chosen as the odd one out in every triad in which it appeared). In the right panel of Figure 3.1 it is clear that the more study encounters someone had with an animal name, the more likely that animal name was to be recalled. Here the Bayes factors in favor of a non-zero regression line slope were over 1000 for all three groups, which is strong evidence for better recall for animal names that were chosen more often as the odd one out.

3.3 An Extension of the SIMPLE Model

The triadic comparison task acts as the study phase for the surprise free recall task. However, the design of this task is very different from the standard study phase of a typical free recall task. Words are presented in triads, and each word is repeated a total of four times so that it appears exactly once in a triad with every other word. The analysis summarized in Figure 3.1 suggests a model of people's free recall in this situation needs to be able to account for recency effects and allow for variable repetition of items at study.

As argued above, the SIMPLE model of memory may be able to handle these challenges. Each encounter with an animal name can be represented as a memory trace, including the temporal structure of a sequence of trials, each with three options followed by a response. SIMPLE can also represent elapsed time between the triadic comparison task and the free recall task. SIMPLE has four basic assumptions: First, each memory trace is represented in




Figure 3.2: An illustration of the triadic comparison task. The task begins on the left side, and each tick on the timeline represents an encounter with an animal name. In this example, the blue tick marks represent all of the traces in which the word “beaver” was encountered. Note that beaver was chosen as the odd one out in the first and the last triad in which it appeared. After the 60 second retention interval, a surprise free recall task is presented.

psychological space along a time dimension, and potentially other dimensions that relate to properties of the items. The time dimension is logarithmically compressed, which allows for SIMPLE to be applied to a wide range of timescales. Second, a memory trace is easier to retrieve to the extent that it is more easily discriminated from other traces within psychological space. Third, the temporal discriminability of memory traces from each other is a function of the ratio of their distances from the time of retrieval. Finally, the probability of retrieval of a specific memory trace is an increasing function of that trace’s discriminability over the total discriminability of the other traces in psychological space.

As an example, consider Figure 3.2, which represents the triadic comparison task as it is completed in real time. Each tick mark represents the timing of an encounter with an animal name, either as part of a triad read out loud by the clinician or the odd-one-out choice made by the participant. Here we assume, consistent with clinical practice, that animal names are read aloud at a rate of one per second and the person responds with their odd-one-out choice within two seconds. Figure 3.2 highlights the triads in which the word “beaver” appears, and the blue tick marks represent the traces in which this word is encountered. Following a

60 second retention interval, the person must recall as many of the animal names as possible, in any order.

Figure 3.2 focuses on temporal information about the encountered items, but it makes sense to include additional representational dimensions corresponding to the semantic similarities between the animals. In free recall tasks in particular, people tend to recall items in clusters of semantic similarity (Bousfield, 1953; Howard & Kahana, 2002b; Romney et al., 1993). To incorporate the effect of semantic similarity, we inferred the pairwise semantic similarity of the 21 animal names from the triadic comparison choices of the Healthy group using the model presented by Westfall & Lee (2021). These pairwise similarity values were standardized to range from 0 to 1 with each animal having maximal similarity with itself. Further details are provided in the Appendix. Given the addition of this semantic dimension to the psychological space, the discriminability of memory traces can depend on either their temporal nearness, their semantic similarity, or some combination of both. This corresponds to the assumption that, depending on the attention weight people give to the dimensions of its psychological space, animal names can be confused either because they were presented close in time, or because they represent similar animals.

3.3.1 Model Description

The temporal similarity of two memory traces is a function of their separation in time and the acuity with which they are accessed. The temporal distance η_{ij} between two memory traces i and j , encountered at times T_i and T_j relative to the time of retrieval is represented as

$$\eta_{ij} = \exp(-\lambda |\ln(T_i) - \ln(T_j)|), \quad (3.1)$$

which incorporates the key logarithmic compression assumption. The parameter $\lambda > 0$ is a temporal distinctiveness parameter, representing a person’s memory acuity. The temporal discriminability of an item i , with trace T_i , relative to a probe j , aiming to retrieve the item with trace T_j , is then calculated as

$$\pi_{ij}^T = \frac{\eta_{ij}}{\sum_{k=1}^n \eta_{ik}}. \quad (3.2)$$

The pairwise semantic similarity of the two traces, π_{ij}^S , is incorporated into the model via a weighting parameter ω , illustrated in Figure 3.3, that determines the relative importance of temporal versus semantic similarity to the overall evaluation of similarity of the two traces,

$$\pi_{ij} = \omega \pi_{ij}^T + (1 - \omega) \pi_{ij}^S. \quad (3.3)$$

The model transforms the retrieval probabilities via a logistic function,

$$\pi_{ij}^* = \frac{1}{1 + \exp[-\beta(\pi_{ij} - \tau)]}, \quad (3.4)$$

where π_{ij}^* represents the probability that probe j will yield trace i . This function serves as a mechanism to allow for study words to be omitted during recall. Parameter $\tau \in (0, 1)$ is the threshold value, above which a trace is successfully retrieved. Parameter $\beta > 0$ represents the scale, or the noisiness of the threshold value. A large β value indicates that all traces above the threshold are retrieved and all traces below the threshold are not, whereas a smaller β value indicates a more gradual transition from low to high retrieval probabilities. Following the correction noted by Lee & Pooley (2013), the probability that a trace i will be retrieved at least once is calculated as

$$\theta_i^T = 1 - \prod_{j=1}^n (1 - \pi_{ij}^*). \quad (3.5)$$

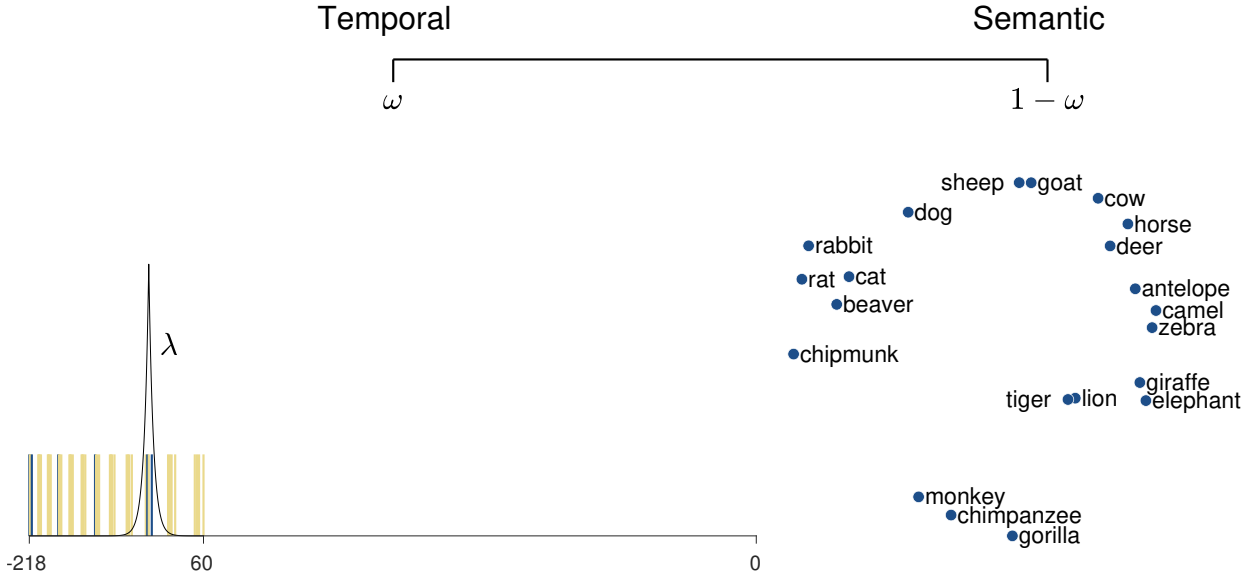


Figure 3.3: The role of the temporal distinctiveness λ and temporal weight ω parameters in the extended SIMPLE model. The left side of the figure represents SIMPLE’s logarithmically compressed temporal dimension. Parameter λ represents the acuity with which a person is able to access the memory traces representing encounters with animal names. The right side of the figure is a multidimensional scaling representation of the semantic relationships between the animal names, with more similar animals located closer together. Parameter ω represents the relative weight given to temporal versus semantic similarity as memory cues for recall.

Previous applications of SIMPLE have not considered the possibility of multiple traces representing the same item, as happens for the animal names in the triadic comparison task. This extension, however, is straightforward. The probability that animal a will be recalled is naturally measured as the probability that at least one trace representing that animal is retrieved. That probability is given by

$$\theta_a = 1 - \prod_{i=1}^n (1 - \theta_i^T)^{z_{ia}} \quad (3.6)$$

where $z_{ia} = 1$ if animal a occurred in trace i , and $z_{ia} = 0$ otherwise.

3.3.2 Application to Clinical Data

The recall data take the form $y_{pa} = 1$ if person p recalled animal a and $y_{pa} = 0$ if they did not. For each person, this probability is given by Equation 3.6 in the basic model above, based on individual-specific parameters for their temporal acuity λ_p , temporal versus semantic weighting ω_p , threshold noise β_p , and threshold τ_p . We assume that temporal acuity and the temporal versus semantic weight vary according to the level of impairment. Accordingly, we use a hierarchical model with group means μ_λ^g and μ_ω^g and standard deviations σ_λ^g and σ_ω^g , where g is the indicator for the Healthy, MCI, or Dementia group. Individual person parameter values are modeled as

$$\begin{aligned}\lambda_p &\sim \text{Gaussian}_+(\mu_\lambda^g, \frac{1}{(\sigma_\lambda^g)^2}) \\ \omega_p &\sim \text{Gaussian}_{(0,1)}(\mu_\omega^g, \frac{1}{(\sigma_\omega^g)^2}),\end{aligned}\tag{3.7}$$

depending on their group.¹

Based on our previous work (Westfall & Lee, 2021), we do not expect the structure of semantic memory to differ because of cognitive impairment. In particular, we expect the *access* to semantic information to differ by cognitive impairment, but we do not expect the semantic *distance* itself to change with impairment. Accordingly, we assume that the thresholding mechanism does not change with impairment, so that there are single means

¹Note that we parameterize the Gaussian distribution in terms of its mean and precision, consistent with the JAGS software we use to implement the models.

μ_β and μ_τ and standard deviations σ_β and σ_τ that apply to everybody, so that

$$\begin{aligned}\beta_p &\sim \text{Gaussian}_+(\mu_\beta, \frac{1}{\sigma_\beta^2}) \\ \tau_p &\sim \text{Gaussian}_{(0,1)}(\mu_\tau, \frac{1}{\sigma_\tau^2}).\end{aligned}\tag{3.8}$$

These modeling assumptions focus the explanation of changes in free recall performance on the acuity of memory and the use of semantic information, while still allowing for individual differences in the thresholding mechanism.

The triadic-recall task involves 48 memory traces. These assumptions specify the 48 T_i trace values as 218, 217, 216, 214, 204, 203, 202, 200, \dots , 64, 63, 62, 60, corresponding to seconds from recall. Consistent with previous applications of the SIMPLE model, we also assume that retrieval involves probing all 48 memory traces.

Our model is completed by placing priors on the parameters. Following the idea that priors should capture theoretical assumptions related to parameter values (Lee & Vanpaemel, 2018), we choose a uniform prior for the μ_λ^g parameters based on the timings of the traces. Specifically, we bound the value to be above the minimum required to ensure that the two temporally closest traces also have temporal similarity of at least 0.25 and bound the value below so that the two most temporally distant traces have similarity below 0.5. These prior constraints formalize the theoretical assumption that recall is not based on pure rote learning, but involves generalization, and that probing memory produces some level of meaningful signal rather than activating all possible traces. For our T_i values this leads to the prior $\mu_\lambda^g \sim \text{uniform}(1/2, 300)$. The priors on μ_ω^g and μ_τ are naturally set as $\text{uniform}(0, 1)$ and we set $\mu_\beta \sim \text{uniform}(0, 20)$ to allow for both very noisy and very precise thresholds. We set uniform priors on the standard deviations over a large range of plausible values (Gelman, 2006).

When memory is probed, the temporal similarity between a probe and all the traces is calculated using Equation 3.1. The temporal discriminability, overall similarity incorporating semantic information, and probability of trace retrieval then follow from Equations 3.2–3.4, using the individual-specific parameter values. Finally, the application of Equation 3.6 gives θ_{pa} , the probability participant p recalls animal a , so the behavioral data are modeled as

$$y_{pa} \sim \text{Bernoulli}(\theta_{pa}). \quad (3.9)$$

We implemented the model in JAGS (Plummer, 2003), which provides a high-level scripting language for implementing probabilistic models using Markov-chain Monte Carlo sampling methods (Lee & Wagenmakers, 2013). The results are based on 4 chains of 3000 posterior samples collected after 5000 discarded burn-in samples and a thinning factor of 10. We assessed convergence of chains by visual inspection and through the \hat{R} statistic (Brooks & Gelman, 1998).

3.4 Results

3.4.1 Descriptive Adequacy

We evaluated the descriptive adequacy of the model by comparing the model-described probability of recall to the actual rate of recall observed in the data. This is shown in the lower-right panel of Figure 3.4. The gray diagonal line in the figure indicates the model-described probabilities of recall matching the probabilities observed in the data (e.g., among the animal names for which the model gave a 70% probability of recall, 70% were in fact recalled). The data are binned into deciles, and error bars represent the standard error of the mean of the proportions. While there is some small discrepancy, the model generally describes the data well.

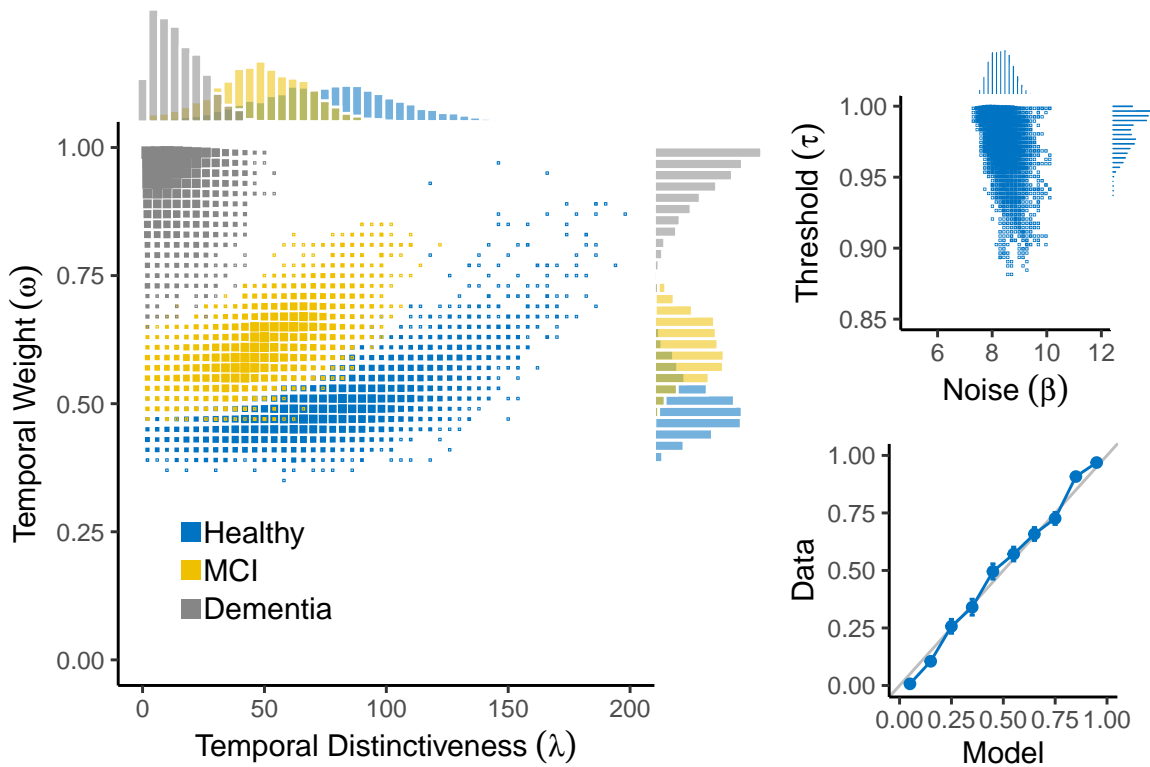


Figure 3.4: Modeling results. The *left panel* shows the joint and marginal posterior distributions for temporal distinctiveness μ_λ and temporal weight μ_ω . These posteriors indicate that as impairment increased, the temporal distinctiveness of memory traces decreased and people relied more on cues of temporal similarity than semantic similarity. The *top-right panel* shows that people used a relatively high threshold μ_τ and a noise μ_β consistent with gradual change from low to high recall probabilities. The *bottom right panel* compares the model-described probabilities of recall to the observed probability of recall, and indicates the model is descriptively adequate.

3.4.2 Modeling Results

The modeling results for μ_β and μ_τ are presented in the upper-right panel of Figure 3.4. The value of the threshold parameter μ_τ is near 1.0 with a mean of 0.98 and a 95% credible interval of [0.93, 1.00]. This high threshold indicates that recall was only possible for traces with high retrieval probabilities. We believe this is a consequence of the combinatorics of probing memory with all 48 traces that repeatedly represent just nine unique items. The noise parameter μ_β had more moderate values, mean 8.36 and 95% CrI [7.64, 9.29], indicating a gradual and somewhat noisy transition from low to high retrieval probabilities.

The most important results involve the μ_λ and μ_ω parameters, since they detail how memory acuity and the use of semantic similarity change between the groups. These inferences are presented in the larger left panel of Figure 3.4. We begin with a discussion of the temporal distinctiveness parameter, μ_λ . The Healthy group had a mean of 76.97, 95% CrI [19.64, 137.7]; the MCI group had a mean of 48.28, 95% CrI [8.82, 84.13]; and the Dementia group had a mean of 14.59, 95% CrI [1.11, 36.92]. The marginal distributions of μ_λ suggest that greater temporal distinctiveness is associated with groups that have less cognitive impairment. The posterior probability that μ_λ for the Healthy group is greater than μ_λ for the MCI group is 0.79, and the posterior probability that μ_λ for the MCI group is greater than μ_λ for the Dementia group is 0.94. While there is some overlap in the marginal distributions of μ_λ , we interpret these results as evidence that higher values of μ_λ are associated with less cognitive impairment, and accordingly, the distributions of μ_λ approach smaller values as cognitive impairment increases, indicating decreased temporal distinctiveness among the memory traces with increasing impairment.

The weight parameter μ_ω indicates how much weight is placed on temporal versus semantic similarity when attempting to recall words. A μ_ω of 0.5 can be interpreted as meaning that approximately equal weight is placed on temporal and semantic cues at retrieval. The

Healthy group had a mean of 0.50, 95% CrI [0.41, 0.63]; the MCI group had a mean of 0.61, 95% CrI [0.51, 0.73]; and the Dementia group had a mean of 0.93, 95% CrI [0.78, 1.00]. As groups become more impaired, the marginal distributions for μ_ω approach larger values. The posterior probability that μ_ω for the MCI group is greater than μ_ω for the Healthy group is 0.96, and the posterior probability that μ_ω for the Dementia group is greater than μ_ω for the MCI group is near 1.0. We interpret these results as evidence that people in the Dementia group rely more on cues of temporal similarity than semantic similarity in the free recall task. In contrast, both the Healthy and MCI groups place relatively more importance on semantic cues and less importance on temporal cues compared to the Dementia group.

Overall, these findings suggest that as impairment increases, the temporal distinctiveness of memory traces decreases. The use of semantic similarities also decreases with impairment, perhaps because of a loss of access to the relevant semantic information (Westfall & Lee, 2021). Healthy controls and people with MCI give significant emphasis to semantic similarity, consistent with previous research stating that in free recall, people tend to recall items in semantically related clusters both within and between categories (Bousfield, 1953; Howard & Kahana, 2002b; Romney et al., 1993). The loss of both memory acuity and access to useful semantics jointly seem to cause the worsening recall performance, specifically for people in the Dementia group.

3.4.3 Alternative Memory Probing via Temporal Context

In this section, we explore a further extension of the SIMPLE model by implementing an alternative memory probing strategy, based on the idea that temporal context plays a central role in encoding and accessing memory (Howard & Kahana, 2002a; Polyn et al., 2009). The key idea of temporal context theories is that time naturally partitions into different contexts, and these contexts are part of the memory representation for information encountered within

them. The triadic comparison task has obvious temporal partitions in the form of the triads. As the visual clustering in Figure 3.2 makes clear, each presentation of the three animal names and the subsequent answer could be thought of as a context.

The standard assumption of the SIMPLE model, which we used in our analyses, is that every trace is probed. We call this the trace-probing model. The alternative approach suggested by context theory is to probe only once per triad. That is, rather than using all 48 traces to probe memory, only one trace per triad-choice cluster is probed. We call this the context-probing model. It is possible that context probing may be more similar to the way people actually reinstate context to aid memory recall.

There are different possibilities for exactly where a triad is probed. We focus on a model that assumes the final presented animal word in the triad is probed. Intuitively, this seems like a plausible center for the context. It is the end of the stimulus presentation, and the moment at which the respondent must begin thinking about their answer.²

Modeling Results

The results of this analysis are summarized in Figure 3.5. These results are based on three chains of 3000 posterior samples collected after 5000 discarded burn-in samples and a thinning factor of 10. The descriptive adequacy of the model was evaluated in the same way as the trace-probing model. The context-probing model over-predicted recall probabilities for low recall probabilities, but otherwise described the data well.

As discussed earlier, the high value of the threshold parameter μ_τ may be due to the probing of memory with traces that are associated with repeated items. In this new model, we used one quarter of the number of probes. As a result, there were somewhat lower values of

²We also briefly explored alternative assumptions, in which each triad was probed at the first or second presented word, or at the time the answer was provided. Those models did not converge as readily as the one we report here.

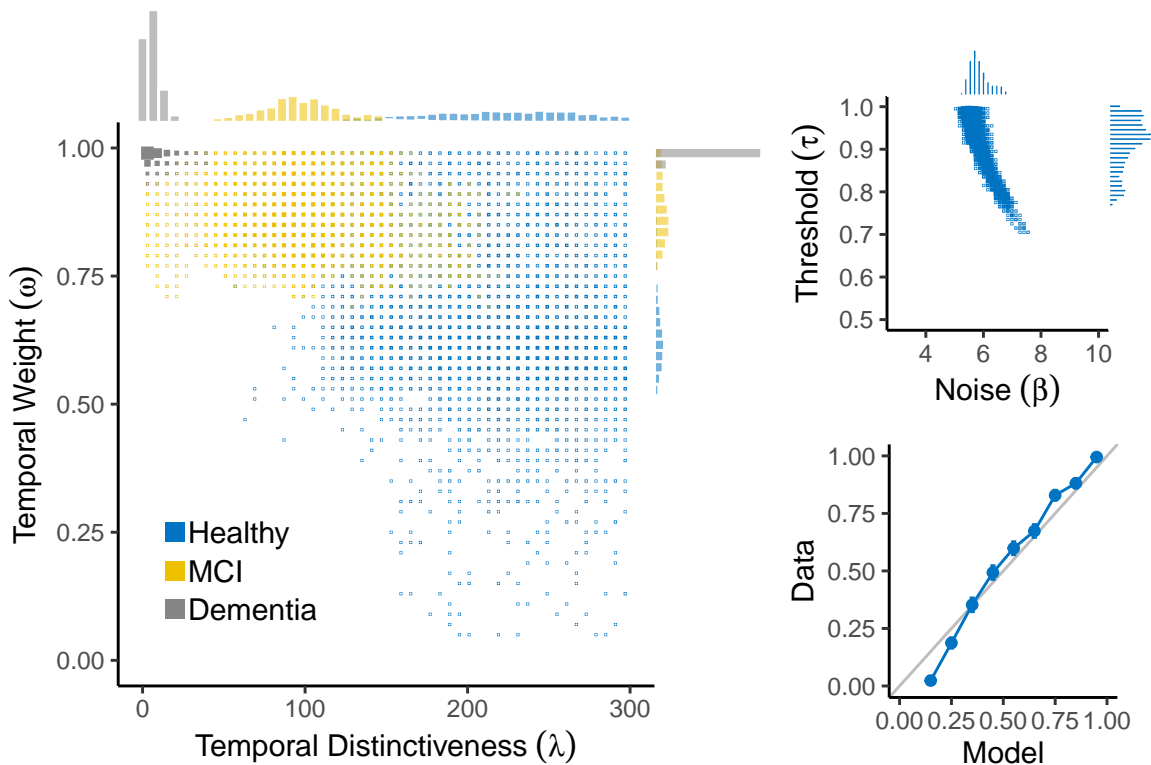


Figure 3.5: Modeling results of the alternative context-probing model. The *left panel* shows the joint and marginal posterior distributions for temporal distinctiveness μ_λ and temporal weight μ_ω . These posteriors indicate that as impairment increases, the temporal distinctiveness of memory traces decreases and people rely more on cues of temporal similarity than semantic similarity. The *top right panel* shows that people again had a relatively high threshold μ_τ and a noise μ_β consistent with a gradual change from low to high recall probabilities, but the joint posterior now shows a trade-off between these two parameters. The *bottom right panel* compares the model-described probabilities of recall to the probability of recall in the data, and indicates the model is descriptively adequate except for over-predicting the recall of infrequently recalled items.

the threshold parameter μ_τ with a mean of 0.91 and 95% CrI [0.78, 1.00], but ultimately retrieval was still only possible for traces with high retrieval probabilities. There were also smaller values of μ_β , with a mean 5.89 and 95% CrI [5.33, 6.78], indicating a slightly noisier transition from low to high retrieval probabilities than the trace-probe model.

For the temporal distinctiveness parameter μ_λ the Healthy group had a mean of 215.40, 95% CrI [124.19, 293.86]; the MCI group had a mean of 96.90, 95% CrI [32.31, 162.96]; and the Dementia group had a mean of 6.11, 95% CrI [0.66, 19.71]. The posterior probability that μ_λ for the Healthy group is greater than μ_λ for the MCI group is nearly 1.0. The posterior probability that μ_λ for the MCI group is greater than μ_λ for the Dementia group is also nearly 1.0. There is clearly very little overlap in these marginal distributions, and they approach smaller values as cognitive impairment increases, indicating decreased temporal distinctiveness among the memory traces with increasing impairment. Compared to the trace-probe model, the distributions for temporal distinctiveness for the Healthy and MCI groups are not only more dispersed, but also include larger values.

For the weight parameter, μ_ω , the Healthy group had a mean of 0.68, 95% CrI [0.42, 0.98]; the MCI group had a mean of 0.88, 95% CrI [0.76, 0.99]; and the Dementia group had a mean of 0.99, 95% CrI [0.97, 1.00]. The posterior probability that μ_ω for the MCI group is greater than μ_ω for the Healthy group is 0.86. The posterior probability that μ_ω for the Dementia group is greater than μ_ω for the MCI group is 0.97. Similar to the observations of the trace-probing model, the marginal distribution of μ_ω for the Dementia group approaches values closer to 1, suggesting that people in the Dementia group rely more on cues of temporal similarity than semantic similarity in the free recall task as compared to the MCI and Healthy groups. However, unlike the full model, the marginal distribution for the MCI group is substantially higher than the Healthy group.

Overall, the results for the context-probe model essentially point to similar conclusions as the trace-probe model. It is clear that people in the Dementia group have very low values of

temporal distinctiveness and high values of temporal weight. However there is much more uncertainty for the MCI group and the Healthy group. It is possible that probing once per triad-choice cluster means that there are multiple ways to account for the behavioral data of these groups. There is also a concern in terms of the descriptive adequacy of the context-probe model, because the recall of infrequently-recalled words is overestimated. Additionally, while the parameters μ_τ and μ_β did not change much in mean values, there is a clear trade-off between these two parameters. Ultimately, we think further development of the context-probe model needs to be motivated by additional independent evidence that people actually probe memory according to the temporal structure provided by the triadic-comparison task.

3.5 Discussion

In this paper we extended the SIMPLE model of memory to account for the free recall of animal names learned incidentally during a triadic comparison task in a real-world clinical data set. We found that the ability to distinguish between temporally similar items – and memory acuity overall – decreased as cognitive impairment increased. Additionally, the extent to which groups used temporal cues relative to semantic cues in memory recall increased with increasing impairment. Jointly, these findings explain the poor recall performance of people in the Dementia group to be due to both lower memory acuity and a reliance on rote learning and temporal cues compared to people in the Healthy and MCI groups.

Our results have both theoretical and clinical implications. The standard SIMPLE model does not incorporate item repetition, but repetition is an essential component of the triadic comparison task, and for the everyday challenge of recalling from past experience. We accounted for the repetition of items on recall by storing each encounter with an item as a separate memory trace. This extension of the model allowed for improved recall for items that were encountered more frequently and therefore were associated with more memory

traces. Our model thus provided an example of how free recall can be predicted when stimuli are encountered in more complicated ways than standard study-test designs.

A decline in episodic memory over and above that associated with healthy aging is one of the first noticeable signs of MCI (Berna et al., 2012), however evidence for differences in semantic memory between healthy controls and patients with MCI is less conclusive. Adlam et al. (2006) found that in a category fluency task, patients with MCI performed similarly to patients with Alzheimer’s disease. Joubert et al. (2010) found a similar pattern of results using a picture-naming task and identifying images of famous people. However, several other studies have found that MCI patients had semantic memory performance similar to healthy controls (Carter et al., 2012; Nakhla et al., 2022; Salmon, 2012) using the Boston Naming Test (Goodglass et al., 1983), a widely used picture-naming task. Our model provides a measure of how semantic memory may influence episodic memory (Howard & Kahana, 2002b). In this case, we do see a difference in the use of semantic memory between healthy controls and patients with MCI. Specifically, people in the MCI group are more likely to use temporal information than semantic information to cue memory recall as compared to the Healthy group. Whether this technique is due to an impaired ability to access semantic information or some other reason, such as the reliance on a non-optimal memory cueing technique, is undetermined for now.

In conclusion, the use of cognitive models like the SIMPLE model of memory to measure latent variables affords an opportunity for a more complete understanding of how memory changes with cognitive impairment, over and above more basic measurements such as recall accuracy. One drawback of SIMPLE is that, while it allows for the prediction of whether an item will be recalled, it does not predict the order in which items are recalled. Future research should focus on the temporal context of recall to try to determine the order in which words are likely to be recalled.

3.6 Publication Note

A previous version of this chapter was published as Westfall, H.A., & Lee, M.D. (2022). A model of free recall for multiple encounters of semantically-related stimuli with an application to understanding cognitive impairment. In J. Culbertson, A. Perfors, H. Rabagliati, & V. Ramenzoni (Eds.), *Proceedings of the 44th Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.

The current version of this chapter has been submitted and is currently under review at the journal *Computational Brain and Behavior*.

CONCLUSION

The research described in this dissertation shows that as cognitive impairment increases, the ability to access semantic memory decreases. In Chapter 1, the odd-one-out choices made by people with Alzheimer's disease were less based on the animal similarities and more random compared to healthy controls. In Chapters 2 and 3, people with MCI and Alzheimer's disease had difficulty using semantic information as a retrieval cue to recall the animal names learned during the odd-one-out triadic comparison task.

Each chapter of the dissertation described a custom-made cognitive model that allowed for the measurement of latent variables not directly observable from behavioral data. The model described in Chapter 1 made it possible to account for non-similarity-based influences on odd-one-out choice behavior, such as response determinism and recency bias. Neither of those parameters could be calculated using only the counts of the odd-one-out choices. The model developed in Chapter 2 provided the insight that the semantic similarity between the animal names and the number of times an animal name was chosen as the odd-one-out were important influences on episodic recall. This model also demonstrated that other item-level traits that are often controlled in free recall tasks such as word frequency, word length, and age of acquisition, were *not* influential on item recall for this particular data set. In Chapter 3, the extended SIMPLE model was able to determine the extent to which different groups relied on semantic versus episodic cues in the free recall task, using the model-based similarities from Chapter 1 and the order of items as they were encountered in the odd-one-

out triadic comparison task. Collectively, these results demonstrate the usefulness of using cognitive models to measure and understand the impact of Alzheimer’s disease on semantic memory.

Nonetheless, results in this dissertation suggest that there are a few promising avenues for detecting impairment in semantic memory from simple behavioral assessments. Based on the findings in Chapter 2, one measure of semantic decline is the similarity of successively recalled words. Specifically, I found that there was a substantial difference between healthy controls and people with MCI in the use of abstract features (i.e., features relating to an animal’s abilities or behavior) as a free recall cue. For a less fine-grained measure of semantic decline (say comparing healthy controls to people with Alzheimer’s dementia) one could examine the recency bias in the triadic comparison task, as described in Chapter 1. In this task, people with Alzheimer’s dementia were much more likely to simply choose the last animal name that they heard as the odd one out.

The particular data set examined in this dissertation is richly informative. However, access to information on recall intrusions and response times would allow for a more complete analysis. The data as it is currently organized codes memory intrusions as errors, but does not record the erroneous word recalled. I would like to see how the types of intrusions are related to the animals that a person actually encountered in the triadic comparison task. For example, if a person was unable to recall the word “chimpanzee” and instead responded “ape”, this might be evidence that a person might be able to recall superordinate categories, but have more difficulty recalling specific animals. There are some theories of semantic decline that propose semantic memory degrades in a “bottom-up” fashion, where Alzheimer’s patients more readily generate category names than specific exemplars (Henry et al., 2004; Martin & Fedio, 1983; Tröster et al., 1989). Information on recall intrusions would allow for theories like this to be directly tested.

Information on response time is not collected in the MCIS due to the manner in which

the task is administered, but access to this information would allow for analyses of global and local search strategies. For example, it would be possible to develop an optimal foraging model of free recall output. Cognitively healthy people tend to generate words in semantically-related categorical clusters. The response times within a cluster tend to be much faster than between clusters (Gruenewald & Lockhead, 1980). This behavior is consistent with a two-stage semantic search process, where people search for clusters within their semantic memory and then produce what is inside of them (Hills et al., 2012). With information on response time, it would be possible to compare the local and global search strategies of the different FAST groups.

The next logical step in this dissertation's line of research would be to develop a model that is able to take the model parameters inferred thus far and determine the expected order of free recall. From the models described in this dissertation, we have the pairwise animal similarity as inferred for healthy controls. For each group, we have a distribution of temporal distinctiveness and a distribution of weight given to semantic versus episodic memory cues. Finally, for each person, we have the order in which items were encountered as part of the triadic comparison task as well as their odd-one-out choices. Together, it should be possible to create a model of the order of free recall. Given the impact Alzheimer's disease seems to have on the use of similarity information in guiding the sequence of recall, and the value of model-based measurement, such a model has the potential to provide a highly diagnostic assessment of semantic memory.

References

- Adlam, A.-L. R., Bozeat, S., Arnold, R., Watson, P., & Hodges, J. R. (2006). Semantic knowledge in mild cognitive impairment and mild Alzheimer's disease. *Cortex*, *42*(5), 675–684. doi: [https://doi.org/10.1016/S0010-9452\(08\)70404-0](https://doi.org/10.1016/S0010-9452(08)70404-0)
- Alzheimer's Association. (2023). *2023 Alzheimer's disease facts and figures*. Retrieved from <https://www.alz.org/media/Documents/alzheimers-facts-and-figures.pdf>
- Au, A., Chan, A. S., & Chiu, H. (2003). Conceptual organization in Alzheimer's dementia. *Journal of Clinical and Experimental Neuropsychology*, *25*(6), 737–750. doi: <https://doi.org/10.1076/jcen.25.6.737.16468>
- Balota, D. A., Burgess, G. C., Cortese, M. J., & Adams, D. R. (2002). The word-frequency mirror effect in young, old, and early-stage Alzheimer's disease: Evidence for two processes in episodic recognition performance. *Journal of Memory and Language*, *46*(1), 199–226. doi: <https://doi.org/10.1006/jmla.2001.2803>
- Balota, D. A., Dolan, P. O., & Duchek, J. M. (2000). Memory changes in healthy older adults. In E. Tulving & F. I. M. Craik (Eds.), *The Oxford handbook of memory* (pp. 395–410). Oxford University Press, Inc.
- Balota, D. A., Yap, M. J., Hutchison, K. A., Cortese, M. J., Kessler, B., Loftis, B., . . . Treiman, R. (2007). The English lexicon project. *Behavior Research Methods*, *39*(3), 445–459. doi: <https://doi.org/10.3758/BF03193014>
- Benton, A. L., Hamsher, K., & Sivan, A. B. (1983). *Multilingual aphasia examination* (3rd ed.). AJA Associates.
- Berna, F., Schönknecht, P., Seidl, U., Toro, P., & Schröder, J. (2012). Episodic autobiographical memory in normal aging and mild cognitive impairment: A population-based study. *Psychiatry Research*, *200*(2-3), 807–812. doi: <https://doi.org/10.1016/j.psychres.2012.03.022>
- Borg, I., & Lingoes, J. (1987). *Multidimensional similarity structure analysis*. New York, NY: Springer-Verlag.
- Bousfield, W. A. (1953). The occurrence of clustering in the recall of randomly arranged associates. *The Journal of General Psychology*, *49*(2), 229–240. doi: <https://doi.org/10.1080/00221309.1953.9710088>

- Brooks, S. P., & Gelman, A. (1998). General methods for monitoring convergence of iterative simulations. *Journal of Computational and Graphical Statistics*, *7*(4), 434–455. doi: <https://doi.org/10.1080/10618600.1998.10474787>
- Brown, G. D. A., Neath, I., & Chater, N. (2007). A temporal ratio model of memory. *Psychological Review*, *114*(3), 539–576. doi: <https://doi.org/10.1037/0033-295X.114.3.539>
- Burton, M. L., & Nerlove, S. B. (1976). Balanced designs for triads tests: Two examples from English. *Social Science Research*, *5*(3), 247–267. doi: [https://doi.org/10.1016/0049-089X\(76\)90002-8](https://doi.org/10.1016/0049-089X(76)90002-8)
- Buschke, H. (1984). Cued recall in amnesia. *Journal of Clinical and Experimental Neuropsychology*, *6*(4), 433–440. doi: <https://doi.org/10.1080/01688638408401233>
- Carter, S. F., Caine, D., Burns, A., Herholz, K., & Lambon Ralph, M. A. (2012). Staging of the cognitive decline in Alzheimer’s disease: Insights from a detailed neuropsychological investigation of mild cognitive impairment and mild Alzheimer’s disease. *International Journal of Geriatric Psychiatry*, *27*(4), 423–432. doi: <https://doi.org/10.1002/gps.2738>
- Chan, A. S., Butters, N., Paulsen, J. S., Salmon, D. P., Swenson, M. R., & Maloney, L. T. (1993). An assessment of the semantic network in patients with Alzheimer’s disease. *Journal of Cognitive Neuroscience*, *5*(2), 254–261. doi: <https://doi.org/10.1162/jocn.1993.5.2.254>
- Chan, A. S., Butters, N., Salmon, D. P., Johnson, S. A., Paulsen, J. S., & Swenson, M. R. (1995). Comparison of the semantic networks in patients with dementia and amnesia. *Neuropsychology*, *9*(2), 177–186. doi: <https://doi.org/10.1037/0894-4105.9.2.177>
- Chan, A. S., Butters, N., Salmon, D. P., & McGuire, K. A. (1993). Dimensionality and clustering in the semantic network of patients with Alzheimer’s disease. *Psychology and Aging*, *8*(3), 411–419. doi: <https://doi.org/10.1037/0882-7974.8.3.411>
- Chan, A. S., Salmon, D. P., & De La Pena, J. (2001). Abnormal semantic network for ‘animals’ but not ‘tools’ in patients with Alzheimer’s disease. *Cortex*, *37*(2), 197–217. doi: [https://doi.org/10.1016/S0010-9452\(08\)70568-9](https://doi.org/10.1016/S0010-9452(08)70568-9)
- Cox, T. F., & Cox, M. A. A. (1994). *Multidimensional scaling*. London, UK: Chapman and Hall.
- De Deyne, S., Navarro, D. J., Perfors, A., Brysbaert, M., & Storms, G. (2019). The “Small World of Words” English word association norms for over 12,000 cue words. *Behavior Research Methods*, *51*(3), 987–1006. doi: <https://doi.org/10.3758/s13428-018-1115-7>
- De Deyne, S., Verheyen, S., Ameel, E., Vanpaemel, W., Dry, M. J., Voorspoels, W., & Storms, G. (2008). Exemplar by feature applicability matrices and other Dutch normative data for semantic concepts. *Behavior Research Methods*, *40*, 1030–1048. doi: <https://doi.org/10.3758/BRM.40.4.1030>

- Delis, D. C., Kramer, J. H., & Kaplan, E. (2000). *California Verbal Learning Test* (2nd ed.). San Antonio, TX: The Psychological Corporation.
- Elvevåg, B., & Storms, G. (2003). Scaling and clustering in the study of semantic disruptions in patients with schizophrenia: A re-evaluation. *Schizophrenia Research*, *63*(3), 237–246. doi: [https://doi.org/10.1016/S0920-9964\(02\)00331-6](https://doi.org/10.1016/S0920-9964(02)00331-6)
- Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models. *Bayesian Analysis*, *1*(3), 515–534. doi: <https://doi.org/10.1214/06-BA117A>
- Goldstone, R. L. (1999). Similarity. In R. A. Wilson & F. C. Keil (Eds.), *The MIT encyclopedia of the cognitive sciences* (pp. 763–765). MIT Press.
- Goodglass, H., Kaplan, E., & Weintraub, S. (1983). *Boston naming test*. Philadelphia, PA: Lea & Febiger.
- Grober, E., & Buschke, H. (1987). Genuine memory deficits in dementia. *Developmental Neuropsychology*, *3*(1), 13–36. doi: <https://doi.org/10.1080/87565648709540361>
- Gruenewald, P. J., & Lockhead, G. R. (1980). The free recall of category examples. *Journal of Experimental Psychology: Human Learning and Memory*, *6*(3), 225–240. doi: <https://doi.org/10.1037/0278-7393.6.3.225>
- Healey, M. K., & Kahana, M. J. (2014). Is memory search governed by universal principles or idiosyncratic strategies? *Journal of Experimental Psychology: General*, *143*(2), 575–596. doi: <https://doi.org/10.1037/a0033715>
- Henry, J. D., Crawford, J. R., & Phillips, L. H. (2004). Verbal fluency performance in dementia of the Alzheimer's type: A meta-analysis. *Neuropsychologia*, *42*(9), 1212–1222. doi: <https://doi.org/10.1016/j.neuropsychologia.2004.02.001>
- Hills, T. T., Jones, M. N., & Todd, P. M. (2012). Optimal foraging in semantic memory. *Psychological Review*, *119*(2), 431–440. doi: <https://doi.org/10.1037/a0027373>
- Howard, M. W., & Kahana, M. J. (1999). Contextual variability and serial position effects in free recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *25*(4), 923–941. doi: <https://doi.org/10.1037/0278-7393.25.4.923>
- Howard, M. W., & Kahana, M. J. (2002a). A distributed representation of temporal context. *Journal of mathematical psychology*, *46*(3), 269–299. doi: <https://doi.org/10.1006/jmps.2001.1388>
- Howard, M. W., & Kahana, M. J. (2002b). When does semantic similarity help episodic retrieval? *Journal of Memory and Language*, *46*(1), 85–98. doi: <https://doi.org/10.1006/jmla.2001.2798>
- Huntley, J. D., & Howard, R. J. (2010). Working memory in early Alzheimer's disease: A neuropsychological review. *International Journal of Geriatric Psychiatry*, *25*(2), 121–132. doi: <https://doi.org/10.1002/gps.2314>

- Huys, Q. J. M., Browning, M., Paulus, M. P., & Frank, M. J. (2021). Advances in the computational understanding of mental illness. *Neuropsychopharmacology*, *46*, 3–19. doi: <https://doi.org/10.1038/s41386-020-0746-4>
- Joubert, S., Brambati, S. M., Ansado, J., Barbeau, E. J., Felician, O., Didic, M., . . . Kergoat, M.-J. (2010). The cognitive and neural expression of semantic memory impairment in mild cognitive impairment and early Alzheimer’s disease. *Neuropsychologia*, *48*(4), 978–988. doi: <https://doi.org/10.1016/j.neuropsychologia.2009.11.019>
- Katz, S., Ford, A. B., Moskowitz, R. W., Jackson, B. A., & Jaffe, M. W. (1963). Studies of illness in the aged: The index of ADL: A standardized measure of biological and psychosocial function. *Journal of the American Medical Association*, *185*(12), 914–919. doi: [10.1001/jama.1963.03060120024016](https://doi.org/10.1001/jama.1963.03060120024016)
- Kavé, G. (2005). Standardization and norms for a Hebrew naming test. *Brain and Language*, *92*(2), 204–211. doi: <https://doi.org/10.1016/j.bandl.2004.06.004>
- Kopelman, M. D., Wilson, B. A., & Baddeley, A. D. (1989). The autobiographical memory interview: A new assessment of autobiographical and personal semantic memory in amnesic patients. *Journal of Clinical and Experimental Neuropsychology*, *11*(5), 724–744. doi: <https://doi.org/10.1080/01688638908400928>
- Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato’s problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, *104*(2), 211–240. doi: <https://doi.org/10.1037/0033-295X.104.2.211>
- Lawton, M. P., & Brody, E. M. (1969). Assessment of older people: Self-maintaining and instrumental activities of daily living. *The Gerontologist*, *9*, 179–186.
- Lee, M. D., Abramyan, M., & Shankle, W. R. (2016). New methods, measures, and models for analyzing memory impairment using triadic comparisons. *Behavior Research Methods*, *48*(4), 1492–1507. doi: <https://doi.org/10.3758/s13428-015-0662-4>
- Lee, M. D., & Pooley, J. P. (2013). Correcting the SIMPLE model of free recall. *Psychological Review*, *120*(1), 293–296. doi: <https://doi.org/10.1037/a0030971>
- Lee, M. D., & Pope, K. J. (2006). Model selection for the rate problem: A comparison of significance testing, Bayesian, and minimum description length statistical inference. *Journal of Mathematical Psychology*, *50*(2), 193–202. doi: <https://doi.org/10.1016/j.jmp.2005.11.010>
- Lee, M. D., & Vanpaemel, W. (2018). Determining informative priors for cognitive models. *Psychonomic Bulletin & Review*, *25*, 114–127. doi: <https://doi.org/10.3758/s13423-017-1238-3>
- Lee, M. D., & Wagenmakers, E.-J. (2013). *Bayesian cognitive modeling: A practical course*. Cambridge, UK: Cambridge University Press.

- Levine, B., Svoboda, E., Hay, J. F., Winocur, G., & Moscovitch, M. (2002). Aging and autobiographical memory: Dissociating episodic from semantic retrieval. *Psychology and Aging, 17*(4), 677–689. doi: <https://doi.org/10.1037/0882-7974.17.4.677>
- Luce, R. D. (1959). *Individual choice behavior: A theoretical analysis*. New York, NY: John Wiley & Sons.
- Lund, K., & Burgess, C. (1996). Producing high-dimensional semantic spaces from lexical co-occurrence. *Behavior Research Methods, Instruments, & Computers, 28*(2), 203–208. doi: <https://doi.org/10.3758/BF03204766>
- Maia, T. V., & Frank, M. J. (2011). From reinforcement learning models to psychiatric and neurological disorders. *Nature Neuroscience, 14*, 154–162. doi: <https://doi.org/10.1038/nn.2723>
- Marshall, P. H., & Werder, P. R. (1972). The effects of the elimination of rehearsal on primacy and recency. *Journal of Verbal Learning and Verbal Behavior, 11*(5), 649–653. doi: [https://doi.org/10.1016/S0022-5371\(72\)80049-5](https://doi.org/10.1016/S0022-5371(72)80049-5)
- Martin, A., & Fedio, P. (1983). Word production and comprehension in Alzheimer’s disease: The breakdown of semantic knowledge. *Brain and Language, 19*(1), 124–141. doi: [https://doi.org/10.1016/0093-934X\(83\)90059-7](https://doi.org/10.1016/0093-934X(83)90059-7)
- McKhann, G. M., Knopman, D. S., Chertkow, H., Hyman, B. T., Jack Jr., C. R., Kawas, C. H., . . . Phelps, C. H. (2011). The diagnosis of dementia due to Alzheimer’s disease: Recommendations from the National Institute on Aging-Alzheimer’s Association workgroups on diagnostic guidelines for Alzheimer’s disease. *Alzheimer’s & Dementia, 7*(3), 263–269. doi: <https://doi.org/10.1016/j.jalz.2011.03.005>
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*. doi: <https://doi.org/10.48550/arXiv.1301.3781>
- Mortamais, M., Ash, J. A., Harrison, J., Kaye, J., Kramer, J., Randolph, C., . . . Ritchie, K. (2017). Detecting cognitive changes in preclinical Alzheimer’s disease: A review of its feasibility. *Alzheimer’s & Dementia, 13*(4), 468–492. doi: <https://doi.org/10.1016/j.jalz.2016.06.2365>
- Murdock, B. B. (1962). The serial position effect of free recall. *Journal of Experimental Psychology, 64*(5), 482–488. doi: <https://doi.org/10.1037/h0045106>
- Nakhla, M. Z., Banuelos, D., Pagán, C., Gavarrete Olvera, A., & Razani, J. (2022). Differences between episodic and semantic memory in predicting observation-based activities of daily living in mild cognitive impairment and Alzheimer’s disease. *Applied Neuropsychology: Adult, 29*(6), 1499–1510. doi: <https://doi.org/10.1080/23279095.2021.1893172>
- Navarro, D. J., & Griffiths, T. L. (2008). Latent features in similarity judgments: A nonparametric Bayesian approach. *Neural computation, 20*(11), 2597–2628. doi: <https://doi.org/10.1162/neco.2008.04-07-504>

- Navarro, D. J., & Lee, M. D. (2004). Common and distinctive features in stimulus similarity: A modified version of the contrast model. *Psychonomic Bulletin & Review*, *11*(6), 961–974. doi: <https://doi.org/10.3758/BF03196728>
- Nebes, R. D. (1989). Semantic memory in Alzheimer's disease. *Psychological Bulletin*, *106*(3), 377–394. doi: <https://doi.org/10.1037/0033-2909.106.3.377>
- Nebes, R. D., & Brady, C. B. (1990). Preserved organization of semantic attributes in Alzheimer's disease. *Psychology and Aging*, *5*(4), 574–579. doi: <https://doi.org/10.1037/0882-7974.5.4.574>
- Nelson, D. L., McEvoy, C. L., & Schreiber, T. A. (2004). The University of South Florida free association, rhyme, and word fragment norms. *Behavior Research Methods, Instruments, & Computers*, *36*(3), 402–407. doi: <https://doi.org/10.3758/BF03195588>
- Newcombe, F. (1969). *Missile wounds of the brain: A study of psychological deficits*. Oxford University Press.
- Peterson, J. C., Abbott, J. T., & Griffiths, T. L. (2018). Evaluating (and improving) the correspondence between deep neural networks and human representations. *Cognitive Science*, *42*(8), 2648–2669. doi: <https://doi.org/10.1111/cogs.12670>
- Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In K. Hornik, F. Leisch, & A. Zeileis (Eds.), *Proceedings of the 3rd International Workshop on Distributed Statistical Computing* (Vol. 124, pp. 1–10). Vienna, Austria.
- Plummer, M. (2016). rjags: Bayesian graphical models using MCMC [Computer software manual]. Retrieved from <https://CRAN.R-project.org/package=rjags> (R package version 4-6)
- Polyn, S. M., Norman, K. A., & Kahana, M. J. (2009). A context maintenance and retrieval model of organizational processes in free recall. *Psychological Review*, *116*(1), 129–156. doi: <https://doi.org/10.1037/a0014420>
- R Core Team. (2018). R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria. Retrieved from <https://www.R-project.org/>
- Razani, J., Chan, A., Nordin, S., & Murphy, C. (2010). Semantic networks for odors and colors in Alzheimer's disease. *Neuropsychology*, *24*(3), 291–299. doi: <https://doi.org/10.1037/a0018269>
- Reisberg, B. (1988). Functional assessment staging (FAST). *Psychopharmacology Bulletin*, *24*(4), 653–659.
- Ribeiro, F., Guerreiro, M., & De Mendonça, A. (2007). Verbal learning and memory deficits in mild cognitive impairment. *Journal of Clinical and Experimental Neuropsychology*, *29*(2), 187–197. doi: <https://doi.org/10.1080/13803390600629775>

- Romney, A. K., Brewer, D. D., & Batchelder, W. H. (1993). Predicting clustering from semantic structure. *Psychological Science*, *4*(1), 28–34. doi: <https://doi.org/10.1111/j.1467-9280.1993.tb00552.x>
- Romney, A. K., & Moore, C. C. (1998). Toward a theory of culture as shared cognitive structures. *Ethos*, *26*(3), 314–337. doi: <https://doi.org/10.1525/eth.1998.26.3.314>
- Salmon, D. P. (2012). Loss of semantic knowledge in mild cognitive impairment. *American Journal of Psychiatry*, *169*(12), 1226–1229. doi: <https://doi.org/10.1176/appi.ajp.2012.12101262>
- Shankle, W. R., Mangrola, T., Chan, T., & Hara, J. (2009). Development and validation of the Memory Performance Index: Reducing measurement error in recall tests. *Alzheimer's & Dementia*, *5*(4), 295–306. doi: <https://doi.org/10.1016/j.jalz.2008.11.001>
- Shepard, R. N. (1980). Multidimensional scaling, tree-fitting, and clustering. *Science*, *210*(4468), 390–398. doi: <https://doi.org/10.1126/science.210.4468.390>
- Shepard, R. N., & Arabie, P. (1979). Additive clustering: Representation of similarities as combinations of discrete overlapping properties. *Psychological Review*, *86*(2), 87–123. doi: <https://doi.org/10.1037/0033-295X.86.2.87>
- Storms, G., Dirikx, T., Saerens, J., Verstraeten, S., & De Deyn, P. P. (2003). On the use of scaling and clustering in the study of semantic deficits. *Neuropsychology*, *17*(2), 289–301. doi: <https://doi.org/10.1037/0894-4105.17.2.289>
- Toppino, T. C., & Gerbier, E. (2014). About practice: Repetition, spacing, and abstraction. In B. H. Ross (Ed.), (Vol. 60, p. 113-189). Academic Press. doi: <https://doi.org/10.1016/B978-0-12-800090-8.00004-4>
- Trenkle, D. L., Shankle, W. R., & Azen, S. P. (2007). Detecting cognitive impairment in primary care: Performance assessment of three screening instruments. *Journal of Alzheimer's Disease*, *11*(3), 323–335. doi: <https://doi.org/10.3233/jad-2007-11309>
- Tröster, A. I., Salmon, D. P., McCullough, D., & Butters, N. (1989). A comparison of the category fluency deficits associated with Alzheimer's and Huntington's disease. *Brain and Language*, *37*(3), 500–513. doi: [https://doi.org/10.1016/0093-934X\(89\)90032-1](https://doi.org/10.1016/0093-934X(89)90032-1)
- Tversky, A. (1977). Features of similarity. *Psychological Review*, *84*(4), 327–352. doi: <https://doi.org/10.1037/0033-295X.84.4.327>
- Voorspoels, W., Storms, G., Longenecker, J., Verheyen, S., Weinberger, D. R., & Elvevåg, B. (2014). Deriving semantic structure from category fluency: Clustering techniques and their pitfalls. *Cortex*, *55*, 130–147. doi: <https://doi.org/10.1016/j.cortex.2013.09.006>
- Welsh, K. A., Butters, N., Mohs, R. C., Beekly, D., Edland, S., Fillenbaum, G., & Heyman, A. (1994). The consortium to establish a registry for Alzheimer's disease (CERAD). Part V. A normative study of the neuropsychological battery. *Neurology*, *44*(4), 609–614. doi: <https://doi.org/10.1212/WNL.44.4.609>

- Westfall, H. A., & Lee, M. D. (2021). A model-based analysis of the impairment of semantic memory. *Psychonomic Bulletin & Review*, *28*, 1484–1494. doi: <https://doi.org/10.3758/s13423-020-01875-9>
- Westfall, H. A., & Lee, M. D. (2022). A model of free recall for multiple encounters of semantically related stimuli with an application to understanding cognitive impairment. In J. Culbertson, A. Perfors, H. Rabagliati, & V. Ramenzoni (Eds.), *Proceedings of the annual meeting of the cognitive science society* (Vol. 44).
- Wetzels, R., Grasman, R. P. P. P., & Wagenmakers, E.-J. (2010). An encompassing prior generalization of the Savage–Dickey density ratio. *Computational Statistics & Data Analysis*, *54*(9), 2094–2102. doi: <https://doi.org/10.1016/j.csda.2010.03.016>
- Whatmough, C., & Chertkow, H. (2002). Category-specific recognition impairments in Alzheimer’s disease. In E. M. E. Forde & G. W. Humphreys (Eds.), *Category specificity in brain and mind* (pp. 181–210). Psychology Press.
- White, A., Voorspoels, W., Storms, G., & Verheyen, S. (2014). Problems of reliability and validity with similarity derived from category fluency. *Psychiatry Research*, *220*(3), 1125–1130. doi: <https://doi.org/10.1016/j.psychres.2014.10.001>
- Wilson, R. S., Bacon, L. D., Fox, J. H., Kramer, R. L., & Kaszniak, A. W. (1983). Word frequency effect and recognition memory in dementia of the Alzheimer type. *Journal of Clinical and Experimental Neuropsychology*, *5*(2), 97–104. doi: <https://doi.org/10.1080/01688638308401157>
- World Health Organization. (2020, September 21). *Dementia*. Retrieved from <https://www.who.int/en/news-room/fact-sheets/detail/dementia>
- Yiannopoulou, K. G., & Papageorgiou, S. G. (2020). Current and future treatments in Alzheimer disease: An update. *Journal of Central Nervous System Disease*, *12*, 1–12. doi: <https://doi.org/10.1177/1179573520907397>
- Zeigenfuse, M. D., & Lee, M. D. (2010). Finding the features that represent stimuli. *Acta Psychologica*, *133*(3), 283–295. doi: <https://doi.org/10.1016/j.actpsy.2009.07.014>

Appendix A

Model-Based Inference of Pairwise Similarities

In this appendix we describe the method we used to infer the pairwise semantic similarity between the animal names. Standard methods for estimating similarities either measure them directly on Likert scales or calculate them directly from observed choice behavior (Borg & Lingoes, 1987; Cox & Cox, 1994). For odd-one-out choices in a triadic comparison task, the natural data-driven way to estimate the similarity between stimuli a and b is the proportion of times they are *not* chosen in all triads in which they appear (Westfall & Lee, 2022).

A model-based approach to inferring similarities can potentially improve measurement by taking account of the details of the cognitive processes that produce the choice behavior. We used the model of triadic comparison developed by Westfall & Lee (2021). The model assumes that the similarity, s_{ab} , of animal a and animal b is equal to the sum of the weights

of their shared features (Shepard & Arabie, 1979; Tversky, 1977)

$$s_{ab} = \sum_k w_k f_{ak} f_{bk}, \quad (\text{A.1})$$

where f_{ak} is a binary indicator variable that determines whether animal a has feature k , and w_k represents the salience or weight of feature k . The Leuven concept database provides possible features (De Deyne et al., 2008) and a truncated Gaussian prior is placed on the feature weights,

$$w_k \sim \text{Gaussian}(0, 1) \text{ T}(0,). \quad (\text{A.2})$$

The probability π_a of choosing animal a as the odd one out in a triad of animals a , b , and c is determined by their relative similarities, using an exponentiated form of the Luce (1959) choice rule with a response determinism parameter γ . Additionally, the decision model include a recency bias β , corresponding to the strategy of choosing the last animal name presented in the triad. Overall, the choice probabilities are defined as

$$\begin{aligned} \pi_a &= \frac{1 - \beta}{2} \left(\frac{s_{bc}^\gamma}{s_{ac}^\gamma + s_{bc}^\gamma + s_{ab}^\gamma} \right) \\ \pi_b &= \frac{1 - \beta}{2} \left(\frac{s_{ac}^\gamma}{s_{ac}^\gamma + s_{bc}^\gamma + s_{ab}^\gamma} \right) \\ \pi_c &= \beta \left(\frac{s_{ab}^\gamma}{s_{ac}^\gamma + s_{bc}^\gamma + s_{ab}^\gamma} \right). \end{aligned} \quad (\text{A.3})$$

The prior $\gamma \sim \text{gamma}(2, 1)$ is assumed for response determinism, and $\beta \sim \text{uniform}(1/3, 1)$ is assumed for for recency bias. The observed choice of participant i on trial t is modeled as

$$y_{it} = \text{categorical}(\pi_a, \pi_b, \pi_c). \quad (\text{A.4})$$

Given this generative model, inferring the w_k weights provides a measure of the pairwise similarities between animal using Equation A.1. We implemented the model in JAGS to infer the weights, and applied it to the odd-one-out triadic choices of the Healthy controls. The resulting similarity measures were unit-standardized and self-similarities were set to 1.0.

The model-based approach controls for the influence of response determinism and recency bias — which are not related to stimulus similarity — on observed choice behavior. Statistically, in making inferences about w_k , the γ and β parameters are treated as nuisance parameters to be marginalized out. Psychologically, the model-based approach serves to decontaminate the inferred similarities from the impact of response determinism and recency bias.